

MFCA'08



Mathematical Foundations of Computational Anatomy

**Geometrical and Statistical Methods
for Biological Shape Variability Modeling**

September 6th, 2008, Kimmel Center, New York, USA.

<http://www.inria.fr/sophia/asclepios/events/MFCA08>

Editors:

Xavier Pennec (Asclepios, INRIA Sophia-Antipolis, France)
Sarang Joshi (SCI, University of Utah, USA)



Preface

The goal of computational anatomy is to analyze and to statistically model the anatomy of organs in different subjects. Computational anatomic methods are generally based on the extraction of anatomical features or manifolds which are then statistically analyzed, often through a non-linear registration. There are nowadays a growing number of methods that can faithfully deal with the underlying biomechanical behavior of intra-subject deformations. However, it is more difficult to relate the anatomies of different subjects. In the absence of any justified physical model, diffeomorphisms provide a general mathematical framework that enforces topological consistency. Working with such infinite dimensional space raises some deep computational and mathematical problems, in particular for doing statistics. Likewise, modeling the variability of surfaces leads to rely on shape spaces that are much more complex than for curves. To cope with these, different methodological and computational frameworks have been proposed (e.g. smooth left-invariant metrics, focus on well-behaved subspaces of diffeomorphisms, modeling surfaces using courants, etc.)

The goal of the **Mathematical Foundations of Computational Anatomy (MFCA)** workshop is to foster the interactions between the mathematical community around shapes and the MICCAI community around computational anatomy applications. It targets more particularly researchers investigating the combination of statistical and geometrical aspects in the modeling of the variability of biological shapes. The workshop aims at being a forum for the exchange of the theoretical ideas and a source of inspiration for new methodological developments in computational anatomy. A special emphasis is put on theoretical developments, applications and results being welcomed as illustrations. Following the very successful first edition of this workshop in 2006 (see <http://www.inria.fr/sophia/asclepios/events/MFCA06/>), the second edition was held in New-York on September 6, in conjunction with MICCAI 2008.

Contributions were solicited in Riemannian and group theoretical methods, Geometric measurements of the anatomy, Advanced statistics on deformations and shapes, Metrics for computational anatomy, Statistics of surfaces. 34 submissions were received, among which 9 were accepted to MICCAI and had to be withdrawn from the workshop. Each of the remaining 25 papers was reviewed by three members of the program committee. To guarantee a high level program, 16 papers only were selected.

August 2008

Xavier Pennec
General Chair
MFCA'08

Organization

Workshop Chairs

Xavier Pennec	(INRIA Sophia-Antipolis, France)
Sarang Joshi	(SCI, University of Utah, USA)

Program Committee

Rachid Deriche	(INRIA, France)
Ian L. Dryden	(University of Nottingham, UK)
Tom Fletcher	(University of Utah, USA)
James Gee	(Univ. of Pennsylvania, USA)
Guido Gerig	(University of Utah, USA)
Polina Golland	(CSAIL, MIT, USA)
Stephen Marsland	(Massey University, New-Zeeland)
Michael I. Miller	(John Hopkins University, USA)
Mads Nielsen	(IT University of Copenhagen, Denmark)
Salvador Olmos	(University of Saragossa, Spain)
Bruno Pelletier	(University Montpellier, France)
Jerry Prince	(Johns Hopkins University, USA)
Anand Rangarajan	(University of Florida, USA)
Daniel Rueckert	(Imperial College London, UK)
Guillermo Sapiro	(University of Minnesota, USA)
Martin Styner	(UNC Chapel Hill, USA)
Anuj Srivastava	(Florida State University, USA)
Paul Thompson	(University of California Los-Angeles, USA)
Alain Trouvé	(ENS-Cachan, France)
Carole Twining	(University of Manchester, UK)
William M. Wells III	(CSAIL, MIT, and B&W Hospital, Boston, USA)

Table of Contents

Registration

Unbiased Volumetric Registration via Nonlinear Elastic Regularization ..	1
<i>Igor Yanovsky, Carole Le Guyader, Alex Leow, Arthur Toga, Paul Thompson, and Luminita Vese</i>	
A new algorithm for the computation of the group logarithm of diffeomorphisms	13
<i>Matias Bossa and Salvador Olmos</i>	
Comparing algorithms for diffeomorphic registration: Stationary LDDMM and Diffeomorphic Demons	24
<i>Monica Hernandez, Salvador Olmos, and Xavier Pennec</i>	

Morphometry

Brain Mapping with the Ricci Flow Conformal Parameterization and Multivariate Statistics on Deformation Tensors	36
<i>Yalin Wang, Xiaotian Yin, Jie Zhang, Xianfeng Gu, Tony F. Chan, Paul M. Thompson, and Shing-Tung Yau</i>	
Multi-Atlas Tensor-Based Morphometry and its Application to a Genetic Study of 92 Twins	48
<i>Natasha Leporé, Caroline Brun, Yi-Yu Chou, Agatha D. Lee, Marina Barysheva, Greig I. de Zubicaray, Matthew Meredith, Katie L. McMahon, Margaret J. Wright, Arthur W. Toga, and Paul M. Thompson</i>	
Shape Registration with Spherical Cross Correlation	56
<i>Boris Gutman, Yalin Wang, Tony Chan, Paul M. Thompson, and Arthur W. Toga</i>	

Building Atlases

A Forward Model to Build Unbiased Atlases from Curves and Surfaces ..	68
<i>Stanley Durrleman, Xavier Pennec, Alain Trounevé, and Nicholas Ayache</i>	

IV

MAP Estimation of Statistical Deformable Templates Via Nonlinear Mixed Effects Models : Deterministic and Stochastic Approaches	80
<i>Stéphanie Allasonnière, Estelle Kuhn, and Alain Trouvé</i>	
Semiparametric estimation of rigid transformations on compact Lie groups	92
<i>Jérémie Bigot, Jean-Michel Loubes and Myriam Vimond</i>	

Shapes and Surfaces

Geodesic Shape Spaces of Surfaces of Genus Zero	105
<i>Xiuwen Liu, Washington Mio, Yonggang Shi, and Ivo Dinov</i>	
Characterization of Anatomical Shape Based on Random Walk Hitting Times	117
<i>Grace Vesom, Nathan D. Cahill, Lena Gorelick, and J. Alison Noble</i>	
Tiling Manifolds with Orthonormal Basis	128
<i>Moo K. Chung, Anqi Qiu, Brendon M. Nacewicz, Seth Pollak, Richard J. Davidson</i>	

Statistics on Manifolds and Shapes

Diffusion Tensor Imaging and Deconvolution on Spaces of Positive Definite Symmetric Matrices	140
<i>Peter T. Kim and Donald St. P. Richards</i>	
Tubular Surface Evolution for Segmentation of the Cingulum Bundle From DW-MRI	150
<i>Vandana Mohan, Ganesh Sundaramoorthi, John Melonakos, Marc Niethammer, Marek Kubicki, and Allen Tannenbaum</i>	
Modeling the Remaining Flexibility of Partially Fixed Statistical Shape Models	160
<i>Thomas Albrecht, Reinhard Knothe, and Thomas Vetter</i>	
A Hypothesis Testing Framework for High-Dimensional Shape Models . . .	170
<i>Joshua Cates, P. Thomas Fletcher, Ross Whitaker</i>	

Author Index	182
------------------------	-----

Unbiased Volumetric Registration via Nonlinear Elastic Regularization

Igor Yanovsky¹, Carole Le Guyader², Alex Leow³, Arthur Toga³,
Paul Thompson³, and Luminita Vese¹

¹ Department of Mathematics, University of California, Los Angeles, USA,

² Institute of Mathematical Research of Rennes, France,

³ Laboratory of Neuro Imaging, UCLA School of Medicine, USA. *

Abstract. In this paper, we propose a new large-deformation nonlinear image registration model in three dimensions, based on nonlinear elastic regularization and unbiased registration. Both the nonlinear elastic and the unbiased functionals are simplified introducing, in the modeling, a second unknown that mimics the Jacobian matrix of the displacement vector field, reducing the minimization to involve linear differential equations. In contrast to recently proposed unbiased fluid registration method, the new model is written in a unified variational form and is minimized using gradient descent on the corresponding Euler-Lagrange equations. As a result, the new unbiased nonlinear elasticity model is computationally more efficient and easier to implement than the unbiased fluid registration. The model was tested using three-dimensional serial MRI images and shown to have some advantages for computational neuroimaging.

1 Introduction

Given two images, the source and target, the goal of image registration is to find an optimal diffeomorphic spatial transformation such that the deformed source image is aligned with the target image. In the case of non-parametric registration methods (the class of methods we are interested in), the problem can be phrased as a functional minimization problem whose unknown is the displacement vector field \mathbf{u} . Usually, the devised functional consists of a distance measure (intensity-based, correlation-based, mutual-information based [1] or metric-structure-comparison based [2]) and a regularizer that guarantees smoothness of the displacement vector field. Several regularizers have been investigated (see Part II of [1] for a review). Generally, physical arguments motivate the selection of the regularizer. Among those currently used is the linear elasticity smoother first introduced by Broit [3]. The objects to be registered are considered to be observations of the same elastic body at two different times, before and after being subjected to a deformation as mentioned in [1]. The smoother, in this case, is the linearized elastic potential of the displacement vector field. However, this model is unsuitable for problems involving large-magnitude deformations.

* This work was funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 RR021813 entitled Center for Computational Biology (CCB).

In [4], Christensen *et al.* proposed a viscous fluid model to overcome this issue. The deforming image is considered to be embedded in viscous fluid whose motion is governed by Navier-Stokes equations for conservation of momentum:

$$\mu \Delta \mathbf{v}(\mathbf{x}, t) + (\nu + \mu) \nabla (\nabla \cdot \mathbf{v}(\mathbf{x}, t)) = \mathbf{f}(\mathbf{x}, \mathbf{u}(\mathbf{x}, t)), \quad (1)$$

$$\mathbf{v}(\mathbf{x}, t) = \mathbf{u}_t(\mathbf{x}, t) + \nabla \mathbf{u}(\mathbf{x}, t) \cdot \mathbf{v}(\mathbf{x}, t). \quad (2)$$

Here, equation (2), defining material derivative of \mathbf{u} , nonlinearly relates the velocity and displacement vector fields.

One drawback of this method is the computational cost. Numerically, the image-derived force field $\mathbf{f}(\mathbf{x}, \mathbf{u}(\mathbf{x}, t))$ is first computed at time t . Fixing the force field \mathbf{f} , linear equation (1) is solved for $\mathbf{v}(\mathbf{x}, t)$ numerically using the successive over-relaxation (SOR) scheme. Then, an explicit Euler scheme is used to advance \mathbf{u} in time. Recent works [5–7] applied Riemannian nonlinear elasticity priors to deformation velocity fields. These alternating frameworks, however, are time-consuming, which motivates the search for faster implementations (see for instance [8] or [9] in which the instantaneous velocity \mathbf{v} is obtained by convolving \mathbf{f} with a Gaussian kernel).

In this paper, which is inspired from related works on segmentation [10] and on two-dimensional registration [11], we propose an alternative approach to fluid registration. The new model is derived from a variational problem which is not in the form of a two-step algorithm and which can also produce large-magnitude deformations. For that purpose, a nonlinear elasticity smoother is introduced in three dimensions. As will be seen later, the computation of the Euler-Lagrange equations in this case is cumbersome. We circumvent this issue by introducing a second unknown, a matrix variable V , which approximates the Jacobian matrix of \mathbf{u} . The nonlinear elastic regularizer is now applied to V , removing the nonlinearity in the derivatives of the unknown \mathbf{u} in the Euler-Lagrange equations. The Euler-Lagrange equations are straightforwardly derived and a gradient descent method is used.

Also, allowing large deformations to occur may yield non-diffeomorphic deformation mappings. In [4], Christensen *et al.* proposed a regridding technique that re-samples the deforming image and re-initializes the process once the value of the deformation Jacobian drops below a certain threshold. In [12], Haber and Modersitzki introduced an elastic registration model subjected to volume-preserving constraints. To ensure that the transformation $\mathbf{g}(\mathbf{x}) = \mathbf{x} - \mathbf{u}(\mathbf{x})$ is volume-preserving (that is, for any domain Ω , $\int_{\Omega} d\mathbf{x} = \int_{\mathbf{g}(\Omega)} d\mathbf{x}$), they proposed the following pointwise constraint: $\det(\mathcal{I} - D\mathbf{u}(\mathbf{x})) - 1 = 0$. Pursuing in the same direction in [13], the authors introduced a minimization problem under inequality constraints on the Jacobian.

Here we use an information-theoretic approach previously introduced in [14]. In [14], the authors considered a smooth deformation \mathbf{g} that maps domain Ω bijectively onto itself. Consequently, \mathbf{g} and \mathbf{g}^{-1} are bijective and globally volume-preserving. Probability density functions can thus be associated with the deformation \mathbf{g} and its inverse \mathbf{g}^{-1} . The authors then proposed to quantify the magnitude of the deformation by means of the symmetric Kullback-Leibler distance between the probability density functions associated with the deformation and the identity mapping. This distance, when rewritten using skew-symmetry properties, is viewed as a cost function and is combined with the viscous fluid model for registration, which leads to an unbiased fluid registration

model. Unlike the unbiased fluid registration model, the unbiased nonlinear elasticity method, introduced here, allows the functional to be written “in closed form”. The new model also does not require expensive Navier-Stokes solver (or its approximation) at each step as previously mentioned.

2 Method

Let Ω be an open and bounded domain in \mathbb{R}^3 . Without loss of generality, we assume that the volume of Ω is 1, i.e. $|\Omega| = 1$. Let $I_1, I_2 : \Omega \rightarrow \mathbb{R}$ be the two images to be registered. We seek the transformation $\mathbf{g} : \Omega \rightarrow \Omega$ that maps the source image I_2 into correspondence with the target image I_1 . In this paper, we will restrict this mapping to be differentiable, one-to-one, and onto. We denote the Jacobian matrix of a deformation \mathbf{g} to be $D\mathbf{g}$, with Jacobian denoted by $|D\mathbf{g}(\mathbf{x})| = \det(D\mathbf{g}(\mathbf{x}))$ (thus we will use the notation $|V| := \det(V)$ for any 3×3 matrix V). The displacement field $\mathbf{u}(\mathbf{x})$ from the position \mathbf{x} in the deformed image $I_2 \circ \mathbf{g}(\mathbf{x})$ back to $I_2(\mathbf{x})$ is defined in terms of the deformation $\mathbf{g}(\mathbf{x})$ by the expression $\mathbf{g}(\mathbf{x}) = \mathbf{x} - \mathbf{u}(\mathbf{x})$ at every point $\mathbf{x} \in \Omega$. Thus, we consider the problems of finding \mathbf{g} and \mathbf{u} as equivalent.

In general, nonlinear image registration models may be formulated in a variational framework. The minimization problems often define the energy functional E as a linear combination of image matching term F and the regularizing term R : $\inf_{\mathbf{u}} \{E(\mathbf{u}) = F(\mathbf{u}) + \lambda_0 R(\mathbf{u})\}$. Here, $\lambda_0 > 0$ is a weighting parameter.

2.1 Registration metrics

In this paper, the matching functional F takes the form of the L^2 norm (the sum of squared intensity differences), $F = F_{L^2}$, and the mutual information, $F = F_{MI}$.

L^2 -norm: The L^2 -norm matching functional is suitable when the images have been acquired through similar sensors (with additive Gaussian noise) and thus are expected to present the same intensity range and distribution. The L^2 distance between the deformed image $I_2 \circ \mathbf{g}(\mathbf{x}) = I_2(\mathbf{x} - \mathbf{u}(\mathbf{x}))$ and target image $I_1(\mathbf{x})$ is defined as

$$F_{L^2}(\mathbf{u}) = \frac{1}{2} \int_{\Omega} (I_2(\mathbf{x} - \mathbf{u}(\mathbf{x})) - I_1(\mathbf{x}))^2 d\mathbf{x}. \quad (3)$$

Mutual Information: Mutual information can be used to align images of different modalities, without requiring knowledge of the relationship of the two registered images [15, 16]. Here, the intensity distributions estimated from $I_1(\mathbf{x})$ and $I_2(\mathbf{x} - \mathbf{u}(\mathbf{x}))$ are denoted by p^{I_1} and $p_{\mathbf{u}}^{I_2}$, respectively, and an estimate of their joint intensity distribution by $p_{\mathbf{u}}^{I_1, I_2}$. We let $i_1 = I_1(\mathbf{x})$, $i_2 = I_2(\mathbf{x} - \mathbf{u}(\mathbf{x}))$ denote intensity values at point $\mathbf{x} \in \Omega$. Given the displacement field \mathbf{u} , the mutual information computed from I_1 and I_2 is provided by $MI_{\mathbf{u}}^{I_1, I_2} = \int_{\mathbb{R}^2} p_{\mathbf{u}}^{I_1, I_2}(i_1, i_2) \log[p_{\mathbf{u}}^{I_1, I_2}(i_1, i_2)/(p^{I_1}(i_1)p_{\mathbf{u}}^{I_2}(i_2))] di_1 di_2$. We seek to maximize the mutual information between $I_2(\mathbf{x} - \mathbf{u}(\mathbf{x}))$ and $I_1(\mathbf{x})$, or equivalently, minimize the negative of $MI_{\mathbf{u}}^{I_1, I_2}$:

$$F_{MI}(I_1, I_2, \mathbf{u}) = -MI_{\mathbf{u}}^{I_1, I_2}. \quad (4)$$

2.2 Nonlinear Elastic Regularization

The theory of elasticity is based on the notion of strain. Strain is defined as the amount of deformation an object experiences compared to its original size and shape. In three spatial dimensions, the strain tensor, $\mathcal{E} = [\varepsilon_{ij}] \in \mathbb{R}^{3 \times 3}$, $1 \leq i, j \leq 3$, is a symmetric tensor used to quantify the strain of an object undergoing a deformation. The nonlinear strain is defined as $\varepsilon_{ij}(\mathbf{u}) = \frac{1}{2}(\partial_j u_i + \partial_i u_j + \sum_{k=1}^3 \partial_i u_k \partial_j u_k)$, with the nonlinear strain tensor matrix given by

$$\mathcal{E}(\mathbf{u}) = \frac{1}{2}(D\mathbf{u}^t + D\mathbf{u} + D\mathbf{u}^t D\mathbf{u}). \quad (5)$$

Stored energy (Saint Venant-Kirchhoff material) is defined as

$$W(\mathcal{E}) = \frac{\nu}{2}(\text{trace}(\mathcal{E}))^2 + \mu \text{trace}(\mathcal{E}^2),$$

where ν and μ are Lamé elastic material constants. The regularization for nonlinear elasticity becomes

$$R_E(\mathbf{u}) = \int_{\Omega} W(\mathcal{E}(\mathbf{u})) d\mathbf{x}.$$

The regularization term $R_E(\mathbf{u})$ can be minimized with respect to \mathbf{u} . However, since the regularization term is written in terms of partial derivatives of components of \mathbf{u} , the Euler-Lagrange equations become complicated and are computationally expensive to minimize. Instead, following earlier theoretical work [17], we minimize an approximate functional by introducing the matrix variable

$$V \approx D\mathbf{u} \quad (6)$$

and thus consider a new form of nonlinear elasticity regularization functional

$$R_E(\mathbf{u}, V) = \int_{\Omega} W(\widehat{V}) d\mathbf{x} + \frac{\beta}{2} \int_{\Omega} \|V - D\mathbf{u}\|_F^2 d\mathbf{x}, \quad (7)$$

where $\widehat{V} = \frac{1}{2}(V^t + V + V^t V)$, β is a positive constant, and $\|\cdot\|_F$ denotes the Frobenius norm. For β large enough, $R_E(\mathbf{u})$ is well approximated by $R_E(\mathbf{u}, V)$. In the limit as $\beta \rightarrow +\infty$, we obtain $D\mathbf{u} \approx V$ in the L^2 -topology.

The idea of duplication of variables was also used in previous work [18] by Cachier *et al.* but for a different problem. In their case, the registration energy depends on two variables that are both vector fields. The first vector field \mathbf{C} is a set of correspondences between points based on intensity similarity, while the second, denoted by \mathbf{T} , is a smooth vector field constrained by the regularization energy and attracted by the set of correspondences \mathbf{C} .

2.3 Unbiased Registration Constraint

In [14], the authors proposed an unbiased fluid image registration approach. Contrary to classical methods for which the term *unbiased* is used in the sense of symmetric registration, in [14], *unbiased* means that the Jacobian determinants of the deformations recovered between a pair of images follow a log-normal distribution, with zero mean after log-transformation. The authors argued that this distribution is beneficial when recovering change in regions of homogeneous intensity, and in ensuring symmetrical results when the order of two images being registered is switched. As derived in [14] using information theory, the unbiased regularization term is given as

$$R_{UB}(\mathbf{u}) = \int_{\Omega} (|D(\mathbf{x} - \mathbf{u}(\mathbf{x}))| - 1) \log |D(\mathbf{x} - \mathbf{u}(\mathbf{x}))| d\mathbf{x}. \quad (8)$$

It is important to note that R_{UB} generates inverse-consistent deformation maps. The inverse-consistent property of the unbiased technique was shown in a validation study of the unbiased fluid registration methods [19]. Also, to see why minimizing equation (8) leads to unbiased deformation in the logarithmic space, we observe that the integrand is always non-negative, and only evaluates to zero when the deformation \mathbf{g} is volume-preserving everywhere ($|D\mathbf{g}| = 1$ everywhere). Thus, by treating it as a cost, we recover zero-change by minimizing this cost when we compare images differing only in noise.

Given equation (6), we have $D\mathbf{g} = \mathcal{I} - D\mathbf{u} \approx \mathcal{I} - V$, where \mathcal{I} is the 3×3 identity matrix. Therefore, as in subsection 2.2, to simplify the discretization, we introduce

$$R_{UB}(V) = \int_{\Omega} (|\mathcal{I} - V| - 1) \log |\mathcal{I} - V| d\mathbf{x}. \quad (9)$$

Recall that here $|\mathcal{I} - V| = \det(\mathcal{I} - V)$.

2.4 Unbiased Nonlinear Elasticity Registration

The total energy functional employed in this work, is given as a linear combination of the similarity measure F (which is either F_{L^2} from (3) or F_{MI} from (4)), nonlinear elastic regularization R_E in (7), and unbiased regularization R_{UB} in (9):

$$E(\mathbf{u}, V) = F(\mathbf{u}) + R_E(\mathbf{u}, V) + \lambda R_{UB}(V). \quad (10)$$

The explicit weighting parameter is omitted in front of $R_E(\mathbf{u}, V)$, since this term is weighted by Lamé constants ν and μ . We solve the Euler-Lagrange equations in \mathbf{u} and V using the gradient descent method, parameterizing the descent direction by an artificial time t ,

$$\frac{\partial \mathbf{u}}{\partial t} = -\partial E_{\mathbf{u}}(\mathbf{u}, V) = -\partial_{\mathbf{u}} F(\mathbf{u}) - \partial_{\mathbf{u}} R_E(\mathbf{u}, V), \quad (11)$$

$$\frac{\partial V}{\partial t} = -\partial E_V(\mathbf{u}, V) = -\partial_V R_E(\mathbf{u}, V) - \lambda \partial_V R_{UB}(V), \quad (12)$$

which gives systems of three and nine equations, respectively. Explicit expressions for the gradients in these equations are given in Section 3.

Remark 1. The regularization on the deformation \mathbf{g} proposed in this work can be expressed in a general form $R(\mathbf{g}) = \int_{\Omega} R_1(D\mathbf{g})d\mathbf{x} + \int_{\Omega} R_2(|D\mathbf{g}|)d\mathbf{x}$, with $|D\mathbf{g}| := \det(D\mathbf{g})$. For the minimization, an auxiliary variable can also be introduced to simplify the numerical calculations, removing the nonlinearity in the derivatives.

3 Implementation

3.1 The Energy Gradients

Computing the first variation of functional $R_E(\mathbf{u}, V)$, in equation (7), with respect to \mathbf{u} gives the following components of gradient $\partial_{\mathbf{u}} R_E(\mathbf{u}, V)$:

$$\partial_{u_k} R_E(\mathbf{u}, V) = \beta(\partial_1 v_{k1} + \partial_2 v_{k2} + \partial_3 v_{k3} - \Delta u_k), \quad k = 1, 2, 3.$$

The first variation of $R_E(\mathbf{u}, V)$ with respect to V , with $V = [v_{ij}]$, gives $\partial_V R_E(\mathbf{u}, V)$:

$$\begin{aligned} \partial_{v_{11}} R_E(\mathbf{u}, V) &= \beta(v_{11} - \partial_1 u_1) + \nu c_1(1 + v_{11}) + \mu(c_2(1 + v_{11}) + c_5 v_{12} + c_6 v_{13}), \\ \partial_{v_{12}} R_E(\mathbf{u}, V) &= \beta(v_{12} - \partial_2 u_1) + \nu c_1 v_{12} + \mu(c_3 v_{12} + c_5(1 + v_{11}) + c_7 v_{13}), \\ \partial_{v_{13}} R_E(\mathbf{u}, V) &= \beta(v_{13} - \partial_3 u_1) + \nu c_1 v_{13} + \mu(c_4 v_{13} + c_6(1 + v_{11}) + c_7 v_{12}), \\ \partial_{v_{21}} R_E(\mathbf{u}, V) &= \beta(v_{21} - \partial_1 u_2) + \nu c_1 v_{21} + \mu(c_2 v_{21} + c_5(1 + v_{22}) + c_6 v_{23}), \\ \partial_{v_{22}} R_E(\mathbf{u}, V) &= \beta(v_{22} - \partial_2 u_2) + \nu c_1(1 + v_{22}) + \mu(c_3(1 + v_{22}) + c_5 v_{21} + c_7 v_{23}), \\ \partial_{v_{23}} R_E(\mathbf{u}, V) &= \beta(v_{23} - \partial_3 u_2) + \nu c_1 v_{23} + \mu(c_4 v_{23} + c_6 v_{21} + c_7(1 + v_{22})), \\ \partial_{v_{31}} R_E(\mathbf{u}, V) &= \beta(v_{31} - \partial_1 u_3) + \nu c_1 v_{31} + \mu(c_2 v_{31} + c_5 v_{32} + c_6(1 + v_{33})), \\ \partial_{v_{32}} R_E(\mathbf{u}, V) &= \beta(v_{32} - \partial_2 u_3) + \nu c_1 v_{32} + \mu(c_3 v_{32} + c_5 v_{31} + c_7(1 + v_{33})), \\ \partial_{v_{33}} R_E(\mathbf{u}, V) &= \beta(v_{33} - \partial_3 u_3) + \nu c_1(1 + v_{33}) + \mu(c_4(1 + v_{33}) + c_6 v_{31} + c_7 v_{32}), \end{aligned}$$

where

$$\begin{aligned} c_1 &= v_{11} + v_{22} + v_{33} + \frac{1}{2}(v_{11}^2 + v_{21}^2 + v_{31}^2 + v_{12}^2 + v_{22}^2 + v_{32}^2 + v_{13}^2 + v_{23}^2 + v_{33}^2), \\ c_2 &= 2v_{11} + v_{11}^2 + v_{21}^2 + v_{31}^2, & c_5 &= v_{21} + v_{12} + v_{11}v_{12} + v_{21}v_{22} + v_{31}v_{32}, \\ c_3 &= 2v_{22} + v_{12}^2 + v_{22}^2 + v_{32}^2, & c_6 &= v_{31} + v_{13} + v_{11}v_{13} + v_{21}v_{23} + v_{31}v_{33}, \\ c_4 &= 2v_{33} + v_{13}^2 + v_{23}^2 + v_{33}^2, & c_7 &= v_{32} + v_{23} + v_{12}v_{13} + v_{22}v_{23} + v_{32}v_{33}. \end{aligned}$$

We can compute the first variation of (9), obtaining $\partial_V R_{UB}(V)$. We first simplify the notation, letting $J = |\mathcal{I} - V|$. Also, denote $L(J) = (J - 1) \log J$. Hence, $L'(J) =$

$dL(J)/dJ = 1 + \log J - 1/J$. Thus,

$$\begin{aligned}
\partial_{v_{11}} R_{UB}(V) &= -((1 - v_{22})(1 - v_{33}) - v_{32}v_{23})L'(J), \\
\partial_{v_{12}} R_{UB}(V) &= -(v_{23}v_{31} + v_{21}(1 - v_{33}))L'(J), \\
\partial_{v_{13}} R_{UB}(V) &= -(v_{21}v_{32} + (1 - v_{22})v_{31})L'(J), \\
\partial_{v_{21}} R_{UB}(V) &= -(v_{32}v_{13} + v_{12}(1 - v_{33}))L'(J), \\
\partial_{v_{22}} R_{UB}(V) &= -((1 - v_{11})(1 - v_{33}) - v_{13}v_{31})L'(J), \\
\partial_{v_{23}} R_{UB}(V) &= -(v_{12}v_{31} + v_{32}(1 - v_{11}))L'(J), \\
\partial_{v_{31}} R_{UB}(V) &= -(v_{12}v_{23} + v_{13}(1 - v_{22}))L'(J), \\
\partial_{v_{32}} R_{UB}(V) &= -(v_{21}v_{13} + v_{23}(1 - v_{11}))L'(J), \\
\partial_{v_{33}} R_{UB}(V) &= -((1 - v_{11})(1 - v_{22}) - v_{12}v_{21})L'(J).
\end{aligned}$$

3.2 Algorithm

We are now ready to give the algorithm for the unbiased registration via nonlinear elastic regularization.

Algorithm 1 Unbiased Registration via Nonlinear Elastic Regularization

- 1: Initialize $t = 0$, $\mathbf{u}(\mathbf{x}, 0) = 0$, and $V(\mathbf{x}, 0) = 0$.
 - 2: Calculate $V(\mathbf{x}, t)$ using equation (12).
Steps 3-5 describe the procedure for solving equation (11).
 - 3: Calculate the perturbation of the displacement field $\mathbf{R}(\mathbf{x}) = -\partial E_u(\mathbf{u}, V)$.
 - 4: Time step Δt is calculated adaptively so that $\Delta t \cdot \max(\|\mathbf{R}\|_2) = \delta u$, where δu is the maximal displacement allowed in one iteration. Results in this work are obtained with $\delta u = 0.1$.
 - 5: Advance equation (11), i.e. $\partial \mathbf{u}(\mathbf{x}, t)/\partial t = \mathbf{R}(\mathbf{x})$, in time, with time step from step 4, solving for $\mathbf{u}(\mathbf{x}, t)$.
 - 6: If the cost functional in (10) decreases by sufficiently small amount compared to the previous iteration, then stop.
 - 7: Let $t := t + \Delta t$ and go to step 2.
-

4 Results and Discussion

We tested the proposed *unbiased nonlinear elastic* registration model and compared the results to those obtained with the *unbiased fluid registration* method [14], where the unbiased regularization constraint (8) was coupled with the L^2 matching functional (3) and fluid regularization (1), (2). Here, both methods were coupled with the L^2 and mutual information (MI) based similarity measures. In our experiments, we used a pair of serial MRI images ($220 \times 220 \times 220$) from the Alzheimer's Disease Neuroimaging Initiative (ADNI). Since the images were acquired one year apart, from a subject with




Fig. 1. Serial MRI images from the ADNI follow-up dataset (images acquired one year apart) are shown. Volumes I_1 (row 1) and I_2 (row 2) are depicted as a brain volume (column 1) and from sagittal (column 2), axial (column 3), and coronal (column 4) views. Nonrigid registration aligns volume I_2 into correspondence with volume I_1 .

Alzheimer's disease, real anatomical changes are present, which allows methods to be compared in the presence of true biological changes.

In the tests performed using unbiased nonlinear elasticity coupled with L^2 matching, values of $\beta = 20000$ in equation (7) and $\lambda = 2000$ in equation (10) were chosen. For MI matching, $\beta = 80$ and $\lambda = 8$ were used. The values of the Lamé coefficients were chosen to be equal, $\mu = \nu$, in all experiments. Bigger values of μ and ν allow for more smoothing. For unbiased fluid registration model, described in [14], $\lambda = 500$ was chosen for L^2 matching, and $\lambda = 5$ for MI matching.

Figure 2 shows the images being registered along with the resulting Jacobian maps. Results generated using the fluid and nonlinear elasticity based unbiased models are similar, both suggesting a mild volume reduction in gray and white matter and ventricular enlargement that is observed in Alzheimer's disease patients. The advantages of the unbiased nonlinear elasticity model is its more locally plausible reproduction of atrophic changes in the brain and its robustness to original misalignment of brain volumes, which is especially noticeable on the brain surface. The unbiased nonlinear elasticity model coupled with L^2 matching generated very similar results to those obtained with the MI similarity measure, partly because difference images typically contain only noise after registration. Unbiased fluid registration method, however, is more effective in modeling the regional neuroanatomical changes, showing more clearly which parts of the volume have undergone largest tissue changes, such as ventricular enlargement as shown in Figure 2.

Figure 3 shows deformed grids generated with unbiased fluid and unbiased nonlinear elastic registration models. Figure 4 shows the energy decrease per iteration for both models.

In Figure 5, we examined the *inverse consistency* of the mappings [20] generated using unbiased nonlinear elastic registration. Here, the deformation was computed in




Fig. 2. Nonrigid registration was performed on the Serial MRI images from the ADNI Follow-up dataset using unbiased fluid registration and unbiased nonlinear elasticity registration, both coupled with L^2 and MI matching. Jacobian maps are superimposed on the target volume.

both directions (time 2 to time 1, and time 1 to time 2) using mutual information matching. The forward and backward Jacobian maps were concatenated (in an ideal situation, this operation should yield the identity), with the products of Jacobians having values close to 1.

The unbiased nonlinear elasticity model does not require expensive Navier-Stokes solver (or its approximation), which is employed at each iteration for fluid flow models. Hence, in our experiments, unbiased nonlinear elasticity iteration (based on explicit scheme) took 15-20% less time than the unbiased fluid step. Convergence was obtained after roughly the same number of iterations for both methods, resulting in better performance for the unbiased nonlinear elasticity model.

To conclude, we have provided an alternative unified minimization approach to the unbiased fluid registration model and have compared both models. The proposed method proves to be easier to implement and is less computationally intensive. Also, a key benefit of the variational framework and of the numerical scheme of the unbiased




Fig. 3. Results obtained using unbiased fluid registration and unbiased nonlinear elasticity registration, both coupled with L^2 and MI matching. The generated grids are superimposed on top of 2D cross-sections of the 3D volumes (row 1) and are shown separately (row 2).




Fig. 4. Energy per iteration for the unbiased fluid registration and unbiased nonlinear elasticity registration, both coupled with L^2 and MI matching.

nonlinear elastic registration model is its robustness to numerical constraints such as CFL conditions. The method allows to remove the nonlinearity in the derivatives of the unknown \mathbf{u} in the Euler-Lagrange equations. Future studies will examine the registration accuracy of the different models where ground truth is known, and will compare each model's power for detecting inter-group differences or statistical effects on rates of atrophy.

References

1. Modersitzki, J.: Numerical Methods for Image Registration (Numerical Mathematics and Scientific Computation). Oxford University Press, New York (2004)
2. Lord, N., Ho, J., Vemuri, B., Eisenschenk, S.: Simultaneous registration and parcellation of bilateral hippocampal surface pairs for local asymmetry quantification. IEEE Transactions on Medical Imaging **26**(4) (2007) 471–478




Fig. 5. This figure examines the inverse consistency of the unbiased nonlinear elastic registration. Here, the model is coupled with mutual information matching. Jacobian maps of deformations from time 2 to time 1 (column 1) and time 1 to time 2 (column 2) are superimposed on the target volumes. The products of Jacobian maps, shown in column 3, have values close to 1, suggesting inverse consistency.

3. Broit, C.: Optimal Registration of Deformed Images. PhD thesis, University of Pennsylvania (1981)
4. Christensen, G., Rabbitt, R., Miller, M.: Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing* **5**(10) (1996) 1435–1447
5. Brun, C., Lepore, N., Pennec, X., Chou, Y., Lopez, O., Aizenstein, H., Becker, J., Toga, A., Thompson, P.: Comparison of standard and Riemannian elasticity for tensor-based morphometry in HIV/AIDS. *International Conference on Medical Image Computing and Computer Assisted Intervention* (2007)
6. Pennec, X., Stefanescu, R., Arsigny, V., Fillard, P., Ayache, N.: Riemannian elasticity: A statistical regularization framework for non-linear registration. In: *International Conference on Medical Image Computing and Computer Assisted Intervention*. (2005) 943–950
7. Pennec, X.: Left-invariant riemannian elasticity: A distance on shape diffeomorphisms? *International Workshop on Mathematical Foundations of Computational Anatomy* (2006)

- 1–13
8. Bro-Nielsen, M., Gramkow, C.: Fast fluid registration of medical images. In: Visualization in Biomedical Computing. (1996) 267–276
9. D’Agostino, E., Maes, F., Vandermeulen, D., Suetens, P.: A viscous fluid model for multimodal non-rigid image registration using mutual information. *Medical Image Analysis* **7** (2003) 565–575
10. Le Guyader, C., Vese, L.: A combined segmentation and registration framework with a nonlinear elasticity smoother. *UCLA Computational and Applied Mathematics Report* 08-16 (2008)
11. Lin, T., Lee, E.F., Dinov, I., Le Guyader, C., Thompson, P., Toga, A., Vese, L.: A landmark based nonlinear elasticity model for mouse atlas registration. *IEEE International Symposium on Biomedical Imaging* (2008) 788–791
12. Haber, E., Modersitzki, J.: Numerical methods for volume preserving image registration. *Inverse problems*, Institute of Physics Publishing **20**(5) (2004) 1621–1638
13. Haber, E., Modersitzki, J.: Image registration with guaranteed displacement regularity. *International Journal of Computer Vision* **71**(3) (2007) 361–372
14. Yanovsky, I., Thompson, P., Osher, S., Leow, A.: Topology preserving log-unbiased nonlinear image registration: Theory and implementation. *IEEE Conference on Computer Vision and Pattern Recognition* (2007) 1–8
15. Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., Marchal, G.: Automated multi-modality image registration based on information theory. In Bizais, Y., Barillot, C., Di Paola, R., eds.: *Information Processing in Medical Imaging*. Volume 3., Kluwer Academic Publishers, Dordrecht (1995) 264–274
16. Viola, P., Wells, W.: Alignment by maximization of mutual information. *International Conference on Computer Vision* (1995) 16–23
17. Negron-Marrero, P.: A numerical method for detecting singular minimizers of multidimensional problems in nonlinear elasticity. *Numerische Mathematik* **58**(1) (1990) 135–144
18. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: The PASHA algorithm. *Computer Vision and Image Understanding* **89** (2003) 272–298
19. Yanovsky, I., Thompson, P., Osher, S., Hua, X., Shattuck, D., Toga, A., Leow, A.: Validating unbiased registration on longitudinal MRI scans from the Alzheimer’s Disease Neuroimaging Initiative (ADNI). *IEEE International Symposium on Biomedical Imaging* (2008) 1091–1094
20. Christensen, G., Johnson, H.: Consistent image registration. *IEEE Transactions on Medical Imaging* **20**(7) (2001) 568–582

A new algorithm for the computation of the group logarithm of diffeomorphisms

Matias Bossa and Salvador Olmos

GTC, I3A, University of Zaragoza, Spain, {bossa,olmos}@unizar.es *

Abstract. There is an increasing interest on computing statistics of spatial transformations, in particular diffeomorphisms. In the Log-Euclidean framework proposed recently the group exponential and logarithm are essential operators to map elements from the tangent space to the manifold and vice versa. Currently, one of the main bottlenecks in the Log-Euclidean framework applied on diffeomorphisms is the large computation times required to estimate the logarithm. Up to now, the fastest approach to estimate the logarithm of diffeomorphisms is the Inverse Scaling and Squaring (ISS) method. This paper presents a new method for the estimation of the group logarithm of diffeomorphisms, based on a series in terms of the group exponential and the Baker-Campbell-Hausdorff formula. The proposed method was tested on 3D MRI brain images as well as on random diffeomorphisms. A performance comparison showed a significant improvement in accuracy-speed trade-off vs. the ISS method.

1 Introduction

Computational Anatomy is an emerging research field in which anatomy are characterized by means of large diffeomorphic deformation mappings of a given template [1]. The transformation is obtained by non-rigid registration, minimizing a cost function that includes an image matching term, and a regularization term that penalizes large and non-smooth deformations. Several approaches have been proposed in order to analyze the information contained in the transformation. Some methods consist in introducing a right-invariant Riemannian distance between diffeomorphisms, yielding methods with high computational load [2, 3]. Recently, an alternative framework was proposed [4] and consists in endowing the group of transformations with a Log-Euclidean metric. Although this metric is not translation invariant (with respect to the diffeomorphism composition), geodesics are identified with one-parameter subgroups, which can be obtained faster and more easily than the geodesics of a right-invariant Riemannian metric.

One-parameter subgroups of diffeomorphisms $\varphi_t(x)$ are obtained as solutions of the stationary Ordinary Differential Equation (ODE)

$$\frac{d\varphi_t(x)}{dt} = v \circ \varphi_t \equiv v(\varphi_t(x)). \quad (1)$$

* This work was funded by research grants TEC2006-13966-C03-02 from CICYT, Spain

A diffeomorphism $\phi \equiv \phi(x) \equiv \varphi_1(x)$ is defined as the value of the flow φ_t at time one. Any velocity vector field can be written as a linear expansion $v(x) = \sum_{i=1}^D v_i(x) e_i$, where $\{e_i\}_{i=1}^D$ is an orthogonal basis of \mathbb{R}^D . If the components $v_i(x)$ are analytic then the solution of Eq. (1) is also analytic, and is given by the following formal power series (*a.k.a.* Gröbner's Lie Series) [5]:

$$\varphi_t(x) = e^{tV} x = \sum_{n=0}^{\infty} \frac{t^n}{n!} V^n x, \quad (2)$$

where $V \equiv \sum_{i=1}^D v_i(x) \frac{\partial}{\partial x_i}$ is a differential operator and V^n denotes the n -fold self-composition of V .

The Log-Euclidean framework to compute statistics on diffeomorphisms consists in defining a distance between two diffeomorphisms ϕ_1 and ϕ_2 via a norm $\|\cdot\|$ on vector fields: $\text{dist}(\phi_1, \phi_2) = \|v_1 - v_2\|$, where $\phi_i = \exp(v_i)$. Assuming that such a v_i exists we call it the logarithm of ϕ_i , $v_i = \log(\phi_i)$. This metric is equivalent to a bi-invariant Riemannian metric defined on the (abelian) group with the following composition rule: $\phi_1 \odot \phi_2 = \exp(\log(\phi_1) + \log(\phi_2))$. With such a group structure, the distances in the space of diffeomorphisms is computed as the Euclidean distance in the space of vector fields. This distance is inversion-invariant, *i.e.* $\text{dist}(\phi_1, \phi_2) = \text{dist}(\phi_1^{-1}, \phi_2^{-1})$ since $\log(\phi_1^{-1}) = -\log(\phi_1)$ (in fact, it is invariant with respect to the exponentiation to any real power $\neq 0$), and invariant with respect to the new group product, *i.e.* $\text{dist}(\phi_1 \odot \phi_3, \phi_2 \odot \phi_3) = \text{dist}(\phi_1, \phi_2)$, but is not invariant under the standard composition, *i.e.* $\text{dist}(\phi_1 \circ \phi_3, \phi_2 \circ \phi_3) = \text{dist}(\phi_1(\phi_3(x)), \phi_2(\phi_3(x))) \neq \text{dist}(\phi_1, \phi_2)$. Assuming that the logarithm and exponential can be (fast and accurately) computed, any standard statistical analysis can be performed directly on vector fields v_i . This provides a simple way of computing statistics on transformations that avoids the problems of the small deformation frameworks, such as the likely occurrence of non-invertible mappings, and the ones of a right-invariant Riemannian framework, such as the intensive computation cost [6, 7].

Regarding to the computation of the exponential, it was recently proposed to extend the well known Scaling and Squaring (SS) method for computing the matrix exponential to diffeomorphisms [4]. This method basically consist in squaring (self-composing) recursively N times $x + v/2^N \approx \exp(v/2^N) = \exp(v)^{-2^N}$. In a recent study [8], we presented a detailed performance comparison of several methods to compute the group exponential of diffeomorphisms, including the SS method, the forward Euler method and the direct application of the Lie series (2). The SS method achieved the best speed-accuracy trade-off, though two main drawbacks were found: first, the transformation must be computed in the whole domain, contrary to the forward Euler method and the Lie series expansion, that can be computed at a single point; and secondly, there exists an intrinsic lower bound in the accuracy due to the interpolation scheme and the finite size of the sampling grid. Despite of this lower bound, the SS method seems to be fast and accurate enough for most medical image analysis applications.

Regarding to the group logarithm of diffeomorphisms, it was proposed to apply the Inverse Scaling and Squaring (ISS) method [4], based on the following

approximation $v \approx 2^N(\exp(v)^{2^{-N}} - x)$, where the square root of ϕ must be recursively estimated N times. The ISS method is much slower (about 100 times) than the SS method, as the computation of the square root involves an energy functional minimization. In the cases where a diffeomorphism can be written as a composition of two exponentials, $\phi = \exp(v_1) \circ \exp(v_2)$, the logarithm can be estimated with the Baker-Campbell-Hausdorff (BCH) formula, which is a series in terms of the Lie Bracket. In [7] it was tested the BCH formula applied to diffeomorphisms and it was shown that it provides similar accuracy than the ISS method, but with a much lower computational time. In a general case, where the diffeomorphism is neither an exponential of a known vector field, nor a composition of known exponentials, the ISS method seems to be the only available method for estimating the logarithm.

In this work, we propose a new method of computing the logarithm of arbitrary diffeomorphisms in any dimension based on a series involving the group exponential and the BCH formula.

2 A series for the logarithm of diffeomorphisms

The Lie series of the diffeomorphism exponential in Eq. (2) is a generalization of the Taylor expansion of the scalar exponential. However, to our knowledge, the Taylor expansion of the scalar logarithm can not be generalized to the logarithm of diffeomorphisms in the same way. In fact, there exist diffeomorphisms (even infinitely closed to the identity) that cannot be written as the exponential of any vector field in the tangent space [9], *i.e.* the exponential $\phi = \exp(v)$ is not a local diffeomorphism at $v = 0$, therefore a Lie series for the logarithm can not exist. Nevertheless, we will talk about the logarithm v of a diffeomorphism ϕ , and define it as the vector field v whose exponential is closer to ϕ .

The basic idea is that given an initial guess v_0 for v (being v the 'true' logarithm of ϕ), $\exp(-v_0)$ is close to ϕ^{-1} , therefore $\exp(-v_0) \circ \phi$ is close to the identity and can be approximated by $\exp(-v_0) \circ \phi \equiv \exp(\delta v_0) \approx x + \delta v_0$. Then $\delta v_0 \approx \tilde{\delta} v_0 \equiv \exp(-v_0) \circ \phi - x$ and v_0 can be corrected with $\tilde{\delta} v_0$ in order to get a better estimation of v :

$$\begin{aligned} \phi \equiv \exp(v) &= \exp(v_0) \circ (\exp(-v_0) \circ \phi) \\ &= \exp(v_0) \circ \exp(\delta v_0) \\ &\approx \exp(v_0) \circ \exp(\tilde{\delta} v_0) \end{aligned}$$

Recalling that the set of diffeomorphisms is a noncommutative group, v can be approximated by the BCH formula [7]: $v = v_0 + \delta v_0 + 1/2[v_0, \delta v_0] + \dots \approx v_0 + \tilde{\delta} v_0 + 1/2[v_0, \tilde{\delta} v_0] + \dots$, where $[v, w] \equiv \sum_i w_i \partial_i v - v_i \partial_i w$ is the Lie bracket. Finally, we will show that the sequence $v_i = v_{i-1} + \tilde{\delta} v_{i-1} + 1/2[v_{i-1}, \tilde{\delta} v_{i-1}] + \dots$, with $\tilde{\delta} v_{i-1} = \exp(-v_{i-1}) \circ \phi - x$, quickly converges to v . Before going to the more general case of diffeomorphisms, a convergence analysis is presented for the scalar case.

Proposition 1. *Let be $f = e^v$, $v \in \mathbb{R}$, and let $v_n(f)$ be defined by*

$$\begin{aligned} v_0 &= 0 \\ v_n &= v_{n-1} + fe^{-v_{n-1}} - 1 \end{aligned} \quad (3)$$

then the sequence¹ v_n converges to $\lim_{n \rightarrow \infty} v_n(f) = v$ and the error in the n -th term, $\delta_n = v - v_n$, decreases with n as

$$\delta_n \propto \mathcal{O}(\|f - 1\|^{2^n}) \quad (4)$$

Proof. Replacing $f = e^v$ in (3) and expanding the exponential in its power series we get

$$\begin{aligned} v_n &= v_{n-1} + e^{v-v_{n-1}} - 1 = v_{n-1} + e^{\delta_{n-1}} - 1 \\ &= v_{n-1} + 1 + v - v_{n-1} + \sum_{k=2}^{\infty} \frac{\delta_{n-1}^k}{k!} - 1 \\ -\delta_n &= \sum_{k=2}^{\infty} \frac{\delta_{n-1}^k}{k!} \propto \mathcal{O}(\|\delta_{n-1}\|^2) \end{aligned} \quad (5)$$

Recalling that $\delta_1 = v - v_1 = v - (f - 1)$, and expanding v in its power series $v = \log(f) = \sum_{k=1}^{\infty} \frac{(f-1)^k}{k} (-1)^{k+1}$ we have $\delta_1 = (f - 1) - 1/2(f - 1)^2 + \mathcal{O}(\|f - 1\|^3) - (f - 1)$, and with (5) we get (4). \square

In fact, the reader can check that the expansion of v_n in power series of f is

$$\begin{aligned} v_1 &= f - 1 \\ v_2 &= (f - 1) - \frac{(f - 1)^2}{2} + \frac{(f - 1)^3}{3} - \frac{(f - 1)^4}{8} + \frac{(f - 1)^5}{30} + \mathcal{O}((f - 1)^6) \\ v_3 &= \sum_{k=1}^7 \frac{(f - 1)^k}{k} (-1)^{k+1} - \frac{15}{128}(f - 1)^8 + \frac{13}{144}(f - 1)^9 + \mathcal{O}((f - 1)^{10}) \\ &\vdots \\ v_n &= \sum_{k=1}^{2^n-1} \frac{(f - 1)^k}{k} (-1)^{k+1} + \mathcal{O}((f - 1)^{2^n}) \end{aligned}$$

Note that the first $2^k - 1$ terms of the Taylor expansion of the k -th element of the sequence are equal to the Taylor expansion of the logarithm.

Of course it is not practical to compute the logarithm of a scalar number as the limit of a sequence where an exponential must be computed for each term. However, in the case of diffeomorphisms there is no Taylor expansion (or an alternative method except for the ISS) available for the logarithm, and the

¹ Or equivalently the series $v_n = \sum_{i=0}^{n-1} (g^i(f) - 1)$, where $g(f) = e^{1-f}f$ and $g^n(f)$ is the n -fold self-composition of $g(f)$, i.e. $g^0(f) = f$, $g^1(f) = g(f)$ and $g^n(f) = g(g^{n-1}(f))$.

exponential is not very expensive to compute for the usual numerical accuracy required in medical image analysis.

Diffeomorphism logarithm. Let's assume that a diffeomorphisms ϕ can be written as $\phi = \exp(v)$, for some v , in the sense of the formal power series (2). And let's also assume that, for a given vector field $\tilde{\delta}v_n$ close enough to 0, the BCH formula can be applied to compute $v_{n+1} = \log(\exp(v_n) \circ \exp(\tilde{\delta}v_n))$,

$$v_{n+1} = v_n + \tilde{\delta}v_n + 1/2 [v_n, \tilde{\delta}v_n] + 1/12 [v_n, [v_n, \tilde{\delta}v_n]] + 1/12 [[v_n, \tilde{\delta}v_n], \tilde{\delta}v_n] + 1/48 [[v_n, [v_n, \tilde{\delta}v_n]], \tilde{\delta}v_n] + 1/48 [v_n, [[v_n, \tilde{\delta}v_n], \tilde{\delta}v_n]] + \mathcal{O}((\|v_n\| + \|\tilde{\delta}v_n\|)^5)$$

where $[v, w] = \sum_i (w_i \partial v / \partial x_i - v_i \partial w / \partial x_i)$ is the Lie bracket, then the following proposition can be stated:

Proposition 2. *The sequence*

$$\begin{aligned} v_0 &= 0 \\ v_n &= v_{n-1} + \tilde{\delta}v_{n-1} + 1/2 [v_{n-1}, \tilde{\delta}v_{n-1}] + \dots \end{aligned} \quad (6)$$

with $\tilde{\delta}v_{n-1} = \exp(-v_{n-1}) \circ \phi - x$, converges to v with error

$$\delta_n \equiv \log(\exp(v) \circ \exp(-v_n)) \propto \mathcal{O}(\|\phi - x\|^{2^n}). \quad (7)$$

Proof. Eq. (6) is equivalent to

$$\begin{aligned} \exp(v_n) &= \exp(v_{n-1}) \circ \exp(\tilde{\delta}v_{n-1}) \\ \exp(v_n) &= \exp(v_{n-1}) \circ \exp(\exp(-v_{n-1}) \circ \exp(v) - x) \end{aligned}$$

where we used $\phi = \exp(v)$. Now, multiplying on the right by $\phi^{-1} = \exp(-v)$ and expanding $\exp(\exp(-v_{n-1}) \circ \exp(v) - x)$ in its power series we have

$$\begin{aligned} \exp(v_n) \circ \exp(-v) &= \exp(v_{n-1}) \circ \exp(\exp(-v_{n-1}) \circ \exp(v) - x) \circ \exp(-v) \\ \exp(-\delta_n) &= \exp(v_{n-1}) \circ \left(x + (\exp(-v_{n-1}) \circ \exp(v) - x) + \right. \\ &\quad \left. + \sum_{k=2}^{\infty} \frac{(\exp(-v_{n-1}) \circ \exp(v) - x)^k}{k!} \right) \circ \exp(-v) \\ &= x + \exp(v_{n-1}) \circ (\exp(-v_{n-1}) \circ \exp(v) - x)^2 \circ \exp(-v)/2 \\ &\quad + \sum_{k=3}^{\infty} \frac{\exp(v_{n-1}) \circ (\exp(-v_{n-1}) \circ \exp(v) - x)^k \circ \exp(-v)}{k!} \end{aligned}$$

It is not difficult to see that the last term of r.h.s. is of order $\mathcal{O}(\delta_{n-1}^3)$ and $(\exp(-v_{n-1}) \circ \exp(v) - x)^2 = (\exp(-v_{n-1}) \circ \exp(v) - x) \circ (\exp(-v_{n-1}) \circ \exp(v) - x) = \exp(-v_{n-1}) \circ \exp(v) \circ \exp(-v_{n-1}) \circ \exp(v) - 2 \exp(-v_{n-1}) \circ \exp(v) + x$, and left multiplying by $\exp(v_{n-1})$ and right multiplying by $\exp(-v)$ gives $\exp(v) \circ$

$\exp(-v_{n-1}) - 2x + \exp(v_{n-1}) \circ \exp(-v) = \exp(\delta_{n-1}) - 2x + \exp(-\delta_{n-1}) = \delta_{n-1}^2 + \mathcal{O}(\delta_{n-1}^4)$, therefore

$$\begin{aligned} \exp(-\delta_n) &= x + \delta_{n-1}^2/2 + \mathcal{O}(\delta_{n-1}^3) \\ x - \delta_n + \mathcal{O}(\delta_n^2) &= x + \delta_{n-1}^2/2 + \mathcal{O}(\delta_{n-1}^3) \\ \delta_n &\propto \mathcal{O}(\delta_{n-1}^2) \end{aligned} \tag{8}$$

Recalling (7) the initial error $\delta_1 \equiv \log(\exp(v) \circ \exp(-v_1))$, where $v_1 = \tilde{\delta}_0 = \phi - x = \sum_{k=1}^{\infty} \frac{v^k}{k!}$, and v commutes with v^k for all k , therefore $\delta_1 = v - \sum_{k=1}^{\infty} \frac{v^k}{k!} \propto \mathcal{O}(v^2) \propto \mathcal{O}((\phi - x)^2)$. Together with (8) we get (7). \square

In the estimation of the error (7) it was assumed that an infinite number of terms in the BCH formula was used. It can be argued that when a finite number N^{BCH} of terms is used, $\delta_n \propto \mathcal{O}(\|\phi - x\|^{N^{BCH}+1})$, as far as $2^n > N^{BCH}$. Therefore, in practice, N^{BCH} will limit the accuracy of the estimation.

3 Implementation details

The algorithm is initialized with $v_1 = \phi - x$ and then updated following (6), where only 1 or 2 terms of the BCH formula are used. The computation of the Lie Bracket $[\cdot, \cdot]$ involves first order partial derivatives with respect to the spatial coordinates x_i that was implemented as centered finite differences after Gaussian filtering. The filtering is required because the noise in $\tilde{\delta}v_k$ is quickly magnified after successive derivations. The filter width can be estimated using the following rule [8]: $\nu_\phi \leq \nu_v \exp(\max(dv/dx))$, where ν_ϕ (ν_v) is the cut-off frequency of ϕ (v). In our implementation there were still some isolated points in v_k where the second derivative blown up, and a median filter was applied to these points. The exponential followed by a composition $\exp(-v_k) \circ \phi$ present in $\tilde{\delta}v_k$ was not computed with the SS method because, as explained in [8], both the composition and the SS methods introduce errors due to interpolation and the finite grid size. Instead, an integration scheme such as the Forward Euler method, starting at the locations defined by $\phi(x_i)$, being x_i the grid points, is much more accurate.

The gradient descent method required to compute the square roots in the ISS method was based in a simpler gradient than in [4], in particular avoiding the estimation of the inverse diffeomorphism. This implementation provided a faster and more accurate convergence. It might be possible that the original proposal could provide more stable results for large diffeomorphisms.

4 Results

Firstly, a 60x60x60 smoothed random vector field v was exponentiated with the forward Euler method (step size 1/500) providing a diffeomorphism ϕ . We computed the logarithm $\tilde{v} = \log(\phi)$ using (6) ($N^{BCH} = 0, 1$ and 2), and the

ISS method. Accuracy was assessed by velocity vector field error $e = (v - \tilde{v})$ and its corresponding diffeomorphism error $E = (\phi - \exp(\tilde{v}))$. Computations were performed using a 1.83GHz Core 2 Duo processor within a 2GB memory standard computer running Matlab 7.2 under Linux. Linear interpolation was implemented as C source mex files. Computation time was assessed with 'cputime' Matlab function. Figure 1 illustrates the accuracy-speed trade-off and a slice of the corresponding deformed grid. Each estimation method is described by two curves: a dashed/solid line corresponding to error in v and ϕ respectively. Note the large amplitude of the deformation. Figure 2 shows a zoom detail of the error distributions and vector fields.




Fig. 1. Left: Error vs. CPU time in the estimation of the logarithm corresponding to a random simulation. Solid/dashed lines correspond to E and e respectively. The horizontal lines correspond to the small deformation approximation: $v(x) \approx \phi(x) - x$. Right and bottom: Illustration of the deformation grid. Fig. 2 will show the error distribution inside the red square.

Regarding to the accuracy in the estimation of the logarithm v , which is actually our target, the ISS method only provided a midway accuracy between small deformation approximation and the proposed method for $N^{BCH} = 1, 2$. However, the corresponding diffeomorphism had similar accuracy for all methods. Regarding computation time, the proposed method with $N^{BCH} = 1$ was about 10 times faster than ISS method. From Fig. 2 it can be seen that the error was not due to outliers but in spatially correlated regions and far from the boundary.

A second set of experiments were performed on 3D MRI brain data sets. Two $181 \times 217 \times 181$ brain images with isotropic 1mm resolution were randomly selected from LPBA40 database from LONI UCLA [10]. Two non-rigid registra-




Fig. 2. Detail of the spatial distribution of E (left) and e (center) within the red square in Fig. 1. Black/red arrows denote proposed ($N^{BCH} = 1$, $n = 15$) and ISS method ($n = 6$) respectively. Right: Velocity vector fields divided by 10 (black: proposed method; red: ISS; blue: ground truth).

tion methods were used: a diffeomorphic non-rigid registration [11] that provided a vector field v as outcome; and Elastix [12] which is a registration method that provides a deformation field parameterized with B-Splines. In the later there is no warranty of the existence of v .

Left panels in Figures 3 and 4 show the error vs. computation time for the case of diffeomorphic and Elastix registration, respectively. In figure 4 only errors in ϕ are available. Additionally, a representative axial slice of the source image and the corresponding deformed grid is shown at right panel in both figures.




Fig. 3. Left: Error vs. CPU time in the estimation of the logarithm corresponding to a diffeomorphism computed with stationary LDDMM. Solid/dashed lines correspond to E and e respectively. The horizontal lines correspond to the small deformation approximation. Right: Illustration of the deformation grid superimposed on the brain image.




Fig. 4. Left: Error E vs. CPU time in the estimation of the logarithm for a transformation computed with Elastix. The horizontal line corresponds to the small deformation approximation. Right: Illustration of deformation grids superimposed on the brain image.

It is worthy to note that the error e curve of the ISS method in Figures 1 and 3 are very different from the curve shown in [4]. We hypothesized that this behaviour could be explained by the large amplitude of the deformations. In order to verify this possibility the same experiment was performed on the same vector field v divided by a factor of 10. Left panel of Figure 5 shows the error curves and right panel shows a detail of the deformed grid and the corresponding vector field. For this particular case of very small deformations, the ISS method was much more accurate than the logarithm series. Now the shape of the error curve was similar to the one in [4], with smaller error values. Note that all the error values, even for the small deformation approximation, are negligible for medical image analysis applications. When deformations are so small, $v \approx \phi(x) - x$ is accurate enough for standard statistical analysis.




Fig. 5. Left: Idem Fig. 3 for $v/10$. Right: Illustration of the deformation grid and the corresponding velocity vector field.

In our opinion, the accuracy of the ISS method for large diffeomorphisms was limited by the fact that the right way to interpolate diffeomorphisms is unknown. Interpolation of diffeomorphisms is performed in the squaring (self-composition) operation. The composition of diffeomorphisms using a kernel interpolation scheme can provide non-diffeomorphic mappings. In contrast, velocity vector fields belong to a linear vector space, therefore they can be summed or interpolated without leaving the space.

5 Conclusion

We presented a new algorithm for the estimation of the group logarithm of arbitrary diffeomorphisms based on a series in terms of the Lie bracket and the group exponential. This method provided a much better accuracy-speed trade-off than the ISS method to estimate the vector field v defining a diffeomorphism. In particular, at least one term of the BCH formula was essential for the series to provide a significant improvement vs. the ISS method.

Once a fast algorithm to compute the logarithm is available, statistics of the spatial transformations mapping image instances to a given atlas can be easily computed by means of standard multivariate statistics on the tangent space assuming the Log-Euclidean framework. This will be the topic of future studies.

References

1. Grenander, U., Miller, M.I.: Computational anatomy: an emerging discipline. *Q. Appl. Math.* **LVI**(4) (1998) 617–694
2. Miller, M.I., Troune, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. *Ann. Rev. Biomed. Eng* **4** (2002) 375–405
3. Beg, M.F., Miller, M.I., Troune, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* **61** (2) (2005) 139–157
4. Arsigny, V., Commonwick, O., Pennec, X., Ayache, N.: Statistics on diffeomorphisms in a Log-Euclidean framework. *MICCAI* **4190** (2006) 924 – 931
5. Winkel, R.: An Exponential Formula for Polynomial Vector Fields II. Lie Series, Exponential Substitution, and Rooted Trees. *Advances in Mathematics* **147** (1999) 260–303
6. Arsigny, V.: Processing Data in Lie Groups: An Algebraic Approach. Application to Non-Linear Registration and Diffusion Tensor MRI. PhD Thesis, École polytechnique (November 2006)
7. Bossa, M., Hernandez, M., Olmos, S.: Contributions to 3d diffeomorphic atlas estimation: Application to brain images. In: *MICCAI*. (2007) 667–674
8. Bossa, M., Zacur, E., Olmos, S.: Algorithms for computing the group exponential of diffeomorphisms: performance evaluation. In: *MMBIA Workshop at CVPR*. (2008)
9. Glockner, H.: Fundamental problems in the theory of infinite-dimensional Lie groups. *Journal of Geometry and Symmetry in Physics* **5** (2006) 24–35

10. Shattuck, D.W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K.L., Poldrack, R.A., Bilder, R.M., Toga, A.W.: Construction of a 3d probabilistic atlas of human cortical structures. *NeuroImage* **39**(3) (2008) 1064–1080
11. Hernandez, M., Bossa, M.N., Olmos, S.: Registration of anatomical images using geodesic paths of diffeomorphisms parameterized with stationary vector fields. In: MMBIA Workshop at ICCV. (2007) 1–8
12. Staring, M., Klein, S., Pluim, J.P.: Nonrigid Registration with Tissue-Dependent Filtering of the Deformation Field. *Physics in Medicine and Biology* **52**(23) (December 2007) 6879 – 6892

Comparing algorithms for diffeomorphic registration: Stationary LDDMM and Diffeomorphic Demons

Monica Hernandez¹, Salvador Olmos¹, and Xavier Pennec²

¹ Communication Technologies Group (GTC)
Aragon Institute of Engineering Research (I3A)
University of Zaragoza, Spain

² Asclepios, INRIA Sophia-Antipolis, France

Abstract. The stationary parameterization of diffeomorphisms is being increasingly used in computational anatomy. In certain applications it provides similar results to the non-stationary parameterization alleviating the computational charge. With this characterization for diffeomorphisms, two different registration algorithms have been recently proposed: stationary LDDMM and diffeomorphic Demons. To our knowledge, their theoretical and practical differences have not been analyzed yet. In this article we provide a comparison between both algorithms in a common framework. To this end, we have studied the differences in the elements of both registration scenarios. We have analyzed the sensitivity of the regularization parameters in the smoothness of the final transformations and compared the performance of the registration results. Moreover, we have studied the potential of both algorithms for the computation of essential operations for further statistical analysis. We have found that both methods have comparable performance in terms of image matching although the transformations are qualitatively different in some cases. Diffeomorphic Demons shows a slight advantage in terms of computational time. However, it does not provide as stationary LDDMM the vector field in the tangent space needed to compute statistics or exact inverse transformations.

Key words: Computational Anatomy, diffeomorphic registration, stationary parameterization, LDDMM, diffeomorphic Demons

1 Introduction

Computational Anatomy aims at the study of the statistical variability of anatomical structures [1]. Anatomical information is encoded by the spatial transformations existing between anatomical images and a template selected as reference [2]. The analysis of these transformations allows modeling the anatomical variability of a population. In particular, statistical inference can be used in order to identify anatomical differences between healthy and diseased groups or improve the diagnosis of pathologies [3–5]. In the absence of a justified physical model

for inter-subject variability, diffeomorphisms (i.e. differentiable maps with differentiable inverse) provide a convenient mathematical framework to perform this statistical analysis [6, 7].

The Large Deformation Diffeomorphic Metric Mapping (LDDMM) has been considered the reference paradigm for diffeomorphic registration in Computational Anatomy [8, 9]. Diffeomorphisms are represented as end point of paths parameterized by time-varying vector fields defined on the tangent space of a convenient Riemannian manifold. Despite the solid foundations of the mathematical framework, the high computational requirements have made this methodology not much attractive for clinical applications where more efficient registration algorithms are usually preferred.

Recently, an alternative parameterization using stationary vector fields was proposed [7]. This parameterization has been applied for diffeomorphic registration in the variational problem studied in the LDDMM framework [10, 11] and diffeomorphic Demons algorithm [12]. Stationary LDDMM is embedded into the theoretical complexity of the LDDMM framework although it has resulted into a much more efficient algorithm while providing similar registration results. Diffeomorphic Demons is intended as an extension of original Demons algorithm suitable for practical applications due to its efficiency and the quality of registration results.

Although both methods have arisen from different backgrounds, they consider non-rigid registration as a diffusion process [13]. Moreover, they fit into the same variational framework with the same image matching metric and similar characterizations for the diffeomorphic transformations. To our knowledge, the theoretical and practical differences between both methods have not been analyzed yet. In this article, we provide a comparison between both algorithms in this common framework. The elements of the registration scenario (transformation parameterization, image metric, regularization and optimization scheme) have been studied for both methods. In the experimental section we have analyzed the influence of the regularization parameters on the smoothness of the final transformations and compared the performance of the registration results. Moreover, we have studied the potential of both algorithms for the computation of the inverse transformation and the logarithm which constitute essential operations for further statistical analysis.

The rest of the article is divided as follows. In Section 2 we study the elements of stationary-LDDMM and diffeomorphic Demons. Results are presented in Section 3. Finally, Section 4 presents the main concluding remarks.

2 Stationary LDDMM and Diffeomorphic Demons

In Computational Anatomy, diffeomorphic registration is defined as a variational problem involving the characterization of diffeomorphic transformations, an image metric to measure the similarity between the images after registration, a regularization constraint to favor stable numerical solutions, and an optimization technique to search for the optimal transformation in the space of valid

diffeomorphisms. The transformation that deforms the source I_0 into the target I_1 is computed from the minimization of the energy functional

$$E(\varphi) = \frac{1}{\sigma_{\text{reg}}^2} \cdot E_{\text{reg}}(\varphi) + \frac{1}{\sigma_{\text{sim}}^2} \cdot E_{\text{sim}}(I_0, I_1, \varphi) \quad (1)$$

where the weighting factors $1/\sigma_{\text{reg}}^2$ and $1/\sigma_{\text{sim}}^2$ balance the energy contribution between regularization and matching. In this section we study the elements of this registration scenario for stationary LDDMM [10, 11] and diffeomorphic Demons [12].

2.1 Characterization of diffeomorphic transformations

In the LDDMM framework [14, 8], transformations are assumed to belong to a group of diffeomorphisms (i.e. differentiable maps $\varphi : \Omega \rightarrow \Omega$ with differentiable inverse) endowed with a Hilbert differentiable structure of Riemannian manifold, $\text{Diff}(\Omega)$. The tangent space V is a set of Sobolev class vector fields in Ω . The Riemannian metric is defined from the scalar product $\langle v, w \rangle_V = \langle \mathcal{L}v, \mathcal{L}w \rangle_{L^2}$ where \mathcal{L} is a linear invertible differentiable operator. Diffeomorphic transformations are represented by the end point $\varphi = \phi(1)$ of paths of diffeomorphisms $\phi(t)$ parameterized by time-varying flows $v(t)$ of vector fields in V from the solution of the transport equation $\dot{\phi}(t) = v(t, \phi(t))$. The Sobolev structure in V guarantees the existence of diffeomorphic solutions for these equations if $\int_0^1 \|v(t)\|_V^2 dt < \infty$.

In stationary LDDMM [10, 11], paths of diffeomorphisms are parameterized by constant-time flows of vector fields in V . This stationary parameterization is closely related to the group structure defined in $\text{Diff}(\Omega)$ as the paths starting at the identity parameterized using stationary vector fields are exactly the one-parameter subgroups. Diffeomorphisms belonging to one-parameter subgroups can be computed from the group exponential map $\text{Exp} : V \rightarrow \text{Diff}(\Omega)$

$$\varphi = \text{Exp}(w) \quad (2)$$

where w constitutes the infinitesimal generator of the subgroup [7]. Thus, stationary LDDMM restricts transformations to diffeomorphisms belonging to one-parameter subgroups. It has been shown that the set of diffeomorphisms obtained with the stationary parameterization do not comprise all diffeomorphisms in $\text{Diff}(\Omega)$ [15]. Nevertheless, the stationary parameterization has shown to provide a performance similar to the more general non-stationary parameterization on the registration of MRI brain anatomical images [11].

In diffeomorphic Demons [12], transformations are assumed to belong to a group of diffeomorphisms $\text{Diff}(\Omega)$. In contrast to the LDDMM framework, no Riemannian structure is explicitly considered in $\text{Diff}(\Omega)$. Diffeomorphic transformations are represented as the composition of

$$\varphi = \psi \circ \text{Exp}(u) \quad (3)$$

where ψ is an element in $Diff(\Omega)$ and u is a vector field in Ω belonging to a convenient space of vector fields that guarantees the existence of the exponential map and that the composition $\psi \circ \text{Exp}(u)$ remains in $Diff(\Omega)$. This characterization restricts transformations to any element in $Diff(\Omega)$ that can be obtained by finite composition of exponentials of smooth vector fields $\varphi = \text{Exp}(u_1) \circ \dots \circ \text{Exp}(u_N)$.

2.2 Image metric

In stationary LDDMM the image matching energy is defined from

$$E_{\text{sim}}(I_0, I_1, \varphi) = \|I_0 \circ \text{Exp}(w)^{-1} - I_1\|_{L^2}^2 \quad (4)$$

This term could be replaced by other energies proposed in non-stationary LDDMM (as mutual information or cross correlation, among others [16, 17]). In general, the inverse of the minimizer of $E_{\text{sim}}(I_0, I_1, \cdot)$ is not minimizing the reciprocal energy $E_{\text{sim}}(I_1, I_0, \cdot)$. Therefore, if the order of inputs is swapped the method does not provide exact inverse transformations. Introducing inverse consistency in the registration is important as the symmetry in the image matching should be guaranteed by the diffeomorphic transformations used in most of Computational Anatomy applications [17]. In stationary LDDMM, $\text{Exp}(-w)$ and $\text{Exp}(w)$ are exact inverse transformations. Therefore, the inverse consistent version of the image matching energy for stationary LDDMM simply corresponds to

$$E_{\text{sim}}(I_0, I_1, \varphi) = \|I_0 \circ \text{Exp}(w)^{-1} - I_1\|_{L^2}^2 + \|I_1 \circ \text{Exp}(w) - I_0\|_{L^2}^2 \quad (5)$$

Diffeomorphic Demons is associated to the minimization of

$$E_{\text{sim}}(I_0, I_1, \varphi) = \|I_0 \circ \psi \circ \text{Exp}(u) - I_1\|_{L^2}^2 \quad (6)$$

The inverse consistent version of the image matching energy corresponds to ³

$$E_{\text{sim}}(I_0, I_1, \varphi) = \|I_0 \circ \psi \circ \text{Exp}(u) - I_1\|_{L^2}^2 + \|I_1 \circ \zeta \circ \text{Exp}(w) - I_0\|_{L^2}^2 \quad (7)$$

subject to $(\psi \circ \text{Exp}(u))^{-1} = \zeta \circ \text{Exp}(w)$. In this case, minimization involves the solution of a constrained optimization problem leading to a more complex algorithm for general expressions of ψ and ζ .

2.3 Regularization energy

In stationary LDDMM the regularization term is defined as the norm in V of the infinitesimal generator w associated to the diffeomorphism φ , $E_{\text{reg}}(\varphi) = \|w\|_V^2 = \|\mathcal{L}w\|_{L^2}^2$. The regularization term favors solutions to belong to one-parameter subgroups with small energy preventing the transformations to be

³ The inverse $(\psi \circ \text{Exp}(u))^{-1} = \text{Exp}(-u) \circ \psi^{-1}$ is written in the form given by Eq. 3

non-diffeomorphic. The regularization term depends on the selection of the operator \mathcal{L} that is usually related to the physical deformation model imposed on Ω . However, it remains an open question how to choose the best model in non-rigid registration algorithms [13, 10]. In this work we use the diffusive model $\mathcal{L} = Id - \alpha \nabla^2$. This selection restricts w to lie on a space of Sobolev class two.

In Demons framework regularization is externally imposed using Gaussian smoothing on φ and u . This way, the physical deformation model assumed on Ω is roughly equivalent to the combination of a diffusive and a fluid model [18]. It can be shown that the effect of this Gaussian smoothing is equivalent to using the harmonic regularization $E_{\text{reg}}(\varphi) = \|D\varphi - I\|_{\text{fro}}^2$ in Eq. 1 .

2.4 Optimization scheme

In stationary LDDMM, optimization is performed on the tangent space V (optimization on Hilbert spaces). Although classical gradient descent is usually used for numerical optimization [9, 11], more efficient and robust second-order techniques have been recently proposed [10, 19]. These methods are based on Newton's iterative scheme

$$w^{k+1} = w^k - \epsilon \cdot H_w E(w^k)^{-1} \cdot \nabla_w E(w^k) \quad (8)$$

although they differ on the space where first and second order Gâteaux (i.e. directional) derivatives are computed and the simplification of the Hessian term used to overcome the numerical problems posed by Newton's method.

In [10], Gâteaux derivatives are computed on the space of square integrable functions and Levenberg - Marquardt Newton's simplification is used. Thus, the expressions for the gradient and the Hessian are given by

$$(\nabla_w E(w))_{L^2} = 2 (\mathcal{L}^\dagger \mathcal{L})w - (I_0 \circ \text{Exp}(w)^{-1} - I_1) \cdot \nabla(I_0 \circ \text{Exp}(w)^{-1}) \quad (9)$$

$$(H_w E(w))_{L^2} = 2 (\mathcal{L}^\dagger \mathcal{L}) + \nabla(I_0 \circ \text{Exp}(w)^{-1})^T \cdot \nabla(I_0 \circ \text{Exp}(w)^{-1}) \quad (10)$$

With this approach, the action of the linear operator $\mathcal{L}^\dagger \mathcal{L}$ has to be formulated using the matrix representation of the convolution. As a consequence, the algorithm results in a high dimensional matrix inversion problem with large computational requirements. Although inversion is approached by solving a sparse system of linear equations combining Gauss-Seidel with multigrid techniques [20], the memory requirements for diffeomorphic registration hinder the execution in standard machines.

As an alternative, it was proposed in [19] to compute Gâteaux derivatives in the space V using a Gauss-Newton simplification, which leads to

$$(\nabla_w E(w))_V = 2 w - (\mathcal{L}^\dagger \mathcal{L})^{-1}((I_0 \circ \text{Exp}(w)^{-1} - I_1) \cdot \nabla(I_0 \circ \text{Exp}(w)^{-1})) \quad (11)$$

$$(H_w E(w))_V = 2 I_{\mathbb{R}^3} + (\mathcal{L}^\dagger \mathcal{L})^{-2}(\nabla(I_0 \circ \text{Exp}(w)^{-1})^T \cdot \nabla(I_0 \circ \text{Exp}(w)^{-1})) \quad (12)$$

With this approach, the action of the operators $(\mathcal{L}^\dagger \mathcal{L})^{-1}$ and $(\mathcal{L}^\dagger \mathcal{L})^{-2}$ can be formulated using convolution and the update of Eq. 8 can be computed using pointwise operations with smaller memory requirements.

Apart from the computational requirements, Beg et al. provided an additional argument supporting optimization on space V rather than on L^2 [9]. The linear operator $\mathcal{K} = (\mathcal{L}^\dagger \mathcal{L})^{-1}$ is a compact operator in V . Using results from F. Riesz's spectral theory of compact operators, there exists an orthonormal basis $(\varpi_n)_{n \in \mathbb{N}}$ in L^2 with corresponding singular values $(\lambda_n)_{n \in \mathbb{N}}$ such that

$$\mathcal{K} = \sum_{n \in \mathbb{N}} (\lambda_n \langle \cdot, \varpi_n \rangle_{L^2}) \cdot \varpi_n \quad (13)$$

and $\lambda_n \rightarrow 0$ as $n \rightarrow \infty$ due to operator compactness. The expansion of the gradient expressions in this basis yields ⁴

$$(\nabla_w E(w))_{L^2} = \sum_{n \in \mathbb{N}} \left(\frac{1}{\lambda_n} \langle 2w, \varpi_n \rangle_{L^2} + \langle -b, \varpi_n \rangle_{L^2} \right) \cdot \varpi_n \quad (14)$$

$$(\nabla_w E(w))_V = \sum_{n \in \mathbb{N}} (\langle 2w, \varpi_n \rangle_{L^2} + \lambda_n \langle -b, \varpi_n \rangle_{L^2}) \cdot \varpi_n \quad (15)$$

where $b = (I_0 \circ \text{Exp}(w)^{-1} - I_1) \cdot \nabla(I_0 \circ \text{Exp}(w)^{-1})$. Therefore, whereas the action of the linear operator $(\mathcal{L}^\dagger \mathcal{L})^{-1}$ in Eq. 11 remains bounded, the action of $(\mathcal{L}^\dagger \mathcal{L})$ Eq. 9 results into a high frequency components amplification leading to numerical instabilities in the computations.

In diffeomorphic Demons optimization is performed on the group of diffeomorphisms $\text{Diff}(\Omega)$ (optimization on Lie groups) using the iterative scheme

$$u^{k+1} = \frac{I_0 \circ \varphi^k - I_1}{\|\nabla(I_0 \circ \varphi^k) \cdot D\varphi^k\|_{L^2}^2 + (I_0 \circ \varphi^k - I_1)^2 / \tau^2} \cdot (\nabla(I_0 \circ \varphi^k) \cdot (D\varphi^k)) \quad (16)$$

$$\varphi^{k+1} = \varphi^k \circ \text{Exp}(\epsilon \cdot u^{k+1}) \quad (17)$$

where second order techniques are used for the computation of u^k [12].

Regularization is performed at the end of each iteration by smoothing the updated u^k and φ^k using Gaussian filters of standard deviation σ_u and σ_φ , respectively. Moreover, the term $(I_0 \circ \varphi^k - I_1)^2 / \tau^2$ also contributes to the regularization by enforcing the numerical stability of the optimization scheme and controlling the maximum update step length. This term can be seen as a Levenberg - Marquardt approximation of Gauss-Newton's method. Leaving aside the common variational formulation provided in this work, an identical optimization scheme can be obtained from a variational formulation resulting from the introduction of a hidden variable that controls the correspondences between φ and the true transformation [21].

Alternative to this usual Gauss-Newton optimization, the efficient second order scheme introduced in [22] was used in [12]. This led to replacing the term $\nabla(I_0 \circ \varphi)$ in Eq. 16 by its symmetric version $\nabla(I_0 \circ \varphi) + \nabla I_1$. This was shown to improve the rate of convergence with respect to the original Gauss-Newton

⁴ Analogous conclusions can be inferred from expanding the bilinear form associated to Hessian expressions in this basis.

Table 1. Stationary LDDMM registration. Average and standard deviation of the RSSD (%) (upper row) and J_{\min} (lower row) for different values of the regularization parameters α and $1/\sigma_{\text{sim}}^2$. The optimal result for each algorithm is outlined in bold-face. Non-diffeomorphic results are outlined in italics. Note that the algorithms do not converge for values α of order 0.0001.

		Inverse consistent LDDMM. $RSSD = \frac{1}{2} \frac{\ I_0 \circ \varphi - I_1\ _2^2 + \ I_1 \circ \varphi^{-1} - I_0\ _2^2}{\ I_0 - I_1\ _2^2}$.						
$1/\sigma_{\text{sim}}^2 \backslash \alpha$		1.0	0.01	0.0050	0.0025	0.0010	0.0001	
1.0e3		91.56 \pm 3.04	30.53 \pm 3.76	21.51 \pm 2.37	17.42 \pm 4.16	<i>12.18 \pm 3.17</i>	100.00 \pm 0.00	
		0.60 \pm 0.24	0.44 \pm 0.11	0.27 \pm 0.07	0.17 \pm 0.06	<i>-0.17 \pm 0.64</i>	1.00 \pm 0.00	
1.0e4		90.97 \pm 3.11	24.70 \pm 3.10	17.55 \pm 2.10	13.88 \pm 4.13	<i>9.72 \pm 3.72</i>	100.00 \pm 0.00	
		0.59 \pm 0.24	0.31 \pm 0.15	0.19 \pm 0.08	0.10 \pm 0.05	<i>-3.97 \pm 12.28</i>	1.00 \pm 0.00	
1.0e5		90.97 \pm 3.11	24.70 \pm 3.10	17.55 \pm 2.10	13.82 \pm 4.00	<i>9.61 \pm 3.57</i>	100.00 \pm 0.00	
		0.59 \pm 0.24	0.31 \pm 0.15	0.19 \pm 0.08	0.10 \pm 0.05	<i>-3.99 \pm 12.28</i>	1.00 \pm 0.00	
		Symmetric gradient LDDMM. $RSSD = \frac{\ I_0 \circ \varphi - I_1\ _2^2}{\ I_0 - I_1\ _2^2}$.						
$1/\sigma_{\text{sim}}^2 \backslash \alpha$		1.0	0.01	0.0050	0.0025	0.0010	0.0001	
1.0e3		91.66 \pm 2.88	30.44 \pm 3.44	22.02 \pm 2.33	15.79 \pm 1.68	<i>10.88 \pm 1.21</i>	100.00 \pm 0.00	
		0.65 \pm 0.22	0.44 \pm 0.09	0.26 \pm 0.08	0.11 \pm 0.07	<i>0.02 \pm 0.02</i>	1.00 \pm 0.00	
1.0e4		91.11 \pm 2.73	28.61 \pm 3.44	20.99 \pm 2.38	14.81 \pm 1.59	<i>10.09 \pm 1.35</i>	100.00 \pm 0.00	
		0.63 \pm 0.23	0.39 \pm 0.12	0.23 \pm 0.10	0.10 \pm 0.07	<i>0.01 \pm 0.01</i>	1.00 \pm 0.00	
1.0e5		91.11 \pm 2.73	28.83 \pm 3.62	21.49 \pm 2.46	15.38 \pm 2.74	<i>10.09 \pm 1.34</i>	100.00 \pm 0.00	
		0.63 \pm 0.23	0.39 \pm 0.12	0.25 \pm 0.10	0.11 \pm 0.07	<i>0.01 \pm 0.01</i>	1.00 \pm 0.00	

scheme. It should be noted that the efficient second order scheme can be also introduced in Gauss-Newton LDDMM optimization by modifying Eqs. 9 to 12. In the experimental section, we will explore its influence in registration results.

3 Results

3.1 Datasets and experimental setting

A set of 18 T1-MRI images from the Internet Brain Segmentation Repository (IBSR) were used for comparing the performance of the registration algorithms. The images size was $256 \times 256 \times 128$ with a voxel size of $0.94 \times 0.94 \times 1.5$. The images were acquired at the Massachusetts General Hospital and are freely available at <http://www.cma.mgh.harvard.edu/ibsr/data.html>.

In our experiments, one of the images was randomly selected as a template and the remaining of the datasets were registered to this template using stationary LDDMM and diffeomorphic Demons algorithms. Both algorithms were stopped when the magnitude of the update was negligible or after a maximum of 100 iterations. The selection of the optimal regularization parameters is presented below. Other parameters were fixed to typical values used in previous works [10–12]. The parameter ϵ controls the step size made along the search direction in both methods. It was estimated using a backtracking inexact line-search strategy starting from $\epsilon = 1$ for each iteration (see [19] for more details).

Stationary LDDMM was implemented with the scheme given in Eq. 11 and 12. Both the inverse consistent version of the algorithm IC-LDDMM (Eq. 5), and the symmetric gradient optimization scheme SG-LDDMM (use of $\nabla I_0 \circ \text{Exp}(w)^{-1} + \nabla I_1$ instead of $\nabla I_0 \circ \text{Exp}(w)^{-1}$ in the optimization scheme) have been considered in the study. Diffeomorphic Demons was run with the symmetric gradient as proposed in [12].

3.2 Regularization parameters selection

The selection of the regularization parameters is crucial in deformable registration. Strong regularization constraints hinder large deformations and provide a poor intensity match. In contrast, parameters leading to a weak regularization do not constrain the deformation enough and often lead to non diffeomorphic results. The criteria for parameter selection depends on the application. In this work, we wanted to find a tradeoff that provided the best intensity match with minimum deformation. The Relative Sum of Squared Differences (RSSD) between the images before and after registration was used to quantify the image match while Deformation was measured using the Jacobian minimum.

In LDDMM regularization parameter α determine the shape of the kernels associated to the linear operators $(\mathcal{L}^\dagger \mathcal{L})^{-1}$ and $(\mathcal{L}^\dagger \mathcal{L})^{-2}$ in the Fourier domain. Therefore, the selection of α is crucial on the smoothness of the velocity field w . The lower values of α the higher frequency components are conserved on w thus allowing larger deformations. As α goes to 0, the linear operators become close to the identity leading to negligible regularization and non-diffeomorphic solutions. In this work we fixed the parameter $1/\sigma_{\text{reg}}^2$ to 1.0 in order to handle the parameters selection more easily and studied the influence of α and $1/\sigma_{\text{sim}}^2$ on registration results. Table 1 shows the metrics for parameter selection for different values of these parameters. This led us to select $\alpha = 0.0025$ and $1/\sigma_{\text{sim}}^2 = 1.0e4$.

In diffeomorphic Demons, parameters σ_φ and σ_u control the smoothness of the diffeomorphism φ and the velocity field u , respectively. Therefore, the lower values of σ s the higher frequency components are conserved on φ and u allowing larger deformations. In addition, the maximum step length is bounded by $\|u\| \leq 0.5 \cdot \tau$. As τ increases, the maximum magnitude of the velocity field u remains unbounded which can lead to non-diffeomorphic solutions. In this work we fixed the parameters $1/\sigma_{\text{reg}}^2$ and σ_φ to 1.0 mm. Table 2 shows the metrics for parameter selection for different values of σ_u and τ . Optimal values are obtained for $\sigma_u = 1.0$ mm. (close to voxel size) and $\tau = 0.5$ mm.

3.3 Registration results

We have measured the quality of the image matching and the transformations after registration for stationary LDDMM and diffeomorphic Demons. The image matching has been assessed from the RSSD associated to φ and φ^{-1} . For the quantification of the transformations quality we have considered the regularization energies associated to both variational problems, $\|\cdot\|_V^2$ and $\|\cdot\|_{\text{fro}}^2$. In

Table 2. Diffeomorphic Demons registration. Average and standard deviation of the RSSD (%) (upper row) and J_{\min} (lower row) for different values of the regularization parameters σ_{sim} and τ . The optimal result is outlined in boldface. Non-diffeomorphic results are outlined in italics.

		Diffeomorphic Demons. $RSSD = \frac{\ I_0 \circ \varphi - I_1\ _2^2}{\ I_0 - I_1\ _2^2}$.				
τ	σ_u	1.0	1.5	2.0	2.5	3.0
0.5		14.88 ± 1.74	20.91 ± 2.50	31.58 ± 3.67	40.71 ± 4.34	48.15 ± 4.61
		0.07 ± 0.03	0.15 ± 0.04	0.31 ± 0.05	0.47 ± 0.06	0.59 ± 0.06
1.0		<i>9.95 ± 1.18</i>	<i>13.99 ± 1.69</i>	<i>21.73 ± 2.73</i>	<i>29.37 ± 3.72</i>	<i>36.19 ± 4.47</i>
		<i>0.00 ± 0.00</i>	<i>0.02 ± 0.02</i>	0.11 ± 0.05	0.23 ± 0.07	0.35 ± 0.09
2.0		<i>10.91 ± 1.79</i>	<i>12.60 ± 1.26</i>	<i>18.06 ± 2.15</i>	24.51 ± 3.01	30.59 ± 3.75
		<i>-0.01 ± 0.02</i>	<i>-0.00 ± 0.01</i>	<i>0.02 ± 0.02</i>	0.10 ± 0.04	0.19 ± 0.07

stationary LDDMM φ and φ^{-1} have been computed from the exponential map of w [7]. In diffeomorphic Demons the inverse diffeomorphism and the vectors in the tangent space have been computed as proposed in [7]. Table 3 presents the results of these metrics. In addition, Figure 1 shows the registration results for the experiment with the largest ventricle deformation. Figure 2 shows some representative examples of the image-based energy curves during optimization.

3.4 Efficiency

Our experiments were performed on a 2.33 GHZ machine with a C++ implementation based on the ITK library. We found that the computation time per iteration was approximately 41.54 seconds for diffeomorphic Demons, 53.23 seconds for SG-LDDMM and 90.64 seconds for IC-LDDMM. However, it should be noted that if we were also interested in computing the inverse diffeomorphism or the logarithm from the output of diffeomorphic Demons, the computation time of the inverse diffeomorphism would take in average $5\,706 \pm 34$ seconds whereas the computation time for the logarithm would take $17\,463 \pm 10\,681$ seconds.

Table 3. Average and standard deviation of the metrics associated to the registration results. With IC-LDDMM and SG-LDDMM we indicate the inverse consistent and the symmetric gradient version of LDDMM, respectively.

	$RSSD_{(I_0 \circ \varphi, I_1)}$ (%)	$RSSD_{(I_1 \circ \varphi^{-1}, I_0)}$ (%)	$\ \cdot\ _V^2$	$\ \cdot\ _{\text{fro}}^2$
IC-LDDMM	13.42 ± 4.23	14.43 ± 4.24	140.52 ± 28.84	0.17 ± 0.05
SG-LDDMM	14.81 ± 1.59	15.47 ± 2.50	166.62 ± 12.05	0.20 ± 0.03
Demons	14.88 ± 1.74	19.00 ± 4.72	2626.70 ± 6069.90	0.13 ± 0.01

4 Discussion and conclusions

In this article we presented a theoretical and experimental comparison of two diffeomorphic registration techniques that use stationary vector fields to compute diffeomorphisms. We analyzed the differences in the elements of both registration scenarios, studied the influence of the regularization parameters on the quality of the final transformations and compared the performance of the registration results. For stationary LDDMM, we considered both the inverse consistent version of the algorithm and the symmetric gradient optimization scheme.




Fig. 1. Visual assessment of diffeomorphic registration algorithms. First row shows axial slices of the reference (left) and the deforming image (middle) together with the differences before registration (right). Second row shows the differences after registration with IC-LDDMM (left), SG-LDDMM (middle) and Demons (right). Third row shows the corresponding diffeomorphisms.

For diffeomorphic Demons we just considered the symmetric gradient optimization scheme. We found that parameters α in LDDMM and τ in diffeomorphic Demons were strongly influencing the smoothness of the final transformation. There even exist combinations of such parameters that provide non diffeomorphic transformations. It seems that non diffeomorphic transformations at RSSD approximately below 13% can be achieved in our datasets. This should be taken into account in parameter selection for specific applications.

We found that both algorithms provided a similar intensity matching (average RSSD ranging from 13.42 to 14.88 %). However, in some cases, both algorithms locally showed different performance. As can be appreciated in Figure 1, larger deformations in stationary LDDMM yielded a higher image matching in locations such as, for example, ventricles tails. SG-LDDMM provided a slightly better consistent inverse image matching than diffeomorphic Demons (average RSSD of 15.5 for stationary LDDMM vs RSSD of 19.0 for Demons). In this case, average RSSD differences between methods were statistically different for diffeomorphic Demons. The regularization energy in V showed to be much higher in the case of diffeomorphic Demons. This may be due to the bad numerical con-




Fig. 2. Representative examples of the image matching curves during optimization.

ditioning of the logarithm map computation or to the absence of smoothness constraints on the second order derivatives of the transformations. As shown in Figure 2, IC-LDDMM provided the highest rate of convergence in all cases. At the initial stages of optimization, diffeomorphic demons showed the worst performance in the great majority of cases although it usually reached SG-LDDMM performance at convergence.

Diffeomorphic Demons was 1.28 times faster than SG-LDDMM and 2.18 times faster than IC-LDDMM. However, it should be noted that stationary LDDMM provides elements on the tangent space instead of transformations as output. This allows to compute exponentials and inverses with a low computational cost. On the contrary, diffeomorphic Demons only provides transformations as output. Therefore, logarithms and inverses have to be estimated using quite computationally expensive iterative algorithms.

In conclusion, both methods may be considered close from a theoretical point of view and equivalent from a practical point of view for registration purposes. Diffeomorphic Demons demonstrated similar intensity matching performances to stationary LDDMM at a slightly lower computational cost. It should be advisable to select this algorithm for registration applications where the efficiency of the algorithm is crucial, while stationary LDDMM should be selected for applications where either the transformation smoothness or the inverse consistency is important, or if the inverse transformations or logarithm maps need to be computed. The selection between SG-LDDMM or IC-LDDMM would again depend on the trade-off between computation time and accuracy for the specific application.

References

1. Grenander, U., Miller, M.: Computational Anatomy: an emerging discipline. *Quart. Appl. Math.* **56** (1998) 617 – 694

2. Grenander, U.: General pattern theory. Oxford University Press (1994)
3. Thompson, P.M., et al.: Detection and mapping of abnormal brain structure with a probabilistic atlas of cortical surfaces. *J. Comp. Ass. Tomography* **21**(4) (1997) 567 – 581
4. Thompson, P.M., et al.: Cortical change in Alzheimer’s disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex* **11**(1) (2001) 1 – 16
5. Miller, M.I.: Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *Neuroimage* **23** (2004) 19–33
6. Miller, M.I., Troune, A., Younes, L.: Geodesic shooting for computational anatomy. *J. Math. Imaging Vis.* **24** (2006) 209–228
7. Arsigny, V., Commonwicks, O., Pennec, X., Ayache, N.: Statistics on diffeomorphisms in a Log-Euclidean framework. *MICCAI’06, LNCS* **4190** (2006) 924 – 931
8. Dupuis, P., Grenander, U., Miller, M.: Variational problems on flows of diffeomorphisms for image matching. *Quart. Appl. Math.* (1998) 587 – 600
9. Beg, M., Miller, M., Troune, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* **61** (2) (2005) 139–157
10. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1) (2007) 95 – 113
11. Hernandez, M., Bossa, M.N., Olmos, S.: Registration of anatomical images using geodesic paths of diffeomorphisms parameterized with stationary vector fields. *MMBIA’07* (2007)
12. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic image registration with the demons algorithm. *MICCAI’07, LNCS* **4792** (2007) 319 – 326
13. Modersitzki, J.: Numerical methods for image registration. Oxford University Press (2004)
14. Troune, A.: Diffeomorphism groups and pattern matching in image analysis. *Int. J. Comput. Vis.* **28** (1998) 213 – 221
15. Grabowski, J.: Free subgroups of diffeomorphism groups. *Fundam. Math.* **131** (1988) 103 – 121
16. Lorenzen, P., Prastawa, M., Davis, B., Gerig, G., Bullitt, E., Joshi, S.: Multi-modal image set registration and atlas formation. *Med. Image. Anal.* **10** (2006) 440 – 451
17. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image. Anal.* **12** (1998) 26 – 41
18. Pennec, X., Cachier, P., Ayache, N.: Understanding the Demons algorithm: non-rigid registration by gradient descent. *MICCAI’99, LNCS* **1679** (1999) 597 – 605
19. Hernandez, M., Olmos, S.: Gauss-Newton optimization in diffeomorphic registration. *ISBI’08* (2008)
20. Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical recipes in C: the art of scientific computing. Cambridge University Press, Cambridge (1992)
21. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: the PASHA algorithm. *Comput. Vis. Image Underst.* **89** (2003) 272 – 298
22. Malis, E.: Improving vision-based control using efficient second-order minimization techniques. *IEEE Int. Conf. Robot. Aut.* **2** (2004) 1843 – 1848

Brain Mapping with the Ricci Flow Conformal Parameterization and Multivariate Statistics on Deformation Tensors

Yalin Wang^{1,2}, Xiaotian Yin³, Jie Zhang⁴, Xianfeng Gu³, Tony F. Chan², Paul M. Thompson¹, and Shing-Tung Yau⁵

¹ Lab. of Neuro Imaging, UCLA School of Medicine, Los Angeles, CA 90095, USA,

² Mathematics Department, UCLA, Los Angeles, CA 90095, USA,

³ Computer Science Dept., SUNY at Stony Brook, Stony Brook, NY 11794, USA,

⁴ Mathematics Department, Zhejiang University, Hangzhou, China

⁵ Department of Mathematics, Harvard University, Cambridge, MA 02138, USA,
{ylwang}@math.ucla.edu

Abstract. By solving the Yamabe equation with the discrete surface Ricci flow method, we can conformally parameterize a multiple boundary surface by a multi-hole disk. The resulting parameterizations do not have any singularities and they are intrinsic and stable. For applications in brain mapping research, first, we convert a cortical surface model into a multiple boundary surface by cutting along selected anatomical landmark curves. Secondly, we conformally parameterize each cortical surface using a multi-hole disk. Inter-subject cortical surface matching is performed by solving a constrained harmonic map in the canonical parameter domain. To map group differences in cortical morphometry, we then compute a manifold version of Hotelling’s T^2 test on the Jacobian matrices. Permutation testing was used to estimate statistical significance. We studied brain morphology in 21 patients with Williams Syndrome (WE) and 21 matched healthy control subjects with the proposed method. The results demonstrate our algorithm’s potential power to effectively detect group differences on cortical surfaces.

1 Introduction

Surface-based modeling is valuable in brain imaging to help analyze anatomical shape, to detect abnormalities of cortical surface folding, and to statistically combine or compare 3D anatomical models across subjects. Even so, a direct mapping between two 3D surfaces from different subjects is challenging to compute. Often, higher order correspondences must be enforced between specific anatomical points, curved landmarks, or subregions lying within the two surfaces. This is often achieved by first mapping each of the 3D surfaces to canonical parameter spaces such as a sphere [1, 2] or a planar domain [3]. A flow, computed in the parameter space of the two surfaces [4, 5], then induces a correspondence field in 3D. This flow can be constrained using anatomic landmark points or curves, by constraining the mapping of surface regions represented implicitly using level

sets [3], or by using currents to represent anatomical variation [6]. Feature correspondence between two surfaces can be optimized by using the L^2 -norm to measure differences in curvature profiles or convexity [1] or by using mutual information to align scalar fields of various differential geometric parameters defined on the surface [7]. Artificial neural networks may also be used to rule out or favor certain types of feature matches [8]. Finally, correspondences may be determined by using a minimum description length (MDL) principle, based on the compactness of the covariance of the resulting shape model [9]. Anatomically homologous points can then be forced to match across a dataset. Thodberg [10] identified problems with early MDL approaches and extended them to an MDL appearance model, when performing unsupervised image segmentation.

All oriented surfaces have conformal structures. The conformal structure is, in some respects, more flexible than the Riemannian metric but places more restrictions on the surface morphology than the topological structure. The Ricci flow method can conformally map an open boundary surface to a multi-hole disk [11]. Compared with other conformal parameterization methods [12–15], the Ricci flow method can handle cortical surfaces with complicated topologies without singularities. The continuous Ricci flow conformally deforms a Riemannian metric on a smooth surface such that the Gaussian curvature evolves like a heat diffusion process. In the discrete case, with the circle packing metric, the Ricci flow can be formulated in a variational setting and solved by the Newton method [11].

Tensor-based morphometry is widely used in computational anatomy as a means to understand shape variation between structural brain images. Techniques based on Riemannian manifolds to compare deformation tensors or strain matrices were introduced in [16–18]. In [19], the full deformation tensors were used in the context of tensor-based morphometry. In a conformal parameterization, the original metric tensor is preserved up to a constant. The conformal parametrization provides an ideal framework to apply tensor based morphometry on surfaces, to help understand shape variation between structural brain images.

In this paper, we use the Ricci flow method to compute a conformal mapping between cortical surfaces and a multi-hole surface. Then we compute a direct cortical surface correspondence by computing a constrained harmonic map on the parameter domain. We apply multivariate statistics to the Jacobian matrices to study cortical surface variation between a group of patients with Williams syndrome (WS) and a group of healthy control subjects. WS is a genetic disorder in which the cortex develops abnormally, but the scope and type of systematic differences is unknown [20]. In our experimental results, we identified several significantly different areas on the left and right cortical surfaces between WS patients and control subjects.

2 Ricci Flow Conformal Parameterization

In this section, we introduce the theory of Ricci flow in the continuous setting, and then generalize it to the discrete setting.

2.1 Ricci Flow on Continuous Surfaces

Riemannian Metric and Gaussian Curvature All the concepts used here may be found, with detailed explanations, in [21]. Suppose S is a C^2 smooth surface embedded in \mathbb{R}^3 with local parameters (u_1, u_2) . Let $\mathbf{r}(u_1, u_2)$ be a point on S and $d\mathbf{r} = \mathbf{r}_1 du_1 + \mathbf{r}_2 du_2$ be the tangent vector defined at that point, where $\mathbf{r}_1, \mathbf{r}_2$ are the partial derivatives of \mathbf{r} with respect to u_1 and u_2 , respectively. The *Riemannian metric* or the *first fundamental form* is:

$$\langle d\mathbf{r}, d\mathbf{r} \rangle = \sum \langle \mathbf{r}_i, \mathbf{r}_j \rangle du_i du_j, \quad i, j = 1, 2. \quad (1)$$

The Gauss map $G : S \rightarrow \mathbb{S}^2$ from the surface S to the unit sphere \mathbb{S}^2 maps each point p on the surface to its normal $\mathbf{n}(p)$. The *Gaussian curvature* $K(p)$ is defined as the *Jacobian of the Gauss map*. Intuitively, it is the ratio between the infinitesimal area of the image of the Gauss map and the infinitesimal area on the original surface.

The total curvature of a compact surface is determined by the topology of the surface: $\int_S K dA + \int_{\partial S} k_g ds = 2\pi\chi(S)$, where ∂S is the boundary of the surface S , k_g is the geodesic curvature, and $\chi(S)$ is the Euler characteristic of the surface (an integer).




Fig. 1. Properties of Conformal Mapping: Conformal mappings transform infinitesimal circles to infinitesimal circles and preserve the intersection angles among the circles. Here, infinitesimal circles are approximated by finite ones.

Conformal deformation Let S be a surface embedded in \mathbb{R}^3 . S has a Riemannian metric induced from the Euclidean metric of \mathbb{R}^3 , denoted by \mathbf{g} . Suppose $u : S \rightarrow \mathbb{R}$ is a scalar function defined on S . It can be verified that $\bar{\mathbf{g}} = e^{2u}\mathbf{g}$ is also a Riemannian metric on S , and angles measured by \mathbf{g} are equal to those measured by $\bar{\mathbf{g}}$. We say $\bar{\mathbf{g}}$ is a *conformal deformation* from \mathbf{g} . Figure

1 shows that a conformal deformation maps infinitesimal circles to infinitesimal circles and preserves their intersection angles.

When the Riemannian metric is conformally deformed, curvatures will also be changed accordingly. Suppose \mathbf{g} is changed to $\bar{\mathbf{g}} = e^{2u}\mathbf{g}$, the Gaussian curvature will become $\bar{K} = e^{-2u}(-\Delta_{\mathbf{g}}u + K)$, where $\Delta_{\mathbf{g}}$ is the Laplacian-Beltrami operator under the original metric \mathbf{g} . The geodesic curvature will become $\bar{k} = e^{-u}(\partial_{\mathbf{r}}u + k)$, where \mathbf{r} is the tangent vector orthogonal to the boundary.

Smooth Surface Ricci Flow Suppose S is a smooth surface with a Riemannian metric \mathbf{g} . The Ricci flow deforms the metric $\mathbf{g}(t)$ according to the Gaussian curvature $K(t)$ (induced by $\mathbf{g}(t)$ itself), where t is the time parameter

$$\frac{dg_{ij}(t)}{dt} = -2K(t)g_{ij}(t). \quad (2)$$

If we replace the metric in Eq. 2 with $g(t) = e^{2u(t)}g(0)$, then the Ricci flow can be simplified as $du(t)/dt = -2K(t)$, which states that the metric should change according to the curvature.

The Ricci flow can be easily modified to compute a metric with a *user-defined* curvature $\bar{K} : du(t)/dt = 2(\bar{K} - K)$. The resulting metric $\mathbf{g}(\infty)$ will induce the user-defined curvature \bar{K} .

The Ricci flow has been proven to converge. For surfaces with non-positive and positive Euler numbers, the proofs were given by Hamilton [22] and Chow [23] respectively. For a closed surface, if the total area is preserved during the flow, the Ricci flow will converge to a metric such that the Gaussian curvature is constant everywhere.

2.2 Ricci Flow on Discrete Surfaces

In engineering fields, smooth surfaces are often approximated by simplicial complexes (triangle meshes). Key concepts, such as the metric, curvature, and conformal deformation in the continuous setting can be generalized to the discrete setting. We denote a triangle mesh as Σ , the mesh boundary as $\partial\Sigma$, a vertex set as V , an edge set as E , and a face set as F . e_{ij} represents the edge connecting vertices v_i and v_j , and f_{ijk} denotes the face formed by v_i , v_j , and v_k .

Discrete Riemannian Metric and Gaussian Curvature A Riemannian metric on a mesh Σ is a piecewise constant metric with cone singularities at vertices.

The edge lengths of a mesh Σ are sufficient to define the Riemannian metric, $l : E \rightarrow \mathbb{R}^+$, as long as for each face f_{ijk} , the edge lengths satisfy the triangle inequality: $l_{ij} + l_{jk} > l_{ki}$.

The discrete Gaussian curvature K_i on a vertex $v_i \in \Sigma$ can be computed from the angle deficit,

$$K_i = \begin{cases} 2\pi - \sum_{f_{ijk} \in F} \theta_i^{jk}, & v_i \notin \partial\Sigma \\ \pi - \sum_{f_{ijk} \in F} \theta_i^{jk}, & v_i \in \partial\Sigma \end{cases} \quad (3)$$




Fig. 2. Circle Packing Metric (a) Flat circle packing metric (b) Circle packing metric on a triangle.

where θ_i^{jk} represents the corner angle attached to vertex v_i in the face f_{ijk} . The discrete Gaussian curvatures are determined by the discrete metrics. And the total discrete curvature is, similar to the smooth case, a topological invariant: $\sum_{v_i \in V} K_i = 2\pi\chi(M)$.

Discrete Conformal Deformation In the discrete setting, conformal deformation is carried out using the concept of *circle packing metric*, which was introduced by Thurston in [24].

By approximating infinitesimal circles using circles with finite radii, a *circle packing metric* of Σ can be denoted as (Γ, Φ) , where Γ is a vertex function, $\Gamma : V \rightarrow \mathbb{R}^+$, which assigns a radius γ_i to the vertex v_i ; Φ is an edge weight function, $\Phi : E \rightarrow [0, \frac{\pi}{2}]$, which assigns an acute angle (i.e. weight) $\Phi(e_{ij})$ to each edge e_{ij} . Figure 2 illustrates the circle packing metric. Each vertex v_i has a circle whose radius is γ_i . For each edge e_{ij} , the intersection angle ϕ_{ij} is defined through two circles around v_i and v_j , which either intersect or are tangent.

Two circle packing metrics (Γ_1, Φ_1) and (Γ_2, Φ_2) on the same mesh are *conformally equivalent* if $\Phi_1 \equiv \Phi_2$. A *conformal deformation* of a circle packing metric only modifies the vertex radii and preserves the intersection angles of the edges.

Discrete Surface Ricci Flow Suppose (Σ, Φ) is a weighted mesh with an initial circle packing metric. Similar to the smooth setting, if we set $\bar{\mathbf{k}} = (\bar{K}_1, \bar{K}_2, \dots, \bar{K}_n)^T$ to be the user-defined target curvature, the discrete Ricci flow can be defined as :

$$\frac{du_i(t)}{dt} = (\bar{K}_i - K_i), \quad (4)$$

The Discrete Ricci flow can be formulated in the variational setting; namely, it is a negative gradient flow of a special energy form. Let (Σ, Φ) be a weighted mesh with spherical (Euclidean or hyperbolic) background geometry. For arbitrary two vertices v_i and v_j , the following symmetric relation holds: $\partial K_i / \partial u_j = \partial K_j / \partial u_i$. Let $\omega = \sum_{i=1}^n K_i du_i$ be a differential one-form [25]; the symmetric relation guarantees that this one-form is closed (curl free) in the metric space:

$d\omega = 0$. Then by Stokes theorem, the following integration is path independent:

$$f(\mathbf{u}) = \int_{\mathbf{u}_0}^{\mathbf{u}} \sum_{i=1}^n (\bar{K}_i - K_i) du_i, \quad (5)$$

where n is the number of vertices, $\mathbf{u}_i = \log(\gamma_i)$, γ_i is the radius associated with edge i , and \mathbf{u}_0 is an arbitrary initial metric.

The above integration (Eq. 5) is called the *discrete Ricci energy*, which is well-defined. The discrete Ricci energy has been proved to be strictly convex (i.e., its Hessian is positive definite) in [26]. The global minimum uniquely exists, which gives the desired discrete metric that induces $\bar{\mathbf{k}}$. The discrete Ricci flow is the negative gradient flow of this energy, and it converges to the global minimum.

As in [11], the discrete surface Ricci flow method was used to solve the Yamabe equation [11] and conformally map an open boundary cortical surface to a multi-hole disk.

3 Multivariate Statistics on Deformation Tensors

3.1 Derivative Map

Suppose $\phi : S_1 \rightarrow S_2$ is a map from the surface S_1 to the surface S_2 . In order to simplify the formulation, we use the isothermal coordinates of both surfaces for the arguments. Let $(u_1, v_1), (u_2, v_2)$ be the isothermal coordinates of S_1 and S_2 respectively. The Riemannian metric of S_i is represented as $\mathbf{g}_i = e^{2\lambda_i}(du_i^2 + dv_i^2)$, $i = 1, 2$.

On the local parameters, the map ϕ can be represented as $\phi(u_1, v_1) = (\phi_1(u_1, v_1), \phi_2(u_1, v_1))$. The *derivative map* of ϕ is the linear map between the tangent spaces, $d\phi : TM(p) \rightarrow TM(\phi(p))$, induced by the map ϕ . In the local parameter domain, the derivative map is the Jacobian of ϕ ,

$$d\phi = \begin{pmatrix} \frac{\partial \phi_1}{\partial u_1} & \frac{\partial \phi_1}{\partial v_1} \\ \frac{\partial \phi_2}{\partial u_1} & \frac{\partial \phi_2}{\partial v_1} \end{pmatrix}.$$

Let the position vector of S_1 be $\mathbf{r}(u_1, v_1)$. Denote the tangent vector fields as $\frac{\partial}{\partial u_1} = \frac{\partial \mathbf{r}}{\partial u_1}$, $\frac{\partial}{\partial v_1} = \frac{\partial \mathbf{r}}{\partial v_1}$. Because (u_1, v_1) are isothermal coordinates, $\frac{\partial}{\partial u_1}$ and $\frac{\partial}{\partial v_1}$ only differ by a rotation of $\pi/2$. Therefore, we can construct an orthonormal frame on the tangent plane on S_1 as $\{e^{-\lambda_1} \frac{\partial}{\partial u_1}, e^{-\lambda_1} \frac{\partial}{\partial v_1}\}$. Similarly, we can construct an orthonormal frame on S_2 as $\{e^{-\lambda_2} \frac{\partial}{\partial u_2}, e^{-\lambda_2} \frac{\partial}{\partial v_2}\}$.

The derivative map under the orthonormal frames is represented as

$$d\phi = e^{\lambda_2 - \lambda_1} \begin{pmatrix} \frac{\partial \phi_1}{\partial u_1} & \frac{\partial \phi_1}{\partial v_1} \\ \frac{\partial \phi_2}{\partial u_1} & \frac{\partial \phi_2}{\partial v_1} \end{pmatrix}.$$

In practice, smooth surfaces are approximated by triangle meshes. The map ϕ is approximated by a simplicial map, which maps vertices to vertices, edges to

edges and faces to faces. The derivative map $d\phi$ is approximated by the linear map from one face $[v_1, v_2, v_3]$ to another one $[w_1, w_2, w_3]$. First, we isometrically embed the triangle $[v_1, v_2, v_3], [w_1, w_2, w_3]$ onto the plane \mathbb{R}^2 , the planar coordinates of the vertices of v_i, w_j are denoted using the same symbol v_i, w_j . Then we explicitly compute the linear matrix for the derivative map $d\phi$,

$$d\phi = [w_3 - w_1, w_2 - w_1][v_3 - v_1, v_2 - v_1]^{-1}.$$

In our work, we use multivariate statistics on deformation tensors [19], but adapt the concept to surface tensors. Let J be the derivative map and define the deformation tensors as $S = (J^T J)^{1/2}$. Instead of analyzing shape change based on the eigenvalues of the deformation tensor, we consider a new family of metrics, the ‘‘Log-Euclidean metrics’’ [18]. These metrics make computations on tensors easier to perform, as they are chosen such that the transformed values form a vector space, and statistical parameters can then be computed easily using standard formulae for Euclidean spaces.

We apply Hotelling’s T^2 test on the log-Euclidean space of the deformation tensors. Given two groups of n -dimensional vectors $S_i, i = 1, \dots, p, T_j, j = 1, \dots, q$, we use the Mahalanobis distance M to measure the group mean difference,

$$M = (\log \bar{S} - \log \bar{T}) \Sigma^{-1} (\log \bar{S} - \log \bar{T})$$

where \bar{S} and \bar{T} are the means of the two groups and Σ is the combined covariance matrix of the two groups.

4 Experimental Results

We tested our algorithm on brain anatomic surfaces extracted from 3D MRI scans of a group of 21 WS individuals and a group of 21 healthy control subjects. The cerebral cortex and landmark data are the same ones used in [20]. We tested our algorithm with different landmark sets. The first set included four selected landmark curves per hemisphere: the Central Sulcus, Superior Temporal Sulcus, Primary Intermediate Sulcus and Middle Frontal Control Line. A second set of constraint curves included seven selected landmark curves (the three new landmarks are the Precentral Sulcus, Paracentral Sulcus and Subparietal Sulcus). After we cut a cortical surface open along the selected landmark curves, the cortical surface becomes topologically equivalent to an open boundary genus 3 (4 landmarks) or genus 6 (7 landmarks) surface. So the cortical surface can be conformally mapped to a multi-hole disks with 4 and 7 boundaries, respectively. Examples of cortical surfaces with landmark curves overlaid (after cuts introduced) and their parameterization results are shown in Figure 3.

Because of the shape difference between different cortices, the centers and the radii of the inner circles are different. By computing a constrained harmonic map from each individual conformal map to a canonical multi-hole disk in the parameter domain, we can easily compute a direct surface correspondence between each of the cortical surfaces [11]. Currently, the reference canonical multi-hole disk

is arbitrarily chosen. For landmark curve matching, we guaranteed the matching of both curve ends. For other parts, we match curves based on unit speed parameterization on both curves.

Based on the surface matching results, the Jacobian matrices were computed as described in Section 3. For each point on the cortical surface, we ran permutation test with 5000 random assignments of subjects to groups to estimate the statistical significance of the areas with group differences in surface morphometry. Figure 4 and 5 illustrates our experimental results. We compared left and right cortical surface morphology between 21 control subjects and 21 WS patients with mappings constrained by a total of 4 and 7 selected landmark curves. Different sets of landmarks were used as anchors to evaluate the impact of the choice of anatomical constraints on the results. The significance maps, for the left and right hemispheres, show group differences at the voxel level, between WS patients and control subjects. Mappings with 4 and 7 selected landmark curves were computed. We detected few significant shape differences between left and right cortical surfaces (i.e., anatomical asymmetries) in both the control group and the WS group. Even so, we did find significant shape differences for both left and right cortical surfaces between WS and control subjects. We also found the regions with differences detected at the voxel level were consistent for the mappings computed with 4 and 7 selected landmark curves. However, detection power was not as high as expected in regions around the landmarks. Our future research will examine the statistics in the vicinity of the chosen landmarks, and multiple comparison correction methods, for example based on controlling the false discovery rate, will be used to assess the overall significance of the group differences.




Fig. 3. Cortical surfaces with landmark curves and their conformal parameterization results. The first row shows a cortex with 4 landmarks and the second row shows a cortex with 7 landmarks (one landmark is not visible in this view).


5 Conclusions and Future Work

We applied the Ricci flow conformal parameterization for brain cortical surface registration. Based on the derivative map between two matching surfaces, a multivariate statistic on the Jacobian matrices was used to study surface morphometry in WS. Experimental results suggest that the significantly different areas were consistent with respect to the choice of landmark constraints and the algorithm has the potential to detect systematic surface abnormalities associated with disease. In the future, we will further study other possible statistics on the Jacobian matrices and find the optimal statistics for analyzing deformations computed using our Ricci flow conformal parameterization-based surface matching method. We will also validate our algorithm on larger databases of anatomical models.


References

1. Fischl, B., Sereno, M.I., Tootell, R.B., Dale, A.M.: High-resolution inter-subject averaging and a coordinate system for the cortical surface. In: Human Brain Mapping. Volume 8. (1999) 272–84
2. Bakircioglu, M., Joshi, S., Miller, M.: Landmark matching on brain surfaces via large deformation diffeomorphisms on the sphere. In: SPIE Medical Imaging. Volume 3661. (1999) 710–15
3. Leow, A., Yu, C.L., Lee, S.J., Huang, S.C., Nicolson, R., Hayashi, K.M., Protas, H., Toga, A.W., Thompson, P.M.: Brain structural mapping using a novel hybrid implicit/explicit framework based on the level-set method. *NeuroImage* **24**(3) (2005) 910–27
4. Thompson, P.M., Giedd, J.N., Woods, R.P., MacDonald, D., Evans, A.C., Toga, A.W.: Growth patterns in the developing human brain detected using continuum-mechanical tensor mapping. *Nature* **404**(6774) (March 2000) 190–193
5. Davatzikos, C.: Spatial normalization of 3D brain images using deformable models. *J. Comp. Assisted Tomography* **20**(4) (1996) 656–65
6. Durrleman, S., Pennec, X., Troune, A., Thompson, P.M., Ayache, N.: Inferring brain variability from diffeomorphic deformations of currents: An integrative approach. *Medical Image Analysis* (2008) In Press.
7. Wang, Y., Chiang, M.C., Thompson, P.M.: Automated surface matching using mutual information applied to riemann surface structures. In: Med. Image Comp. Comput.-Assist. Intervention, Proceedings, Part II. (Oct. 2005) 666–674
8. Pitiot, A., Delingette, H., Toga, A.W., Thompson, P.M.: Learning object correspondences with the observed transport shape measure. In: IPMI. (2003) 25–37
9. Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J.: A minimum description length approach to statistical shape modeling. In: IEEE TMI. Volume 21. (2002) 525–37
10. Thodberg, H.H.: Minimum description length shape and appearance models. In: IPMI. (2003) 51–62
11. Wang, Y., Gu, X., Chan, T.F., Thompson, P.M., Yau, S.T.: Brain surface conformal parameterization with the Ricci flow. In: IEEE ISBI. (2007) 1312–1315
12. Hurdal, M.K., Stephenson, K.: Cortical cartography using the discrete conformal approach of circle packings. *NeuroImage* **23** (2004) S119–S128


13. Angenent, S., Haker, S., Tannenbaum, A., Kikinis, R.: Conformal geometry and brain flattening. *Med. Image Comput. Comput.-Assist. Intervention* (Sep. 1999) 271–278
14. Gu, X., Wang, Y., Chan, T.F., Thompson, P.M., Yau, S.T.: Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE TMI* **23**(8) (Aug. 2004) 949–958
15. Wang, Y., Lui, L.M., Gu, X., Hayashi, K.M., Chan, T.F., Toga, A.W., Thompson, P.M., Yau, S.T.: Brain surface conformal parameterization using Riemann surface structure. *IEEE TMI* **26**(6) (June 2007) 853–865
16. Commowick, O., Stefanescu, R., Fillard, P., Arsigny, V., Ayache, N., Pennec, X., g. Malandain: Incorporating statistical measures of anatomical variability in atlas-to-subject registration for conformal brain radiotherapy. In: *MICCAI. Volume II*, Palm Springs, CA, USA (2005) 927–934
17. Pennec, X., stefanescu, R., Arsigny, V., Fillard, P., Ayache, N.: Riemannian elasticity: A statistical regularization framework for non-linear registration. In: *MICCAI. Volume II*, Palm Springs, CA, USA (2005) 943–950
18. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors”. In: *Magn. Reson. Med. Volume 56*. (2006) 411–421
19. Lepore, N., Brun, C., Chou, Y.Y., Chiang, M.C., Dutton, R.A., Hayashi, K.M., Luders, E., Lopez, O.L., Aizenstein, H.J., Toga, A.W., Becker, J.T., Thompson, P.M.: Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors. *IEEE TMI* **27**(1) (Jan. 2008) 129–141
20. Thompson, P.M., Lee, A.D., Dutton, R.A., Geaga, J.A., Hayashi, K.M., Eckert, M.A., Bellugi, U., Galaburda, A.M., Korenberg, J.R., Mills, D.L., Toga, A.W., Reiss, A.L.: Abnormal cortical complexity and thickness profiles mapped in Williams syndrome. *J. Neurosciende* (2005) 4146–4158
21. Guggenheimer, H.W.: *Differential Geometry*. Dover Publications (1977)
22. Hamilton, R.S.: *The Ricci flow on surfaces. Mathematics and general relativity* (Santa Cruz, CA, 1986), *Contemp. Math. Amer.Math.Soc. Providence, RI* **71** (1988)
23. Chow, B.: The Ricci flow on the 2-sphere. *J. Differential Geom.* **33**(2) (1991) 325–334
24. Thurston, W.P.: *Geometry and Topology of Three-Manifolds. Princeton lecture notes* (1976)
25. Weitraub, S.H.: *Differential Forms: A Complement to Vector Calculus*. Academic Press (2007)
26. Chow, B., Luo, F.: Combinatorial Ricci flows on surfaces. *Journal Differential Geometry* **63**(1) (2003) 97–129




Group Difference between left and right cortical surfaces
for 21 control subjects (registered with 4 landmarks)



Group Difference between left and right cortical surfaces
for 21 control subjects (registered with 7 landmarks)




Group Difference between left and right cortical surfaces
for 21 WS (registered with 4 landmarks)




Group Difference between left and right cortical surfaces
for 21 WS (registered with 7 landmarks)


Fig. 4. Brain morphology study in 21 Williams Syndrome patients and 21 matched control subjects (intra group study).




Group Difference of left cortical surfaces between
21 controls and 21 WS (registered with 4 landmarks)



Group Difference of left cortical surfaces between
21 controls and 21 WS (registered with 7 landmarks)



Group Difference of right cortical surfaces between
21 controls and 21 WS (registered with 4 landmarks)



Group Difference of right cortical surfaces between
21 controls and 21 WS (registered with 7 landmarks)

Fig. 5. Brain morphology study in 21 Williams Syndrome patients and 21 matched control subjects (inter group study).

Multi-Atlas Tensor-Based Morphometry and its Application to a Genetic Study of 92 Twins

Natasha Leporé¹, Caroline Brun¹, Yi-Yu Chou¹, Agatha D. Lee¹, Marina Barysheva¹, Greig I. de Zubicaray², Matthew Meredith², Katie L. McMahon², Margaret J. Wright³, Arthur W. Toga¹, and Paul M. Thompson¹

¹ Laboratory of Neuro Imaging, UCLA, Los Angeles, CA 90095, USA

² Centre for Magnetic Resonance, University of Queensland, Brisbane, Australia

³ Genetic Epidemiology Lab, Queensland Institute of Medical Research, Brisbane, Australia

Abstract. Here we develop a multi-template analysis for tensor-based morphometry (TBM) which aims to reduce error from the registration step. In conventional TBM, a single template is nonlinearly registered to all images in the study, and the deformation field statistics are computed from the transformations. Using an MRI dataset from 23 monozygotic and 23 dizygotic twin pairs, we instead registered each individual twin image to 9 additional brain templates using a Riemannian fluid algorithm [3]. Average deformation tensors from multiple registrations were computed within each image, using a log-Euclidean framework [1]. To quantify improvements as the number of registration templates increased from 1 to 9, sequential t-tests assessed the significance of any error reduction, as each new template was added. For each number of templates, we also computed two tensor-derived metrics, and maps of the intraclass correlation of local volume differences, to evaluate any power advantages of multi-atlas TBM.

1 Introduction

Template selection is an important step in group analyses of brain MR images. In particular, in tensor-based morphometry (TBM), a set of images is non-linearly registered to a common reference image, and a statistical analysis is performed on the deformation tensors $S = \sqrt{J^T J}$, where J represents the Jacobian matrices derived from the deformation fields.

Detection power depends on several factors, and key among these is the quality of the non-linear registration, which also depends on the common target to which all images are mapped. Typically, nonlinear registration is performed either to one of the controls, or to an average of them [8, 7, 13]. In theory the construction of a mean anatomical atlas can reduce the bias induced by registering images to an individual control subject. In practice however, anatomical boundaries and image gradients are often blurrier in the average image, which may reduce the accuracy of the registration. In [5] for instance, the ICBM53 average brain was compared to a single individual template in a TBM study of

HIV/AIDS patients. Greater effect sizes per voxel were found with the single template when compared with those found using the ICBM53 brain as a registration target. Even so, individual variability may cause a bias in the registration when a single template is used, as the template may have a shape and intensity distribution that is closer to some subjects than to others. For instance, in [15], statistical power for TBM with single template registration was found to depend on which particular individual was selected as reference. Templates that were most *average*- in the sense of inducing the smallest deformation tensors when registered to other brains in the study- tended to generate more powerful statistics.

Our aim here was to design a new method that eliminates the dependence on individual variability, while retaining the sharp features associated with registration to a single target. One way to combine both of these requirements consists of moving the averaging step until after the non-linear registration. The gist of our averaging method consists of the following steps: starting from a set of brain MR images and a set of templates, we non-linearly register all images to all templates individually. We then compute the deformation tensors for each image as an average of those generated from the registration to individual templates. Other solutions to the template selection problem include 'targetless' normalization as in [22], [19], [21], [10], and the selection of the optimal individual target [9, 15].

In the standard version of TBM [19, 2], statistics are performed on the determinants of the Jacobian matrices, $\det J$, or equivalently on the determinants of the deformation tensors generated from the deformation field. In [14], multivariate statistics were computed instead on the full deformation tensors. As deformation tensors do not form a vector space under standard matrix addition and scalar multiplication, computations were performed in the log-Euclidean framework [1]. Since statistics are computed on the deformation tensors or a function of its components, for example the determinant, here we also perform a log-Euclidean averaging on those tensors.

Our analysis was performed on a dataset of MR images from 23 monozygotic (MZ) and 23 dizygotic (DZ) twin pairs, as well as 10 template brain MR images from identically scanned healthy subjects who did not belong to those pairs. We used the intraclass correlation of the $\det J$ as a statistic to characterize the influence of shared genes on local brain volume. Voxelwise permutation statistics were computed to assess the significance of the results.

2 Method

2.1 Data

3D T1-weighted images were acquired from 23 pairs of monozygotic twins (MZ; 20 males/26 females; 25.11.5SD years old) and 23 pairs of same-sex dizygotic twins (DZ; all same-sex pairs; 20 males/26 females; 23.52.2 years), as well as 10 individuals of comparable age, scanned identically. All MR images were collected

using a 4 Tesla Bruker Medspec whole body scanner (Bruker Medical, Ettingen, Germany) at the Center for Magnetic Resonance (University of Queensland, Australia). Three-dimensional T1-weighted images were acquired with an inversion recovery rapid gradient echo (MP-RAGE) sequence to resolve anatomy at high resolution. Acquisition parameters were as follows: inversion time (TI)/repetition time (TR)/echo time (TE) = 1500/2500/3.83 msec; flip angle = 15 degrees; slice thickness = 0.9 mm, with an acquisition matrix of 256 x 256 x 256. The study was approved by the Institutional Review Boards at the University of Queensland and at UCLA; all subjects gave informed consent.

Non-brain tissues were removed from all images using the Brain Surface Extraction tool (BSE) of BrainSuite [18]. The masked image was spatially normalized to the Colin27 standard brain by a 9-parameter (3 translations, 3 rotations, 3 scales) transformation, using the FLIRT software [11].

2.2 Template averaging

Each individual image I was non-linearly registered to each of 9 templates using a fluid (large-deformation) version of a Riemannian registration algorithm [3, 17, 6], which guarantees diffeomorphic mappings. In order to compute averages in a common space, all templates were also registered to a tenth template. The deformation fields resulting from the first registration step were concatenated with those of the second registration to obtain 9 sets of deformation fields in the common space for each image I . The deformation tensors $S^1(x, I)$ for the single template case were then computed at each voxel x .

In the log-Euclidean framework, the deformation tensors are projected to the tangent plane of the manifold on which they are defined, via their matrix logarithm. Tensor addition is performed in this space and the result is projected back to the original manifold. Thus, the n -templates average deformation tensor $S^n(x, I)$ at voxel x on image I is given by:

$$S^n(x, I) = \exp \frac{1}{n} \sum_{j=1}^n \log S_j^1(x, I) \quad (1)$$

2.3 Selecting the number of templates

Two measures were used to determine any improvements from increasing the number of templates. We describe both procedures in this section. Briefly, our first comparison method consists of comparing the total magnitude of deformation tensors before and after the addition of a new template. In effect, this amounts to determining how 'average' the effective template would have to be to generate the given set of deformation tensors. The second method we use begins with the assumption that the registration error diminishes with the number of templates used. Thus we wish to compare the distance between the deformation tensors obtained from a given number of templates to that which would have been found had we had access to an infinite number of templates. In practice, we will use values found with 9 templates as our gold standard.

In the log-Euclidean framework, the distance between two deformation tensors S_1 and S_2 is

$$d(S_1, S_2) = \|\log S_1 - \log S_2\|,$$

where $\|\cdot\|$ denotes a norm, and \log is the matrix logarithm. Following [1], here we use

$$d(S_1, S_2) = (\text{Trace}(\log S_1 - \log S_2)^2)^{1/2}. \quad (2)$$

A measure of the size E_i^n of the tensors $S^n(\cdot, i)$ for n templates integrated over image i is given in this framework by [13]:

$$E^n(i) = \int \|\log S^n(x, i)\|^2 d^3x = \int \text{Tr}(\log(S^n(x, i))^2) d^3x \quad (3)$$

For our first test, the $E^n(i)$ are computed for each image and a t -test is performed between the sets of $E^{n-1}(i)$'s and the $E^9(i)$'s to assess the significance of the results.

For the second test, we start with the following hypothesis

$$\log S^n(x, i) = \log S^\infty(x, i) + e^n(x, i) \quad (4)$$

where $e^n(\cdot, i)$ is the error in the logarithm of the deformation tensors in image i from using n templates, and $S^\infty(\cdot, i)$ is the hypothetical deformation tensor field that would be obtained from averaging with an infinite number of templates. Integrating over the image volume, we obtain

$$\|e^n\|_i = \int \text{Tr}(\log(S^n(x, i) - S^\infty(x, i))^2) d^3x \quad (5)$$

Here a t -test is again performed, in this case between the errors $\|e^{n+1}\|_i$ and $\|e^n\|_i$. In practice, as we do not have the value of S^∞ , we compare all tensors to the ones found using all 9 templates.

2.4 Twin statistics

The determinants $\det J$ were computed at each voxel, to assess local tissue volume differences between individuals, after global brain scale differences across subjects were discounted using 9-parameter registration. $\det J > 1$ and $\det J < 1$ respectively represent larger and smaller local volumes in the subject studied, with respect to the reference (template) image. We use the intraclass correlation (ICC) as a statistic to assess the influence of genes on these parameters. The ICC measures the correlation between unordered pairs and is defined as:

$$ICC = \sigma_b^2 / (\sigma_b^2 + \sigma_w^2). \quad (6)$$

Here σ_b^2 is the pooled variance between pairs, while σ_w^2 is the variance within pairs. The voxelwise intraclass correlation is computed over the whole brain volume, and the significance assessed by comparing the ICC values to a permutation distribution [16], for which we randomly reassign subject labels (5000 permutations).

3 Results

In **Fig. 1a**, $\sum_{i=1}^I E^n(i)$ is plotted against n . The p -values in Table 1 give the significance of the error reduction from adding each successive template, i.e., as n increases. For both the MZs and DZs, the averaging converges for 4 templates, in the sense that no statistically significant energy reduction is detected. The registration error summed over all images, $\sum_{i=1}^I \|e^n\|_i$, is shown in **Fig.1b**. P -values for MZs are below the 0.05 threshold for significance for the first 3 templates, while the DZs converge after the second averaging. In all cases, the size of the improvement decreases with an increasing number of templates.




Fig. 1. Left: Size of the deformation tensors for the MZ (red line) and the DZ (blue line) groups. The x-axis shows the number of templates n included in the averaging, while the y-axis represents $\sum_{i=1:I} E^n(i)$, the size of the deformation tensors summed over all images in either the MZ or DZ dataset. Here I is the total number of subjects in each of the two groups. For each value of n , a Students t -test was performed between E_i^n and E_i^9 to assess the significance of adding one more image to the sample, when compared to the result of using all 9 templates. The p -values derived from these tests are shown in **Table 1**. **Right:** Magnitude of the registration error for the MZ (red line) and the DZ (blue line) groups. The x-axis is the number of templates n , and the error $\sum_{i=1:I} \|e^n\|_i$ is plotted on the y-axis. A t -test was again performed for each successive value of n , this time between $\|e^n\|_i$ and $\|e^{n+1}\|_i$ to assess the significance of adding one more reference image. The p -values are shown in **Table 1**.

In order to verify our results in a TBM analysis, we computed the intra-class correlation for the MZs and the DZs for each number of templates. These analyses would be expected to show higher correlations in cases where less registration error was present, if other factors were equal. **Fig.2** shows the p -values that were found using the optimal number of templates, which we determined to be 4 from **Table 1**. The p -values indicate the statistical significance of the correlation ($p < 0.05$ are shown in red). As expected, the p -values are generally lower in the MZ twins (denoting higher intraclass correlations), as MZ twins share all of their genes, and regional brain volumes are known to be under genetic

no of templates	2	3	4	5	6	7	8	9
$E^n(i)$ p -values for the MZs	> 0.0001	0.053	0.014	0.72	0.23	0.42	0.64	0.85
$E^n(i)$ p -values for the DZs	> 0.0001	0.0065	0.0032	0.92	0.60	0.49	0.71	0.26
$\ e^n\ _i$ p -values for the MZs	> 0.0001	0.0027	> 0.0001	0.77	0.082	0.92	0.91	N/A
$\ e^n\ _i$ p -values for the DZs	0.0016	0.46	0.69	0.87	0.94	0.98	0.71	N/A

Table 1. The first two rows show the p -values from performing a t -test comparing the E_i^n 's to the $E_i^n + 1$ for all n 's. P -values from a t -test comparing the error $\|e^n\|_i$ to $\|e^n\|_i$.

control. This is indeed the case. We also computed the cumulative distribution functions of the p -values for maps derived using 1, 4 and 9 templates, for each of the two groups. For a null distribution, the cumulative distribution function is expected to fall along the $x = y$ line. Larger upward deviations from the diagonal are generally indicative of greater statistical power and larger effect sizes [14]. However, none of the distributions differed significantly from each other, indicating thus indicating that statistical power was not detectably different for the 3 maps using progressively higher number of templates. Despite this, any detected differences may have greater validity as the analysis removes the potential bias resulting from selecting an individual brain as a registration target, which may affect the accurate quantification of regional brain volumes.




Fig. 2. Maps of p -values for the intraclass correlation, shown here in logarithmic scales. **Left:** MZ twins. **Right:** DZ twins. Red p -values indicate ICCs significant at a level of $p = 0.05$.

4 Discussion

Here we presented a new multi-atlas version of Tensor-Based Morphometry. A log-Euclidean averaging procedure was performed on the deformation tensors generated from the multiple registrations, resulting in the reduction of registration error and improved statistical power. At least for the case presented here, the error was significantly reduced up to an averaging of 4 templates. Future studies will assess how much multi-template TBM boosts power relative to other influential factors, including the sample sizes used, scan quality or field strength, the regularization model and data fidelity term [5,3], and the tensor statistics used [14].

References

1. Arsigny V et al., *Log-Euclidean metrics for fast and simple calculus on diffusion tensors*, Mag. Res. in Med. 56, (2006) 411–421.
2. Ashburner J, *A fast diffeomorphic image registration algorithm*, Neuroimage 38, (2007) 95–113.
3. Brun C et al., *Comparison of standard and Riemannian elasticity for Tensor-Based Morphometry in HIV/AIDS*, MICCAI workshop on Statistical Registration: Pair-wise and Group-wise Alignment and Atlas Formation (2007).
4. Chou YY et al., *Automated Ventricular Mapping with Multi-Atlas Fluid Image Alignment Reveals Genetic Effects in Alzheimer's Disease*, accepted for publication in Neuroimage (2007).
5. Chiang MC et al., *3D pattern of brain atrophy in HIV/AIDS visualized using tensor-based morphometry*, Neuroimage 34, (2007) 44–60.
6. Christensen GE et al., *Deformable templates using large deformation kinematics*, IEEE-Trans. Imag. Process. 5, (1996) 1435–1447.
7. Guimond et al. *Average brain models: a convergence study*, Comp. Vis. and Im. Understanding 77, (1999) 192–210.
8. Kochunov P et al., *Regional spatial normalization: toward an optimal target*, J. Comp. Assist. Tomogr. 25, (2001) 805–816.
9. Kochunov P et al. *An optimized individual target brain in the Talairach coordinate system*, Neuroimage 17, (2003) 922–927.
10. Kochunov P et al., *Mapping structural differences of the corpus callosum in individuals with 18q deletions using targetless regional spatial normalization*, Hum. Brain Map. 24, (2005) 325–331.
11. Jenkinson M and Smith S, *A global optimisation method for robust affine registration of brain images*, Med. Imag. Anal. 5, (2001) 143–56.
12. Leow AD et al., *Statistical properties of Jacobian maps and inverse-consistent deformations in non-linear image registration*, IEEE-Trans. Med. Imag. 26, (2007) 822–832.
13. Leporé N et al., *Mean template for Tensor-Based Morphometry using deformation tensors*, MICCAI (2007) 826–833.
14. Leporé N et al., *Multivariate statistics of the Jacobian matrices in Tensor-Based Morphometry and their application to HIV/AIDS*, IEEE-Trans. Med. Imag., in press (2008).
15. Leporé N et al., *Best individual template selection from deformation tensor minimization*, ISBI (2008).

16. Nichols TE, Holmes AP, *Non parametric permutation tests for functional neuroimaging: a primer with examples*, Hum. Brain Mapp. 15, (2001) 1–25.
17. Pennec X et al., *Riemannian elasticity: A statistical regularization framework for non-linear registration*, MICCAI, (2005) 943–950.
18. Shattuck DW and Leahy RM, *BrainSuite: an automated cortical surface identification tool*, Med. Imag. Anal. 8, (2002) 129–141.
19. Studholme C et al., *Deformation tensor morphometry of semantic dementia with quantitative validation*, Neuroimage 21, (2004) 1387–1398.
20. Studholme C, Cardenas V, *A template free approach to volumetric spatial normalization of brain anatomy*, Patt. Recogn. Lett. 25, (2004) 1191–1202.
21. Twining CJ, *A unified information-theoretic approach to groupwise non-rigid registration and model building* MICCAI, (2005) 190–193.
22. Thompson PM et al., *Mathematical/Computational Challenges in Creating Population-Based Brain Atlases*, Hum. Brain Mapp. 9, (2000) 81–92.

Shape Registration with Spherical Cross Correlation

Boris Gutman^{1,2} {bgutman@ucla.edu}, Yalin Wang^{1,2}, Tony Chan¹, Paul M. Thompson², and Arthur W. Toga²

¹ UCLA Department of Mathematics

² UCLA Laboratory of Neuro Imaging

Abstract. We present a framework for shape alignment that generalizes several existing methods. We assume that the shape is a closed genus zero surface. Our framework requires a diffeomorphic surface mapping to the 2-sphere which preserves rotation. Our similarity measure is a global spherical cross-correlation function of surface-intrinsic scalar attributes, weighted by the cross-correlation of the parameterization distortion. The final similarity measure may be customized according to the surface-intrinsic scalar functions used in the application.

1 Introduction

Problems of shape alignment are ubiquitous in medical imaging. While many problem-specific solutions exist for particularly common cases (e.g. cortex, hippocampus) [1–3], high quality general shape alignment remains very much an open problem. Our framework generalizes some existing methods without assuming the existence of any landmarks or data-specific features. We assume that the shape is a closed genus zero surface. Our framework requires a diffeomorphic surface mapping to the 2-sphere which preserves rotation. Our similarity measure is a global spherical cross-correlation function of surface-intrinsic scalar attributes, weighted by the cross-correlation of the parameterization distortion, sometimes known as the conformal factor. The final similarity measure may be customized according to the surface-intrinsic scalar functions used in the application. Higher order scalar functions such as mean and Gaussian curvature may be used in conjunction with low order ones, like distance to mass center, to incorporate more localized shape information in addition to global measures.

In this study we have used a global conformal mapping as the spherical parameterization, and only the distance to mass center as intrinsic scalar shape measure. Using our method, we created atlases and registered shapes from a population of hippocampi.

2 Previous work

There is doubtless a galaxy of existing general shape alignment methods, and due to space limitations we will only mention those most prevalent and those closest

to the present work. Davies [4] has used an information-theoretic framework to formulate shape alignment as a minimum description length (MDL) problem. Here, each point is treated as an independent variable, while the cost function is aimed at reducing the "code length" of each shape's representation in this shape space. The beauty of this statistical approach lies in its ability to register multiple shapes simultaneously without the need to select a "target" shape. Of course, this is also its limitation, as for example when a known atlas exists and all data shapes are to be registered to it. The method requires a spherical diffeomorphism like ours.

Variants of the well-known ICP algorithm are another variety of recent developments in rigid shape alignment. Granger [5] introduced the EM-ICP method. ICP's proclivity for terminating at suboptimal local minima is greatly reduced by treating the problem as a general expectation-maximization problem. A multiscale approach is used: at coarser scales, the blurring factor is sufficient to give crude but correct alignment, which is improved upon in later refined stages, where the algorithm approaches the original ICP. Though quite robust, the method depends on the scale of blurring factor being set correctly to avoid local minima. Thus, avoiding them is not guaranteed. Our method, by contrast, performs a global search non-iteratively. Thus, a global maximum correlation is guaranteed irrespective of the shape's original orientation without the need to tune any parameters.

Much like our algorithm, some previous methods have used spherical harmonics for rigid shape alignment. Among them are the first order ellipsoid (FOE) method, popularized by Brechbuhler [6] and used extensively in medical imaging [7, 8] applications and SHREC, a recent variant of the ICP algorithm. Like ours, these methods make use of rotational properties of spherical harmonics.

FOE alignment uses the fact that a shape reconstructed from only the first order spherical harmonics forms an ellipsoid in object space. The method works well when the ellipsoid's three axes have distinct lengths, which largely depends on the shape itself and the degree to which the spherical parameterization preserves area. The method gives only a crude alignment and fails when two or more axes have similar lengths. Even with a proper ellipsoid, there is a symmetry problem.

SHREC [9] is another variation of the ICP. As in our case, the correspondence search is done iteratively on the sphere via Euler's rotation formula and icosahedral subdivision. The mapping satisfies our conditions, while the similarity measure is the RMSD. Since RMSD depends on the position of the object in space, the algorithm requires an initial pre-alignment in both spaces. Rigid Quaternion transform (RQT) is used to align shapes in \mathbb{R}^3 after each iteration of parametric alignment. The main limitation of this algorithm is that it is not guaranteed to converge to the optimal solution (i.e. a correspondence which, when applied to RQT, minimizes RMSD). This is because the optimization in parameter space depends on the object's position in native space. Thus, though the search is more global than in the original ICP, the parameter space search is still locally biased. The original ICP suffers the same problem.

The last algorithm uses a brute force correlation: it computes a cost function anew for each rotation. To mitigate the cost of this, a hierarchical approach is used. Instead, we reduce parameter space alignment to a global refined search via the FFT. Our numerical scheme separates the effects on computation time of the level of detail used for alignment and the number of rotation samples. In SHREC, these two are tied together owing to their brute-force nature. This means that we can refine rotation space tessellation while maintaining the same level of surface detail without significantly affecting computation time. Our use of orientation-invariant shape attributes in conjunction with scale invariant cross-correlation makes our approach completely independent of changes in object position and size. SHREC, by contrast, requires volume normalization as a pre-processing step.

3 Mathematical preliminaries

3.1 Spherical harmonics

Spherical harmonics are functions $f : \mathbb{S}^2 \rightarrow \mathbb{C}$ which are simultaneously eigenfunctions of the Laplace-Beltrami and the angular momentum operators; they are expressed explicitly as

$$Y_l^m(\theta, \phi) = \sqrt{\frac{(2l+1)(l-m)!}{4(l+m)!}} P_l^m(\cos\theta) e^{im\phi} \quad (1)$$

for degree and order $m, l \in \mathbb{Z}$, $|m| \leq l$, where $P_l^m(x)$ is the associated Legendre polynomial. Spherical harmonics form a countable orthonormal basis for square-integrable functions on the sphere. A projection of a function $f \in L_2(\mathbb{S}^2)$ onto this basis yields the SPH coefficients

$$\hat{f}(l, m) = \langle f, Y_l^m \rangle \quad (2)$$

where $\langle f, g \rangle$ is the usual L_2 inner product.

A key property of spherical harmonics is their behavior under a shift on the sphere. Given an element of the rotation group $R \in SO(3)$, a rotated spherical harmonic is expressed as

$$Y_l^m(\omega) = \sum_{n=-l}^l Y_l^n(R^{-1}\omega) D_{m,n}^l(R) \quad (3)$$

where

$$D_{m,n}^l(R) = e^{-i(m\alpha+n\gamma)} d_{m,n}^l(\beta), \quad (4)$$

α, β, γ are the Euler angles of R and $d_{m,n}^l$ are irreducible representations of $SO(3)$ [10],

$$d_{m,n}^l(\beta) = \sum_t (-1)^t \times \frac{\sqrt{(l+n)!(l-n)!(l+m)!(l-m)!}}{(l+n-t)!(l-m-t)!(t+m-n)!t!} \quad (5)$$

$$\times \left(\cos \frac{\beta}{2}\right)^{2l+n-m-2t} \left(\sin \frac{\beta}{2}\right)^{2t+m-n}.$$

In particular, (3) implies that

$$f(\omega) = g(R^{-1}\omega) \implies \widehat{f(l, m)} = \sum_{n=-l}^l \widehat{g(l, n)} D_{m,n}^l(R). \quad (6)$$

3.2 Discrete spherical cross correlation

Given two functions $f, g \in L_2(\mathbb{S}^2)$, their spherical cross-correlation is defined as

$$C_{f,g}(R) = \int_{\mathbb{S}^2} f(\omega) g(R^{-1}\omega) d\omega. \quad (7)$$

In the special case where $f(\omega) = g(R'^{-1}\omega)$, $C_{f,g}(R)$ is maximized when $R = R'$, assuming that f is not spherically symmetric. For bandlimited functions, i.e. for those functions whose spherical harmonic coefficients vanish for all $l \geq B$ for some bandwidth B , the correlation becomes

$$C_{f,g}(R) = \sum_{l=0}^{B-1} \sum_{m=-l}^l \widehat{f}(l, m) \overline{\widehat{g}(l, m) \Lambda(R)}. \quad (8)$$

Here, $\Lambda(R)$ is the operator associated with the rotation matrix. The expression for shifted spherical harmonic coefficients (6) implies that

$$C_{f,g}(R) = \sum_{l,m,n} \widehat{f}(l, m) \overline{\widehat{g}(l, n)} D_{m,n}^l(R). \quad (9)$$

This expression forms the basis of our similarity measure.

3.3 Fast cross correlation via FFT

The material presented so far has been used in the prior works we mentioned. Now, we present a simple lemma which leads to a great speed up in computing the correlation (9). It suffices to make the observation that any rotation $R(\alpha, \beta, \gamma)$ may be expressed as a product of two rotations:

$$R(\alpha, \beta, \gamma) = R_1(\alpha + \pi/2, \pi/2, 0) R_2(\beta + \pi, \pi/2, \gamma + \pi/2). \quad (10)$$

Now using the fact that

$$D_{m,n}^l(R_1 \cdot R_2) = \sum_k D_{m,k}^l(R_1) D_{k,n}^l(R_2), \quad (11)$$




Fig. 1. Two randomly selected hippo surfaces. The red circle is homologous to the north pole on the sphere; the blue circle, to the south pole; and the blue line, to $\phi = 0$. The initial spherical mappings do not align these two very well.

we substitute (10) into (9) to obtain

$$\begin{aligned}
 C_{f,g}(R(\alpha, \beta, \gamma)) &= \sum_{l,m,n,k} \hat{f}(l, m) \overline{\hat{g}(l, n)} \\
 &\quad \times D_{m,k}^l(\alpha + \pi/2, \pi/2, 0) D_{k,n}^l(\beta + \pi, \pi/2, \gamma + \pi/2) \\
 &= \sum_{l,m,n,k} \hat{f}(l, m) \overline{\hat{g}(l, n)} d_{m,k}^l(\pi/2) d_{k,n}^l(\pi/2) \times e^{i(m(\alpha+\pi/2)+k(\beta+\pi)+n(\gamma+\pi/2))} \\
 &= IFFT\{\sum_l \hat{f}(l, m) \overline{\hat{g}(l, n)} d_{m,k}^l(\pi/2) d_{k,n}^l(\pi/2)\}(\alpha + \pi/2, \beta + \pi, \gamma + \pi/2).
 \end{aligned} \tag{12}$$

This simple result has been shown elsewhere [11], but to the best of the authors' knowledge this is the first time it has been used for shape registration.

4 Shape registration with cross correlation

4.1 Similarity measure

Given a 2-manifold $M \subset \mathbb{R}^3$, a diffeomorphic spherical parameterization $f : \mathbb{S}^2 \mapsto M$ and a family of rotation-invariant shape attributes $s_i : M \mapsto \mathbb{R}$, $0 < i \leq N$, let

$$S_i = s_i \circ f. \tag{13}$$

Then, given two manifolds M_1, M_2 and their corresponding shape attribute maps $\{S_{1,i}\}_{i=0}^N, \{S_{2,i}\}_{i=1}^N$, we define our shape similarity measure as




Fig. 2. Top surface from figure 1 is correlated to the bottom one, and its signal is shifted on the sphere without changing the shape's position in object space. The homologous points now appear to be in good correspondence between the surfaces.

$$\mathcal{S}_{M_1, M_2}(R) = C_{\lambda_1, \lambda_2}(R) \sum_{i=1}^N \kappa_i C_{S_{1,i}, S_{2,i}}(R), \quad (14)$$

where κ_i are user-defined shape attribute weights, and λ_1, λ_2 are spherical maps of the conformal factor of each manifold. These last two are used to mitigate the fact that scalar functions which appear similar on the sphere may in fact represent vastly differently-sized regions on the original surfaces due to varying area distortion of the spherical map.

Because we recover shifts in object space with shifts on \mathbb{S}^2 , we require that the spherical parameterization preserve rotation in the following sense. Suppose $M_2 = R \circ M_1$, and $f_1, f_2 : \mathbb{S}^2 \mapsto M_1, M_2$ are spherical maps. Then

$$f_2(\omega) = R \circ f_1(R^{-1}\omega). \quad (15)$$

Many existing parameterizations satisfy this requirement, e.g. [6, 12].

4.2 Previous methods as special cases

SHREC and FOE are special cases of our method. FOE simply takes the conformal factor λ to be constant and uses spherical harmonics up to order one only. This is equivalent to setting the bandwidth B to 2. The single shape attribute used is the Euclidian distance to the surface average value. This is roughly the same as distance to mass center, especially for area-preserving spherical maps with which FOE is typically used.

SHREC minimizes RMSD, which can be reduced to spherical cross-correlation since $4\pi \text{RMSD}_{M_1, M_2}^2(R) =$

$$\begin{aligned} & \sum_{i \in \{x, y, z\}} \int_{\mathbb{S}^2} \|S_{1,i} - \Lambda(R)S_{2,i}\|^2 \\ &= \sum_{i \in \{x, y, z\}} \int_{\mathbb{S}^2} [S_{1,i}^2 + (\Lambda(R)S_{2,i})^2 - 2S_{1,i}\Lambda(R)S_{2,i}] \end{aligned} \quad (16)$$

$$= \sum_{i \in \{x,y,z\}} \|S_{1,i}\|_2^2 + \|S_{2,i}\|_2^2 - 2C_{S_{1,i},S_{2,i}}(R)$$




Fig. 3. Top surface from figure 1 is initially aligned using FOE (top row) and subsequently aligned using a modified version of SHREC (bottom row). Though this is slightly different from the original SHREC, in principle the two algorithms are the same. Here, we see the local minimum problem suffered by SHREC, typical of an ICP-type algorithm: there is very little change in alignment after initialization.

Thus, minimizing RMSD is equivalent to maximizing the correlation of the two shapes' spatial coordinates. These are, of course, not quite the scalar shape attributes we intend to use in our similarity measure. SHREC's dependence on iterative RQT refinement for correspondence optimization makes it less robust. This is the price of using orientation-dependent features. Further, it is not clear whether the correspondence which, when applied to RQT, minimizes RMSD is truly the best correspondence. One can conceive two shapes with some patches quite similar and other very different. One may then like to align the two objects to get the best correspondence between the similar patches without regard to the different ones. In such a case, cross correlation of invariant features will achieve a better alignment. Still, SHREC could be made faster with the use of FFT-based correlation rather than a brute-force approach.

SHREC requires $O(B^3 N_R)$ operations, where B is again the bandwidth and N_R the number of rotation samples. This is because recomputing $D_{m,n}^l(\alpha, \beta, \gamma)$ and the corresponding shifted spherical harmonic coefficients requires $O(B^3)$ operations. Our method requires $O(B^4 + N_R \log(N_R))$. N_R is roughly also of the order B^3 even with hierarchical sampling; this means we have effectively reduced the order of operations from $O(B^6)$ to $O(B^4)$. This allows us to sample rotation space more finely without a significant change in execution time. It also makes using higher order coefficients and hence greater level of detail for alignment purposes feasible. We see this in experiments below.




Fig. 4. Average of two shapes from figure 1 before correlation.




Fig. 5. Average of two shapes from figure 1 after correlation.

5 Results

We used a population of 45 right hippocampal surfaces extracted from healthy elderly subjects. Our spherical parameterization was the global conformal map of Gu. et al. [12]. Spherical harmonics were computed with the spherical FFT of [13] and cross correlation was computed with the help of the fftw library [14]. $SO(3)$ was sampled at $200 \times 100 \times 200$ samples, which yielded an angle frequency of ~ 1.8 degrees. A bandwidth of 64 was used throughout the testing. Only the distance to mass center was used as an invariant shape attribute for both populations.

As a preliminary experiment, we applied the cross correlation algorithm to a pair of hippocampal surfaces shown in figure 1. Here we see that the initial spherical maps do not provide a very good correspondence. Figure 2 shows the result of a spherical shift based on cross correlation. Figures 4 and 5 show two point-wise averages of the shapes, before and after cross correlation. The improvement is obvious. For $B = 64$, the average running time for above experiment was 44.6 ± 3.6 seconds on a Gateway 7426GX Laptop with a 2.41 GHz AMD processor and 1 GB RAM, tested with 45 hippocampal surfaces. This includes computation of spherical harmonics, rotation matrices $d_{m,n}^l$ and correlation. Shen [9] reports an average of 23.5 seconds running time on a common laptop for SHREC, while using $B = 12$. One would expect SHREC to take $(64/12)^3$ times longer for our bandwidth (see above), or on the order of 60 minutes. As a preliminary comparison, we implemented a modified version of SHREC. The only difference with the original is that at each iteration, our parameter alignment was initially done with cross correlation as described above, and subsequently refined according to the scheme outlined in [9]. Again, a bandwidth of 64 was used. Due to time limitations, we could not run this program on our whole dataset. We only present the results of one subject in figure 3. Here, we can see that the initial alignment

determines the final result to a great degree. This example converged after only 4 iterations. Execution time was close to 7 minutes. Since the execution time reported in [9] was for a MATLAB implementation, while ours is in C++, and because we use a fast cross-correlation, this time is significantly lower than the 60 minute estimate.

To test the effect of higher coefficients on correlation quality, we increased the bandwidth to 128, while keeping the same rotation tessellation and correlated 6 of the 45 subjects in our HP population. Running time increased to 404 ± 10.4 seconds, while the shape distance to the target hippocampus decreased on average only 3.9 ± 9.4 %. To test for the effect of rotation sampling frequency, we also decreased N_R to $100 \times 50 \times 100$ samples, while keeping the bandwidth at 64. Running time was 38.2 ± 3.5 seconds, an insignificant improvement in speed. In another experiment, we limited our bandwidth to 16 and reran cross correlation based alignment while keeping angle resolution at $200 \times 100 \times 200$. Execution time was reduced to about 3 seconds.




Fig. 6. Average of 45 right hippocampi with cross correlation, bandwidth = 64.




Fig. 7. Average of 45 right hippocampi with FOE.

We constructed hippocampal averages using FOE and two versions of cross correlation ($B = 16$ and $B = 64$) and compared results. First, a shape was selected, and all remaining shapes registered to it using each of the methods. Then, the shapes were averaged, normalized for volume (after registration, only




Fig. 8. Average of 45 right hippocampi with cross correlation, bandwidth = 16.

for a fair distance comparison) and a rigid quaternion transform applied to each shape to align it both to the target subject and to the volume-normalized average shape. Our shape distance was defined as vertex-wise distance between the surfaces, weighted by product of the sum of areas of adjacent triangles in each mesh. The results are in table 1. The table does not show a notable fact: distance to the target subject was improved by our method for every subject compared to FOE. Minimal improvement was 8%, and maximal 117%. All but one subject registered with $B = 64$ had superior alignment to the result of using $B = 16$. Compared to FOE, $B = 16$ reduced shape distance for all but 5 subjects. Table 2 shows a summary of intra-subject differences by registration method. Using a bandwidth of 16 gives a significant improvement compared to FOE, but the results are still much improved by using a bandwidth of 64. Figures 6 - 8 illustrate the hippocampal averages achieved with the three methods. Note that these averages were computed without spatially aligning the subjects to the target. Doing so would have likely given a more detailed shape.

Method	Distance to Target	Distance to Average
FOE	4.14 ± 1.33	2.89 ± 1.05
Cross cor.	2.64 ± 0.92	2.36 ± 0.75
$B = 16$	3.38 ± 1.06	2.33 ± 0.85

Table 1. Weighted point-wise distance of 45 (44) hippocampal surfaces by registration method.

We should note here that the poor result of the FOE average is due in part to the large area distortion of the conformal map. The tail of the hippocampus is mapped to such a small region on \mathbb{S}^2 that it is very hard to match well. We see in figure 7 that the tail suffers the most. The resulting spherical harmonic representation contains much redundant information; hence, the area distortion partially cancels out the benefits of using high order coefficients. The mapping used in [6, 9] is by contrast area-preserving, and hence better suited for alignment. The results of [6, 8] look closer to the one achieved here with cross correlation. This is, however, indicative of the potential our method has when applied to area-

preserving spherical maps, and we intend to experiment with this idea in the immediate future.

common shape	FOE vs. B = 64	FOE vs. B = 16	B = 16 vs. B = 64
target	61 % \pm 26 %	25 % \pm 23 %	30 % \pm 18 %
average	24 % \pm 29 %	26 % \pm 22 %	-6 % \pm 29 %

Table 2. Intra-subject improvement by registration method, in percentage of the second method’s result.

6 Conclusion

We have presented a framework for shape alignment which generalizes several existing methods. Our method is robust, fast and allows for use of greater detail in alignment than was possible before. The correspondence search is performed globally and no pre-alignment is required; thus, the result and computation time are independent of the shape’s size and initial orientation. Reaching a global maximum is always guaranteed. Our method can be tailored to suit a particular application by selecting the appropriate shape features for a particular data type. We intend to experiment with various shape attributes, apply area-preserving spherical maps to our method and extend the technique to automated patch selection and matching. Lastly, all software used here is available through the LONI Pipeline environment. Please contact the authors for more details.

References

1. Shi, Y., Thompson, P., et al.: Direct mapping of hippocampal surfaces with intrinsic shape context. *Neuroimage* **37**(3) (2007) 792–807
2. Tosun, D., Prince, J.: Cortical surface alignment using geometry driven multispectral optical flow. In: *IPMI 2005*. (2005)
3. Fischl, B., et al.: High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping* **8** (1999) 272–284
4. Davies, R., et al.: A minimum description length approach to statistical shape modeling. *IEEE Transactions on Medical Imaging* **21**(5) (2002) 525–537
5. Granger, S., Pennec, X.: Multi-scale em-icp: A fast and robust approach for surface registration. In: *Lecture Notes in Computer Science*. (2002)
6. Brechbuhler, C., Gerig, G., Kubler, O.: Parameterization of closed surfaces for 3d shape description. *Computer Vision and Image Understanding* **61**(2) (1995) 154–170
7. Gerig, G., Styner, M.: Shape versus size: Improved understanding of the morphology of brain structures. In: *MICCAI 2001*. (2001)
8. Gerig, G., Styner, M., et al.: Shape analysis of brain ventricles using spharm. In: *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MM-BIA’01)*. (2001)

9. Shen, L., Huang, H., Makedon, F., Saykin, A.: Efficient registration of 3d spharm surfaces. In: IEEE 4th Canadian Conference on Computer and Robot Vision. (2007)
10. Talman, J.: Special functions: A group-theoretic approach. W.A.Benjamin, Inc. (1968)
11. Sorgi, L., Daniilidis, K.: Template gradient matching in spherical images. Image Processing: Algorithms and Systems III. Proceedings of the SPIE **5298** (2004) 88–98
12. Gu, X., Wang, Y., et al.: Genus zero surface conformal mapping and its application to brain surface mapping. IEEE Transactions on Medical Imaging **23**(8) (2004) 949
13. Healy, D., Rockmore, D., et al.: Ffts for the 2-sphere-improvements and variations. J. Fourier Anal. Applicat. **9**(4) (2003) 341–385
14. Frigo, M.: A fast fourier transform compiler. In: ACM SIGPLAN Conference on Programming Language Design and Implementation. (1999)

A Forward Model to Build Unbiased Atlases from Curves and Surfaces

Stanley Durrleman^{1,2}, Xavier Pennec¹, Alain Trounev², and Nicholas Ayache¹

¹ INRIA - Asclepios Team-Project, Sophia Antipolis, France

² Centre de Mathématiques et Leurs Applications (CMLA), ENS-Cachan, France

Abstract. Building an atlas from a set of anatomical data relies on (1) the construction of a “mean anatomy” (called template or prototype) and (2) the estimation of the variations of this template within the population. To avoid biases introduced by separate processing, we *jointly* estimate the template and its deformation, based on a consistent statistical model. We use here a forward model that considers data as noisy deformations of an unknown template. This differs from backward schemes which estimate a template by pulling back data into a common reference frame. Once the atlas is built, the likelihood of a new observation depends on the Jacobian of the deformations in the backward setting, whereas it is directly taken into account while building the atlas in the forward scheme. As a result, a specific numerical scheme is required to build atlases. The feasibility of the approach is shown by building atlases from 34 sets of 70 sulcal lines and 32 sets of 10 deep brain structures.

1 Forward vs. Backward Models for Template Estimation

In the medical imaging field, atlases are useful to drive the personalization of generic models of the anatomy, to analyze the variability of an organ, to characterize and measure anatomical differences between groups, etc. Many frameworks have been proposed to build atlases from large database of medical images [1–4], much fewer were proposed for anatomical curves or surfaces [5, 6]. In any case, the underlying idea remains the same: one estimates a “mean anatomy” (called template) and one learns how this mean model deforms within a given population. The most widely used method in medical imaging is based on a *backward* model that deforms every observations back to a common reference frame (See Fig.1). However, we prefer here to base our statistical estimation on a *forward* model, as pioneered in [7, 8], which considers the observations (T_i) as noisy deformations (ϕ_i) of an unknown template (\bar{T}). Formally, the *forward* model can be written as:

$$T_i = \phi_i \cdot \bar{T} + \varepsilon_i \quad (1.1)$$

whereas the *backward* model is:

$$\phi_i \cdot T_i = \bar{T} + \varepsilon_i \quad \Longleftrightarrow \quad T_i = \phi_i^{-1} \bar{T} + \phi_i^{-1} \varepsilon_i \quad (1.2)$$




Fig. 1. In the forward scheme, the physical observations (\mathcal{O}_i) are seen as noisy deformation (ϕ_i) of unknown template ($\bar{\mathcal{O}}$). In the backward scheme, the template is an average of deformed observations. In the forward scheme the noise is removed from the observations whereas it is pulled back in the common frame with the backward scheme.

The backward model considers either (Eq.1.2-left) that the template \bar{T} is noisy and the observations T_i free of noise or (Eq.1.2-right) that the noise added to the observations ($\phi_i^{-1}\varepsilon_i$) **depends** on the observations via an unknown deformation. By contrast, the forward model (Eq.1.1) considers that the template is not blurry, as an “ideal” object, and an independent and **identically distributed** noise ε_i is added to every observations. This models more accurately the physical acquisitions, whereas the backward model relies on less realistic assumptions.

The observations T_i are given as discrete sampled objects. The template \bar{T} models an average “ideal” biological material and it is therefore supposed to be continuous. Since in the backward model sampled observations deform to a continuous template, an extrinsic **interpolation** scheme is required. By contrast, in the forward setting, the deformed template needs only to be **sampled** to be compared to the observations. This does not only reproduce more accurately the real physical acquisition process, but also depends on less arbitrary assumptions.

Assume now that we can define probabilities on objects T (images, curves, surfaces, etc.) and on deformations ϕ . The statistical estimation of an atlas would require at least to compute $p(\bar{T}|T_i)$, the probability of having the template given a training database of T_i . Once the atlas is built, one would like to know how a new observation T_k is compared to the learnt variability model: one needs to compute the likelihood of this observation given the template $p(T_k|\bar{T})$. Because ϕ_i acts differently in Eq.1.1 and in Eq.1.2, the computational cost of these two steps varies significantly. In the backward scheme, computing $p(\bar{T}|T_i)$ should be simpler than computing $p(T_i|\bar{T})$ which depends on the Jacobian of ϕ_i . It is exactly the reverse for the forward scheme. Since it is better to spend more time to build the atlas (which is done once for all) and to keep simple the test of any new available data, the forward model seems better suited even from a computational point of view.

Finally, the forward model is also better understood from a theoretical point of view. For instance, the convergence of the Maximum A Posteriori (MAP) template estimation, when the number of available observations is growing, is proved for images and small deformations [7]. Such proofs for the backward model seem currently out of reach.

For all these reasons, we base here our statistical estimations on the forward model. We show in this paper how the atlas building step, which is the most critical step in this paradigm, is possible in case of curves and surfaces. Compared to images, dealing with shapes requires specific numerical scheme. We take advantage here of a sparse deconvolution scheme we introduced recently [9, 10].

The paper is organized as follows. A general non parametric framework for shape statistics is introduced in Section 2. In Section 3 we detail the optimization procedure to estimate jointly a template and its deformations to the shapes. A sparse deconvolution method is presented to effectively compute the gradient descent. In Section 4, we show templates computed from large sets of sulcal lines and sets of meshes of sub-cortical structures.

2 Non-parametric Representation of Shapes as Currents

As emphasized in [11, 12, 5], a current is a convenient way to model geometrical shapes such as curves and surfaces. The idea is to characterize shapes via vector fields, which are used to probe them. A surface S is characterized by the flux of any vector field ω through it:

$$S(\omega) = \int_S \omega(x) \cdot (u \times v)(x) d\sigma(x) \quad (2.1)$$

where $u \times v$ is the normal of the surface, (u, v) an orthonormal basis of its tangent plane) and $d\sigma$ the Lebesgue measure on the surface. Similarly, a curve L is characterized by the path integral of any vector field ω along it:

$$L(\omega) = \int_L \omega(x) \cdot \tau(x) dx \quad (2.2)$$

where τ is the tangent of the curve. More generally, we define a current T as a linear continuous mapping from a set of *test* vector fields W to \mathbb{R} . This framework enables to define addition, deformation, Gaussian noise on shapes, to measure a distance between shapes, without assuming point correspondences between objects. We recall here some properties of currents to give a rigorous sense of Eq.1.1 in case of shapes. We refer the reader to [11, 13, 9] for more details.

As mappings from W to \mathbb{R} , the currents build a **vector space**, denoted W^* : $(T_1 + T_2)(\omega) = T_1(\omega) + T_2(\omega)$ and $(\lambda.T)(\omega) = \lambda.T(\omega)$. For surfaces, this means that the flux through two surfaces is the sum of the flux through each surface: the addition is equivalent to the union of surfaces. Scaling a surface means scaling the power of the flux through the surface.

Suppose now that we can provide the *test* space W with a norm that measures the regularity of the vector fields. We can define then a **norm of a current** T as the maximum flux through T of any regular vector fields (i.e. $\|\omega\|_W \leq 1$):

$$\|T\|_{W^*}^2 = \text{Sup}_{\|\omega\|_W \leq 1} |T(\omega)| \quad (2.3)$$

The distance between two surfaces $\|S - S'\|_{W^*}$ is therefore obtained for the regular vector field that best separates the two surfaces, in the sense that the

difference of the flux through the two surfaces is the largest possible. *This distance between shapes does not depend on how shapes are parametrized and does not assume point correspondences between shapes.*

For computational purposes, we suppose, from now onwards, that W is a reproducible kernel Hilbert space (r.k.h.s) with kernel K [14]. In this setting, the space of currents is the **dense span** of all Dirac delta currents δ_x^α , which are defined by $\delta_x^\alpha(\omega) = \langle \omega(x), \alpha \rangle_{\mathbb{R}^3}$ for any $\omega \in W$. A Dirac current may be seen as a tangent (resp. normal) α entirely concentrated at point x . Although a curve (resp. a surface) has an infinite number of tangents (resp. normals), a polygonal lines (resp. a mesh) may be approximated in the space of currents by a finite sum $\sum_k \delta_{x_k}^{\alpha_k}$ where x_k is the center of the segment (resp. center of the mesh cell) and α_k the tangent of the line (resp. the normal of the surface) at x_k .

Moreover, based on the Riesz theorem, we can show that for any current T , the vector field (within Eq. 2.3) that achieves $\sup_{\|\omega\|_W \leq 1} |T(\omega)|$ exists and is unique. We denote this vector field $\mathcal{L}_W^{-1}(T)$ and call it the **dual representation** of T . \mathcal{L}_W is isometric: it provides W^* with an **inner product**:

$$\|T\|_{W^*}^2 = \langle T, T \rangle_{W^*} = \langle \mathcal{L}_W^{-1}(T), \mathcal{L}_W^{-1}(T) \rangle_W \quad (2.4)$$

On Diracs we have:

$$\mathcal{L}_W^{-1}(\delta_x^\alpha)(y) = K(y, x)\alpha \quad (2.5)$$

This equation shows that K is the Green kernel of \mathcal{L}_W . Applying Eq. 2.4 and Eq. 2.5 leads to:

$$\langle \delta_x^\alpha, \delta_y^\beta \rangle_{W^*} = \alpha^t K(x, y) \beta \quad (2.6)$$

This gives by linearity *explicit* and *easily tractable* formula to compute the inner product (and then the distance) between two shapes $T = \sum_i \delta_{x_i}^{\alpha_i}$ and $T' = \sum_j \delta_{y_j}^{\beta_j}$: $\langle T, T' \rangle_{W^*} = \sum_i \sum_j \alpha_i^t K(x_i, y_j) \beta_j$. The mapping \mathcal{L}_W^{-1} is a simple convolution operator, as illustrated in Fig.2-a. On the contrary, finding the current T whose associated vector field γ is given ($T = \mathcal{L}_W(\gamma)$) is an ill-posed deconvolution problem that requires specific numerical scheme as the one we will present in Section 3.3.

The kernel K enables also to define an associated **Gaussian noise** on currents ε . In infinite dimension, ε maps every currents to a Gaussian variable such that $\text{Cov}(\varepsilon(\delta_x^\alpha), \varepsilon(\delta_y^\beta)) = \alpha^t K(x, y) \beta$. If we set a grid $\Lambda = \{x_p\}_{p=1\dots N}$, the span of $(\delta_{x_p})_{p \in \Lambda}$ defines a finite dimensional space of currents. The corresponding ε is of the form $\sum_{p \in \Lambda} \delta_{x_p}^{\alpha_p}$ where (α_i) are centered Gaussian variables with covariance matrix \mathcal{K}^{-1} where $\mathcal{K} = (K(x_i, x_j))_{i,j \in \Lambda}$. In this last case, ε has a probability density function (pdf) which is proportionnal to $\exp(-\|\varepsilon\|_{W^*}^2)$. For our applications we will choose a Gaussian kernel: $K(x, y) = \exp(-\|x - y\|^2 / \lambda_W^2) \text{Id}$.

To make sense of the model Eq.1.1, one must still specify how this modelling based on currents may be coupled with a deformation framework. If ϕ is diffeomorphism and S a surface, the flux of ω through the deformed shape $\phi(S)$, denoted for general currents $(\phi_\# \star S)(\omega)$, is equals to the flux through S of the pulled-back vector field $\phi^\# \star \omega$ which is given by the change of variables

formula within the flux integral. This enables to define a general **push-forward action of a diffeomorphism** on any currents. This action replaces for curves and surfaces the usual action on images: $(\phi \star I)(x) = I(\phi^{-1}(x))$. This is here slightly more complex since we do not transport points but tangents or normals (differential 1 and 2-forms to be even more precise). On the basis elements, the push-forward action gives:

$$\phi_{\#} \star \delta_x^\alpha = \delta_{\phi(x)}^{d_x \phi(\alpha)} \quad (2.7)$$

if α is a tangent of a curve and

$$\phi_{\#} \star \delta_x^{u \times v} = \delta_{\phi(x)}^{d_x \phi(u) \times d_x \phi(v)} \quad (2.8)$$

if $u \times v$ is the normal of a surface. One notices that $d_x \phi(u) \times d_x \phi(v) = |d_x \phi| d_x \phi^{-t}(u \times v)$.

In the following, we restrict the deformations ϕ to belong to the group of diffeomorphisms set up in [15, 16]: the diffeomorphisms are obtained by integration of a time-varying vector field v_t : $\partial_t \phi_t = v_t \circ \phi_t$. The geodesic flows $(\phi_t^{v_t})_{t \in [0,1]}$ are completely determined by the initial vector speed v_0 which belongs to a r.k.h.s. V . The final diffeomorphism is denoted ϕ^{v_0} . How to find such a diffeomorphism that best matches two sets of currents is explained in detail in [11, 13].

3 Joint Estimation of Template and Deformations

3.1 A Heuristic Maximum A Posteriori in Infinite Dimension

From a Bayesian point of view, in Eq.1.1 ($T_i = \phi_{i,\#} \star \bar{T} + \varepsilon_i$), T_i are the observations, \bar{T} is unknown, ϕ_i are hidden variables and ε_i independent and identically distributed Gaussian noise with known variance. Suppose now that we can define Gaussian probability density functions (pdf) on the space of Currents W^* and on the space of initial vector fields V : $p_\varepsilon(\varepsilon) = C_\varepsilon \exp(-\|\varepsilon\|_{W^*}^2 / \sigma_W^2)$ and $p_\phi(v) = C_\phi \exp(-\|v\|_V^2 / \sigma_V^2)$. In that case, a Maximum A Posteriori (MAP) estimation for independent observations maximizes $\max_{\bar{T}} \prod_i p(T_i | \bar{T})$. Formally,

$$p(T_i | \bar{T}) = \int p_\varepsilon(T_i | \bar{T}, v_0^i) p_\phi(v_0^i) dv_0 = \int p_\varepsilon(T_i - \phi_i^{v_0^i} \bar{T}) p_\phi(v_0^i) dv_0$$

Since the term within the integral depends on v_0 by a geodesic shooting of diffeomorphisms, there are no closed forms for this likelihood. A usual approximation consists in replacing the integral by the maximum of the distribution within the integral (i.e. its first mode). This leads to: $p(T_i | \bar{T}) \sim \max_{v_0^i} p_\varepsilon(T_i - \phi_i^{v_0^i} \bar{T}) p_\phi(v_0^i)$ which finally gives:

$$(\bar{T}, \phi^{v_0^i}) = \operatorname{argmin}_{\bar{T}, v_0^i} \left\{ \sum_{i=1}^N \frac{1}{\sigma_V^2} \|v_0^i\|_V^2 + \frac{1}{\sigma_W^2} \|T_i - \phi^{v_0^i} \bar{T}\|_{W^*}^2 \right\} \quad (3.1)$$

called Fast Approximation with Mode (FAM). If we already defined Gaussian variables in the space of currents (Section 2), such variables have no pdf because

W^* and V are of infinite dimension. A more rigorous MAP derivation could be done considering finite dimensional parametrization of the v_0 's and of the T_i 's. For instance currents may be projected into a fixed grid; but this would require to adapt the registration scheme of [11] to account for such a discretization. Markov Chain Monte Carlo (MCMC) approaches for sampling the posterior could be also possible along the lines of [17] but this is still challenging and out of the scope of this paper.

3.2 A alternated Minimization Procedure

We solve Eq.3.1 by minimizing it alternatively with respect to the template and to the deformations. When the template \bar{T} is fixed, each term of 3.1 can be minimized separately. For a given observation T_i , minimizing $\frac{1}{\sigma_V^2} \|v_0^i\|_V^2 + \frac{1}{\sigma_W^2} \|T_i - \phi_i^{v_0^i} \cdot \bar{T}\|_{W^*}^2$ with respect to v_0^i is exactly a registration problem, as stated and solved in [11, 13]. This step of the minimization consists therefore of N registrations of the template \bar{T} to each observation T_i .

When the deformations ϕ_i are fixed for every $i = 1 \dots N$, minimizing 3.1 with respect to the template \bar{T} leads to the minimization of:

$$J(\bar{T}) = \frac{1}{2} \sum_{i=1}^N \|\phi_{\sharp, i} \star \bar{T} - T_i\|_{W^*}^2 \quad (3.2)$$

If all $\phi_i = \text{Id}$ (i.e. no deformation), the minimum is reached at the empirical mean: $\bar{T} = \frac{1}{N} \sum_{i=1}^N T_i$. For arbitrary deformations, there is no closed form and we use a gradient descent scheme. The gradient of Eq.3.2 is precisely:

$$\nabla_{\bar{T}} J = \sum_{i=1}^N \phi_{\sharp, i}^* \star (\phi_{\sharp, i} \star \bar{T} - T_i) \quad (3.3)$$

where ϕ_{\sharp}^* is the adjoint action of ϕ_{\sharp} : $\langle \phi_{\sharp}^* \star T, T' \rangle_{W^*} = \langle T, \phi_{\sharp} \star T' \rangle_{W^*}$ for any currents T and T' . This would be a matrix transpose if the action was linear. In this non-linear setting, standard computations give $\phi_{\sharp}^* \star T = \mathcal{L}_W(\phi_{\sharp} \star \mathcal{L}_W^{-1}(T))$. With the backward scheme, Eq.3.2 would be: $J(\bar{T}) = \frac{1}{2} \sum_{i=1}^N \|\bar{T} - \phi_{\sharp, i} \star T_i\|_{W^*}^2$, whose minimum has the closed form $\bar{T} = \frac{1}{N} \sum_{i=1}^N \phi_{\sharp, i} \star T_i$. We see here why the atlas building step is computationally more difficult in the forward setting. However, computing the likelihood of any new observations $p(T_i | \bar{T})$ will be much simpler and faster in this setting.

The input shapes T_i are sampled objects which are approximated as finite set of Dirac currents. As it will appear from this minimization procedure, the template will also always remain a finite set of Dirac currents at every iteration. Therefore, the current $\phi_{\sharp, i} \star \bar{T} - T_i$ is of the form $\sum_k \delta_{y_k^i}^{\beta_k^i}$ which finally leads to $\phi_{\sharp}^* \star \mathcal{L}_W^{-1}(\phi_{\sharp, i} \star \bar{T} - T_i) = \sum_k (d_x \phi_i)^t K(\phi_i(x), y_k^i) \beta_k^i$ in case of curves. Finally, the

vector field associated to the gradient of the energy 3.2 in case of curves can be computed at any point x of the space:

$$\mathcal{L}_W^{-1}(\nabla_{\bar{T}} J)(x) = \sum_{i=1}^N (d_x \phi_i)^t \left(\sum_k K(\phi_i(x), y_k^i) \beta_k^i \right) \quad (3.4)$$

For surfaces, $d_x \phi_i^t$ must be replaced by $|d_x \phi| d_x \phi^{-1}$. At this stage, we see that the gradient descent scheme could not be performed without an efficient numerical algorithm to estimate the current ∇J whose associated vector field is given by Eq.3.4. The sparse deconvolution method of Section 3.3 precisely provides a finite set of Dirac current which approximates ∇J at any arbitrary accuracy, so that the template remains a finite set of Dirac currents at each iteration.

We initialize the algorithm by setting $\phi_i = \text{Id}$, $\bar{T} = \frac{1}{N} \sum_{i=1}^N T_i$ and by computing $\gamma_{\bar{T}} = \mathcal{L}_W^{-1}(\bar{T})$ via a Gaussian convolution. The current \bar{T} is encoded as a list of (position, vectors) that approximate small segments or small triangles. The dense vector field $\gamma_{\bar{T}}$ is discretized at a fixed grid's points: $\Lambda = \{x_p\}$ and therefore encoded as an image of vectors. The following variable grad is also an image of vectors. We then iterate the following loop:

- For $i = 1 \dots N$, $\phi_i \leftarrow$ registration of \bar{T} to T_i .
- Until convergence do
 - grad = 0
 - For $i = 1 \dots N$ do
 - * transport tangents (normals) of \bar{T} with ϕ_i (gives $\phi_{\sharp, i} \star \bar{T}$).
 - * store the (y_k^i, β_k^i) such that $\phi_{\sharp, i} \star \bar{T} - T_i = \sum_k \delta_{y_k^i}^{\beta_k^i}$
 - * Deform Λ with ϕ_i and for each $p \in \Lambda$ compute:
 - $d_{x_p} \phi_i$ by a finite difference scheme
 - $G = \sum_k K(\phi_i(x_p), y_k^i) \beta_k^i$ (convolution)
 - $\text{grad}(p) \leftarrow \text{grad}(p) + 2 d_{x_p} \phi_i^t G$ (or $+2 |d_{x_p} \phi_i| (d_{x_p} \phi_i)^{-1} G$)
 - $\gamma_{\bar{T}} \leftarrow \gamma_{\bar{T}} - \tau \text{grad}$
 - Deconvolution of $\gamma_{\bar{T}}$ to give the new \bar{T} (See section 3.3).

3.3 Sparse Deconvolution by a Matching Pursuit Algorithm

Like any linear combinations of the input currents, the mean current \bar{T} is stored as the set of all weighted tangents (or normals) in the database $\bar{T} = (1/N) \sum_{i=1}^N T_i = \sum_{i=1}^n \delta_{x_k}^{\alpha_k}$ where n grows linearly with the number of observations. This heavy representation makes critical the registration of \bar{T} right after the initialization step. But this representation is often far from being optimal: it may be highly redundant at the scale λ_W . To integrate this redundancy, we compute the vector field $\mathcal{L}_W^{-1}(\bar{T})$ (Fig.2-a) by convolution. Then, an adequate deconvolution scheme could be applied to estimate an adapted basis on which the true solution T is decomposed in $T = \sum_{k'} \delta_{x_{k'}}^{\alpha_{k'}}$ with fast decreasing terms. The first terms of this series will give therefore an approximation of \bar{T} with an increasing accuracy (Fig.2-b).




Fig. 2. Toy example in 2D illustrating the sparse deconvolution scheme. **a-right:** two initial curves (L_1, L_2 in blue), their mean $((L_1 + L_2)/2)$ is represented by a set of 19 Dirac currents in red. **a-left:** a Gaussian convolution (\mathcal{L}_W^{-1}) of the 19 Dirac gives the dual representation of the mean as a dense vector field in W . **b-** The deconvolution method estimates iteratively a set of Dirac currents (**b-right** after 3 steps) whose convolution retrieves the initial vector field with increasing accuracy. The difference between the true vector field and the estimated one is shown in **b-left**.

In Eq.3.4, we make computations on vector fields and we do not know how to retrieve the current T that comes from the resulting vector field. The solution may require theoretically an infinite number of Dirac currents: $T = \sum_{i=0}^{\infty} \delta_{x_i}^{\alpha_i}$. Here again, we need a deconvolution scheme that estimates iteratively the most important terms of this series. The sum of these first terms will provide a sparse approximation of T at a given precision.

The deconvolution scheme we introduced in [9, 10], adapts to our framework based on currents the orthogonal matching pursuit (originally introduced in [18] to decompose images in adapted wavelet bases). Given a vector field $\gamma \in W$, the first step consists in finding the point x_1 that maximizes the projection of $\mathcal{L}_W(\gamma)$ on δ_{x_1} : $\langle \mathcal{L}_W(\gamma), \delta_x^{\varepsilon_k} \rangle_{W^*} = \langle \gamma, K(., x) \varepsilon_k \rangle_W = \gamma(x)_k$ (ε_k is an orthonormal basis of \mathbb{R}^3). x_1 reaches therefore the maximum of γ . A linear set of equation determines α_1 such that $\delta_{x_1}^{\alpha_1}$ is the orthogonal projection of $\mathcal{L}_W(T)$ into $\text{Span}(\delta_{x_1})$. We subtract then $K(., x_1)\alpha_1$ from γ and we iterate on this residue. This builds iteratively a series that has been proved in [10] to converge to the true solution.

4 Experimental Results

4.1 Building Atlases from Curves

We use a database of cortical sulcal landmarks (72 per brain) delineated in a large number of subjects scanned with 3D MRI (age: 51.8 ± 6.2 years). From 34 subjects in the database, we build the template according to our estimation method. We set the parameters $\lambda_W = 12\text{mm}$, $\lambda_V = 25\text{mm}$ and $\gamma = \sigma_W^2/\sigma_V^2 = 0.01$. The diameter of the brain is typically 120mm. Figure 3 shows the estimated template after 2 iterations. During the gradient descent, the bias (in the sense of our forward model Eq.1.1) is removed from the empirical mean (in red), leading




Fig. 3. Estimated template from 34 subjects. Left: initial 34 Sylvian Fissure of the left hemisphere (green), the empirical mean (red) and the estimated template (blue). In black, the mean lines computed from B-spline parametrization of curves [19]. Right: Same curves for 70 sulci. Although results look similar, only the template in blue is not biased in the sense of the model Eq.1.1.

to an unbiased template (in blue). The sparse deconvolution scheme enables also to give a light representation of the template: whereas the database contains 32643 segments (on average 960 segments per subject), the estimated template is represented by only 1361 Dirac currents with an approximation error below 5%. This would be of great interest in the future, for example to register this template toward any new available data.

4.2 Building Atlases from Surfaces

We use a database of 10 segmented sub-cortical structures of the brain (Caudate, Putamen, Globus Pallidus, Amygdala and Hippocampus) in large number of autistics and healthy children scanned with 3D MRI (age 2.7 ± 0.3) [20]. From 25 autistics and 7 controls, we build the template according to our estimation method for surfaces. We use the parameters $\lambda_W = 5\text{mm}$, $\lambda_V = 20\text{mm}$, $\gamma = \sigma_W^2/\sigma_V^2 = 0.001$. The set of 10 sub-cortical structures have typically a diameter of 60mm. Figures 4-a,b show the estimated template after 2 iterations. Each Dirac current δ_x^α is represented by an equilateral triangle whose center of mass is x and non-normalized normal α . Each of the 10 meshes has 2880 cells, so that the total number of normals is $720e^3$. Thanks to the deconvolution scheme, the estimated template is represented by only 1344 normals with an approximation error below 5%. In Fig.4-c, we show the difference between the template of autistics and controls approximated by the proposed matching pursuit algorithm. Note that the anatomical differences between both classes are not only captured by their respective templates but also by the main modes of the deformations ϕ_i .




Fig. 4. Estimated templates from 25 autistics (a) and 7 controls (b) for 10 sub-cortical structures (one per color). In (c), the blue arrows approximate the difference between autistics and controls' template, superimposed with the hippocampus of a control.

5 Conclusion and Perspectives

In this paper, we proposed a statistical model which estimates jointly a template, the deformations of the template to the observations and the noise on these observations. This differs from methods that use extrinsic templates, thus introducing a bias in the analysis of the variability. Our forward model considers observations as noisy deformations of an unknown template. Compared to the more commonly used backward setting, this model requires a specific deconvolution scheme to deal with the non-linear action of diffeomorphisms on curves and surfaces. However, once the model is built, we can decompose any new data into a deformation of the template and a residual noise and then measure how likely this decomposition may be with respect to the model. This offers a way to classify data according to pathologies, to detect outliers in a database, to highlight where a given observation differs from a model, etc. Such a statistical inference would be more difficult within the backward framework since the likelihood of any new observations would involve the Jacobian of the deformations that we take here directly into account while building the atlas.

Numerical experiments show the feasibility and relevance of our approach. In particular, two templates were estimated from deep brain structures of autistics children and controls. The results suggest anatomical differences between both classes. However, such results should be strengthened by rigorous statistical tests. Our future work will also evaluate the model with respect to its prediction and classification capability.

Acknowledgments

We warmly thank Prof. Thompson (UCLA) for providing the sulcal lines database through the Asclepios-LONI associated team **BrainAtlas**, as well as Prof. Gerig (University of Utah) for giving access to the database of deep brain structures. The work was partly supported by the European IP projet Health-e-Child (IST-2004-027749).

References

1. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* **23** (2004) 151–160
2. Avants, B., Gee, J.: Geodesic estimation for large deformation anatomical shape averaging and interpolation. *NeuroImage* **23** (2004) 139–150
3. Marsland, S., Twining, C.: Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images. *Transactions on Medical Imaging* **23**(8) (2004) 1006–1020
4. Zollei, L., Learned-Miller, E., Grimson, E., Wells, W.: Efficient population registration of data. In: *Computer Vision for Biomedical Image Applications*. Volume 3765 of *Lecture Notes in Computer Science.*, Springer (2005) 291–301
5. Glaunès, J., Joshi, S.: Template estimation from unlabeled point set data and surfaces for computational anatomy. In Pennec, X., Joshi, S., eds.: *Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2006)*. (2006)
6. Chui, H., Rangarajan, A., Zhang, J., Leonard, C.: Unsupervised learning of an atlas from unlabeled point-sets. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **26**(2) (2004) 160–172
7. Allasonnière, S., Amit, Y., Trouvé, A.: Towards a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society Series B* **69**(1) (2007) 3–29
8. Ma, J., Miller, M., Trouvé, A., Younes, L.: Bayesian template estimation in computational anatomy. To appear in *NeuroImage* (2008)
9. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Sparse approximation of currents for statistics on curves and surfaces. In: *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI)*. (2008) In press.
10. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Sparse approximation of currents for statistics on curves and surfaces. *Research Report 6571, INRIA (06 2008)*
11. Vaillant, M., Glaunès, J.: Surface matching via currents. In: *Proceedings of Information Processing in Medical Imaging*. Volume 3565 of *Lecture Notes in Computer Science.*, Springer (2005) 381–392
12. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Measuring brain variability via sulcal lines registration: a diffeomorphic approach. In: *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI)*. Volume 4791 of *Lecture Notes in Computer Science.*, Springer (2007) 675–682
13. Glaunès, J.: *Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique*. PhD thesis, Université Paris 13, <http://cis.jhu.edu/~joan/TheseGlaunes.pdf> (2005)
14. Aronszajn, N.: Theory of reproducing kernels. *Transactions of the American Mathematical Society* (68) (1950) 337–404
15. Trouvé, A.: Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision* **28** (1998) 213–221
16. Dupuis, P., Grenander, U., Miller, M.: Variational problems on flows of diffeomorphisms for image matching. *Quarterly of Applied Mathematics* **56**(3) (1998) 587–600
17. Allasonnière, S., Kuhn, E., Trouvé, A.: Bayesian deformable models building via a stochastic approximation algorithm: A convergence study. *Bernoulli Journal* (2008) In revision.

18. Pati, Y., Rezaifar, R., Krishnaprasad, P.: Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: Conf. Record of the 27th Asilomar Conference on Signals, Systems and Computers. Volume 1. (November 1993) 40–44
19. Fillard, P., Arsigny, V., Pennec, X., Hayashi, K., Thompson, P., Ayache, N.: Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines. *NeuroImage* **34**(2) (January 2007) 639–650
20. Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., Gilmore, J., Piven, J.: Magnetic resonance imaging and head circumference study of brain size in autism. *The Archives of General Psychiatry* **62** (2005) 1366–1376

MAP Estimation of Statistical Deformable Templates Via Nonlinear Mixed Effects Models : Deterministic and Stochastic Approaches

Stéphanie Allasonnière¹, Estelle Kuhn², and Alain Trounev³

¹ Center for Imaging Science, Johns Hopkins University, USA,

² LAGA, University Paris 13, France,

³ CMLA, ENS Cachan, France,

stephanie@jhu.edu, kuhn@math.univ-paris13.fr, alain.trouve@cmla.ens-cachan.fr. *

Abstract. In [1], a new coherent statistical framework for estimating statistical deformable templates relevant to computational anatomy (CA) has been proposed. This paper addresses the problem of population average and estimation of the underlying geometrical variability as a MAP computation problem for which deterministic and stochastic approximation schemes have been proposed. We illustrate some of the numerical issues with handwritten digit and 2D medical images and apply the estimated models to classification through maximum likelihood.

1 Introduction

For the last decade, we are witnessing impressive achievements and the emergence of elaborated registration theories [2–4] but the definition of a proper statistical framework for designing and inferring stochastic deformable templates in a principled way is much less mature. Despite a seminal contribution [5] and the fact that deformable templates can be cast into the general Grenander’s Pattern Theory [6], the down-to-earth and fundamental problem of computing *population averages* in presence of *unobserved* warping variables has not received so much attention from a more mathematical statistics perspective. More statistically oriented methods are slowly emerging [7–9] based on penalized likelihood or equivalently MDL approaches. Another line of research is to deal with the problem of population average as an estimation issue of proper stochastic (i.e. generative) models for which *consistency issues* should be addressed. In this direction, nonlinear mixed effects models (NLMM) are common tools in biostatistics and pharmacocinetic [10] to deal with both modelisation and inference of common *population factors* (fixed effects) and *distributions* of *unobserved* individuals factors (random effects). An active realm of research has emerged in the 90’s for designing efficient and consistent estimation algorithms. The importation of such ideas even in the limited context of population average of grey level images in CA is extremely appealing and challenging –both theoretically

* We are thankful to Dr. Craig Stark for providing us with the medical data

and practically— because of the very large (virtually infinite) dimensionality of the related factors (common template and individual warpings). These new avenues have started to be explored and theoretically consistent procedures based on recent advances on stochastic approximation algorithms have been proposed in a series of papers [1, 11, 12]. Since these papers are mainly mathematically focussed papers, we would like in the present paper to address some of the numerical issues of the various “EM-like” algorithms proposed to numerically approximate the Maximum A Posteriori estimator. Some relevant results on the USPS database and 2D medical images are presented, showing the strength of such methods.

The paper is organized as follows. Sections 2, 3 and 4 respectively recall the mixture model and how the estimation is completed and the particular case of the one component model. The last section, Section 5, is devoted to the experiments.

2 The observation model: BME-Templates

Consider a population of n gray level images $(y_i(s))_{s \in \Lambda}$ defined on a discrete grid of pixels Λ and assume that each observation y derives from a noisy sampling at the pixels locations $(x_s)_{s \in \Lambda}$ of an *unobserved* deformation field $z : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ of a *common* continuously defined template $I_0 : \mathbb{R}^2 \rightarrow \mathbb{R}$. This is what we call the Bayesian Mixed Effect Templates (BME-Templates). To keep things simple, we work within the small deformation framework [5] and assume that $y(s) = I_0(x_s - z(x_s)) + \sigma\epsilon(s) = zI_0(s) + \sigma\epsilon(s)$, where ϵ is a Gaussian normalized white noise and σ^2 is the common noise variance. The template I_0 and the deformation z are restricted to belong to subspaces of reproducing kernel Hilbert spaces V_p (resp. V_g) with kernel K_p (resp. K_g). Given $(p_k)_{1 \leq k \leq k_p}$ a fixed set of landmarks covering the image domain, the template function I_0 is parameterized by coefficients $\alpha \in \mathbb{R}^{k_p}$ through: $I_\alpha = \mathbf{K}_p \alpha$, where $(\mathbf{K}_p \alpha)(x) = \sum_{k=1}^{k_p} K_p(x, p_k) \alpha(k)$. Similarly we write $z_\beta = \mathbf{K}_g \beta$ with another set of landmarks $(g_k)_{1 \leq k \leq k_g}$ and a vector $\beta \in \mathbb{R}^{2k_g}$ of coefficients. In order to detect a global geometrical behavior, we consider the parameters β of the deformation field as an unobserved variable which is supposed to be Gaussian centered with covariance matrix Γ_g .

We present a general model based on NLMM defining a Bayesian mixture of m deformable template models (hereafter called components). In order to be able to consider small samples as our training sets, we have chosen to work within the Bayesian framework. In addition to the fact that some of the parameters, as the covariance matrix Γ_g , have been already used in many matching problems giving a first guess of what it could be, the Bayesian approach has its importance in the update formulas as a regularization term. This can particularly be noticed for Γ_g (cf [1]), where it always remains invertible in spite of the small sample size.

The model parameters of each component $t \in \{1, \dots, m\}$ are denoted by $\theta_t = (\alpha_t, \sigma_t^2, \Gamma_g^t)$. We assume that θ belongs to the open parameter space $\Theta \doteq \{ \theta = (\alpha_t, \sigma_t^2, \Gamma_g^t)_{1 \leq t \leq m} \mid \forall t \in \{1, \dots, m\}, \alpha_t \in \mathbb{R}^{k_p}, \sigma_t^2 > 0, \Gamma_g^t \in \Sigma_{2k_g, *}^+(\mathbb{R}) \}$ and $\rho = (\rho_t)_{1 \leq t \leq m}$ to the open simplex ϱ . Here $\Sigma_{2k_g, *}^+(\mathbb{R})$ is the set of strictly

For each component t (fixed effects) :

- ρ_t : probability of the component
- α_t : associated template parameter
- Γ_g^t : associated covariance matrix for deformation parameters
- σ_t^2 : associated additive noise variance

For each observation y_i (random effects) :

- τ_i : associated component
- β_i : deformation parameters
- ϵ_i : additive noise




Fig. 1. Mixed effect structure for our BME-template

positive symmetric matrices. Let $\eta = (\theta, \rho)$, the precise hierarchical Bayesian structure of our model is :

$$\begin{cases} \rho \sim \nu_\rho \\ \theta = (\alpha_t, \sigma_t^2, \Gamma_g^t)_{1 \leq t \leq m} \sim \otimes_{t=1}^m (\nu_p \otimes \nu_g) \mid \rho \\ \tau_1^n \sim \otimes_{i=1}^n \sum_{t=1}^m \rho_t \delta_t \mid \rho, \\ \beta_1^n \sim \otimes_{i=1}^n \mathcal{N}(0, \Gamma_g^{\tau_i}) \mid \tau_1^n, \eta \\ y_1^n \sim \otimes_{i=1}^n \mathcal{N}(z_{\beta_i} I_{\alpha_i}, \sigma_{\tau_i}^2 Id_\Lambda) \mid \beta_1^n, \tau_1^n, \eta \end{cases}$$

with

$$\begin{cases} \nu_\rho(\rho) \propto \left(\prod_{t=1}^m \rho_t \right)^{a_\rho}, \\ \nu_p(d\sigma^2, d\alpha) \propto \left(\exp\left(-\frac{\sigma_0^2}{2\sigma^2}\right) \frac{1}{\sqrt{\sigma^2}} \right)^{a_p} \cdot \exp\left(-\frac{1}{2}\alpha^t(\Sigma_p)^{-1}\alpha\right) d\sigma^2 d\alpha, \\ \nu_g(d\Gamma_g) \propto \left(\exp(-\langle \Gamma_g^{-1}, \Sigma_g \rangle / 2) \frac{1}{\sqrt{|\Sigma_g|}} \right)^{a_g} d\Gamma_g, \end{cases}$$

where the hyper-parameters are fixed. All priors are the natural conjugate priors and assumed independent. A natural choice for the a priori covariance matrices Σ_p and Σ_g is to consider the matrices induced by the metric of the spaces V_p and V_g . Define the square matrices $M_p(k, k') = K_p(p_k, p_{k'}) \forall 1 \leq k, k' \leq k_p$ and $M_g(k, k') = K_g(g_k, g_{k'}) \forall 1 \leq k, k' \leq k_g$, and then set $\Sigma_p = M_p^{-1}$ and $\Sigma_g = M_g^{-1}$, which are typical prior matrices used in many matching algorithms.

3 Estimation of the parameters

The parameter estimates are obtained by maximizing the posterior density on η conditional on y_1^n : $\hat{\eta}_n = \operatorname{argmax}_\eta q(\eta | y_1^n)$. Since the deformation coefficients β_1^n and component labels τ_1^n are unobserved, the natural approach is to use iterative algorithms such as EM [13] to maximize the penalized likelihood given

the observations y_1^n . This likelihood is written as an integral over the hidden variables, making the direct maximization a difficult task. The EM algorithm consists in an iterative procedure to solve this problem. Each iteration of the algorithm is divided into two steps; let l be the current iteration:

E Step: Compute the posterior law on (β_1^n, τ_1^n) as the following distribution:

$$\nu_l(\beta_1^n, \tau_1^n) \propto \prod_{i=1}^n q(y_i | \beta_i, \alpha_{\tau_i, l}) q(\beta_i | \Gamma_{g, l}^{\tau_i}) \rho_{\tau_i, l}$$

M Step: $\eta_{l+1} = \operatorname{argmax}_{\eta} E_{\nu_l} [\log q(y_1^n, \beta_1^n, \tau_1^n, \eta)]$.

In the present context, we initialize the algorithm with the prior model η_0 .

3.1 Fast approximation with modes (FAM)

The expression in the M step requires the computation of the expectation with respect to the posterior distribution of $\beta_1^n, \tau_1^n | y_1^n$, computed in the E step, which is known here up to the re-normalization constant. To overcome this obstacle, given an observation y_i and a label t , the posterior distribution of the random deformation field is approximated at iteration l by a Dirac law on its mode $\beta_{l, i, t}^*$. This yields the following computation :

$$\begin{aligned} \beta_{l, i, t}^* &= \arg \max_{\beta} \log q(\beta | \alpha_{t, l}, \sigma_{t, l}^2, \Gamma_{g, l}^t, y_i) \\ &= \arg \min_{\beta} \left\{ \frac{1}{2} \beta^t (\Gamma_{g, l}^t)^{-1} \beta + \frac{1}{2\sigma_{t, l}^2} |y_i - K_p^\beta \alpha_{t, l}|^2 \right\}, \end{aligned}$$

which is a standard template matching problem with the current parameters. We then approximate the joint posterior on (β_i, τ_i) as a discrete distribution concentrated at the m points $(\beta_{l, i, t}^*)_{1 \leq t \leq m}$ with weights given by: $w_{l, i}(t) \propto q(y_i | \beta_{l, i, t}^*, \alpha_{t, l}) q(\beta_{l, i, t}^* | \Gamma_{g, l}^t) \rho_{t, l}$. The label $\tau_{l, i}$ is then sampled from the distribution $\sum_{t=1}^m w_{l, i}(t) \delta_t$ and the deformation is the mode of the drawn label $\beta_{l, i} = \beta_{l, i, \tau_i}^*$.

The maximization is then done on this approximation of the likelihood.

3.2 Using a stochastic version of the EM algorithm : SAEM-MCMC

An alternative to the computation of the E-step in a complex nonlinear context is to use the stochastic approximation EM algorithm (SAEM) [14] coupled with an MCMC procedure [15] and a truncation on random boundaries. Our model belongs to the exponential density family which means that: $q(y, \beta, \tau, \eta) = \exp[-\psi(\eta) + \langle S(\beta, \tau), \phi(\eta) \rangle]$, where the sufficient statistic S is a Borel function on $\mathbb{R}^{2k_g} \times \{1, \dots, m\}$ taking its values in an open subset \mathcal{S} of \mathbb{R}^m and ψ, ϕ two Borel functions on $\Theta \times \varrho$ (the dependence on y is omitted for sake of simplicity).

We introduce the following function: $L : \mathcal{S} \times \Theta \times \varrho \rightarrow \mathbb{R}$ as $L(s; \eta) = -\psi(\eta) + \langle s, \phi(\eta) \rangle$. Direct generalisation of the proof in [1] to the multicomponent model

gives the existence of a critical function $\hat{\eta} : \mathcal{S} \rightarrow \Theta \times \varrho$ which satisfies: $\forall \eta \in \Theta \times \varrho, \forall s \in \mathcal{S}, L(s; \hat{\eta}(s)) \geq L(s; \eta)$. Then, iteration l of this algorithm consists of the following four steps.

Simulation step: The missing data are drawn using a transition probability of a convergent Markov chain having the posterior distribution as stationary distribution: $(\beta_{l+1}, \tau_{l+1}) \sim \Pi_{\eta_l}((\beta_l, \tau_l), \cdot)$

Stochastic approximation step: Since the model is exponential, the stochastic approximation is done on the sufficient statistics using the simulated values of the missing data: $s_{l+1} = s_l + \Delta_{l+1}(S(\beta_{l+1}, \tau_{l+1}) - s_l)$, where $(\Delta_l)_l$ is a decreasing sequence of positive step-sizes.

Truncation step: A truncation is done on the stochastic approximation.

Maximization step: The parameters are updated: $\eta_{l+1} = \hat{\eta}(s_{l+1})$.

Concerning the choice of Π_η used in the simulation step, as we aim to simulate (β_i, τ_i) through a transition kernel whose stationary distribution is $q(\beta, \tau | y_i, \eta)$, we simulate τ_i with a kernel whose stationary distribution is $q(\tau | y_i, \eta)$ and then β_i through a transition kernel that has $q(\beta | \tau, y_i, \eta)$ as stationary distribution. Given any initial deformation field $\xi_0 \in \mathbb{R}^{2k_g}$, we run, for each component t , J_l iterations of a hybrid Gibbs sampler (for each coordinate of the vector, a Hasting-Metropolis sampling is done given the other coordinates) $\Pi_{\eta, t}$ using the conditional prior distribution $\beta^j | \beta^{-j}$ as the proposal for the j^{th} coordinate, β^{-j} referring to β without its j^{th} coordinate. So that we get J_l elements $\xi_{t,i} = (\xi_{t,i}^{(k)})_{1 \leq k \leq J_l}$ of an ergodic homogeneous Markov chain whose stationary distribution is $q(\cdot | y_i, t, \eta)$. Denoting $\xi_i = (\xi_{t,i})_{1 \leq t \leq m}$, we simulate τ_i through the discrete density with weights given by: $\hat{q}_{\xi_i}(t | y_i, \eta) \propto \left(\frac{1}{J_l} \sum_{k=1}^{J_l} \left[\frac{f_t(\xi_{t,i}^{(k)})}{q(y_i, \xi_{t,i}^{(k)}, t | \eta)} \right] \right)^{-1}$, where f_t is the density of the Gaussian distribution $\mathcal{N}(0, \Gamma_{g,t})$. Then, we update β_i by re-running J_l times the hybrid Gibbs sampler Π_{η, τ_i} starting from a random initial point β_0 . It has been proved in [12], that the sequence $(\eta_l)_l$ generated through this algorithm converges a.s. toward a critical point of the penalized likelihood of the observations.

4 Single component model

We focus here on the single component model ($m = 1$). The unobserved variables are only the deformation fields β and the parameters are reduced to $\theta = (\alpha, \sigma^2, \Gamma_g)$. In this particular setting, denoting by P the distribution governing the observations and by $\Theta_* = \{ \theta_* \in \Theta \mid E_P(\log q(y | \theta_*)) = \sup_{\theta \in \Theta} E_P(\log q(y | \theta)) \}$, it has been proved in [1] that the MAP estimator $\hat{\theta}_n$ exists a.s. and converges toward an element in Θ_* . From the algorithmical viewpoint, the FAM algorithm does not require any changes. Indeed, each E step only corresponds to a single computation of the mode of the posterior density. However, the stochastic algorithm can be simplified. In the simulation step, only a single iteration of the Markov chain (i.e. $J_l = 1, \forall l$) is needed for each iteration of the SAEM algorithm: $\beta_{l+1} \sim \Pi_{\theta_l}(\beta_l, \cdot)$ yielding a non homogeneous Markov chain. It has been

proved in [11], that the sequence $(\theta_l)_l$ generated converges almost surely toward a critical point of the penalized likelihood of the observations.

5 Experiments

5.1 Estimation results

We illustrate this theoretical framework with the USPS handwritten digit database which corresponds to non noisy gray level images. In addition, we compare the two algorithmical approaches on 2D medical images of the corpus calosum (the splenium) and a part of the cerebellum.

Figure 2 shows the templates estimated from a training set (Figure 2-(a)) of 20 or 40 images per digit with both algorithms for the models with one and two components per class respectively. The results are quite similar, in particular the two components present the same features for both algorithms. Topologically different shapes are separated (cf digits 7 and 2) and the other digit clusters are relevant. While estimating a single component, the templates are good representatives of the shapes existing in the training set.

Concerning the geometrical variability, Figure 3, left image, presents some synthetic examples drawn with respect to the model with the estimated parameters. In spite of some artefacts described below, the kind of deformations learnt applied to the estimated templates looks like the elements of the training set which means that the algorithms capture this geometrical variability.

Last but not least, one could wonder how those algorithms deal with noisy images. In [1], this particular case has been shown to fail with the FAM algorithm with a toy example. Whereas, in [11, 12], the authors have proved the theoretical convergence of the two stochastic algorithms (for the mixture and simple models). This supports the fact that the estimated parameters should be less sensitive to the noise that can appear in the data. This is what we show in Figure 1 for a database of 20 images per digit which is partly presented (a). The results are related to the theory. Indeed, the FAM algorithm is stuck in




Fig. 2. (a) Some images of the USPS training set: 20 images per class. (b,c,d): Top row : FAM Algorithm, Bottom row : SAEM-MCMC algorithm. (b): one component prototype. (c-d): 2 component prototypes.




Fig. 3. 40 synthetic examples per class generated with the estimated parameters: 20 with the direct deformations and 20 with the inverse deformations. Left: from the non-noisy database estimated parameters. Right: from the noisy database estimated parameters. Note that the variability of digit is well reproduced, both in the case of highly deformable digits (e.g. 2 and 4) or in more constrained situations (e.g. 7 and 1).




Fig. 4. (a) Two images per digit of the noisy database. (b) Estimated prototypes in a noisy setting $\sigma^2 = 1$. (c) with the FAM algorithm. Right : with the SAEM-MCMC coupling procedure.

some local maximum of the likelihood (b) whereas the stochastic algorithm (c) reaches a better estimator for the parameters. This illustrates the power of the stochastic approach to solve this problem. Both the template and the geometrical distribution are well estimated. The results are presented in Figure 4 and in the right image of Figure 3 where we can notice that the estimation of the photometrical and the geometrical variability is quite robust to addition of a significant amount of noise.

The computational times of both algorithms for the simple model are very similar. The gradient descent required to compute the mode at each iteration lasts as long as one run of the Gibbs sampler used in the simulation step. The estimation takes only a couple of minutes on this dataset. For the general model, the SAEM-MCMC algorithm takes longer (increasing linearly with the number of component times the number of iterations of the Gibbs sampler J_l) since it requires the computation of many iterations of m Markov chains which can




Fig. 5. First row : Ten images of the training set representing the splenium and a part of the cerebellum. Second row : Results from the template estimation. (a) : gray level mean image of the 47 images. Templates estimated (b) : with the FAM (c) : with the stochastic algorithms on the simple model (d,e) : on the two component model. Third row : (f,g,h) : gray level mean image of the 47 images of the edges and estimated templates with the FAM and the stochastic algorithm on the simple model.

actually be easily parallelized. In addition, the number J_l of iterations of the Markov chain can be fixed all along the algorithm in the experiments.

We also test the algorithms on some medical images. The database we consider has 47 2D images, each of them representing the splenium (back of the corpus calosum) and a part of the cerebellum. Some of the training images are shown in Figure (5) first row.

The results of the estimation are presented in Figure 5 where we can see the improvement from the gray level mean (a) to our estimations. Image (b), corresponding to the deterministic algorithm result, shows a well contrasted splenium whereas the cerebellum remains a little bit blurry (note that it is still much better than the simple mean). Image (c), corresponding to the stochastic EM algorithm result, presents some real improvement again. Indeed, the splenium is still very contrasted, the background is not blurry and overall, the cerebellum is well reconstructed with several branches. The two anatomical shapes are relevant representants of the ones observed in the training set.

The estimation has been done while enabling the decomposition of the database into two components. The two estimated templates (using the MCMC-SAEM algorithm) are presented in Figure 5 (d) and (e). The differences can be seen in particular on the shape of the splenium, where the fornix is more or less close to the boundary of the image and the thickness of the splenium varies. The number of branches in the two cerebella also tends to be different from one template to the other (4 in the first component and 5 in the second one). The estimation suffers from the small number of images we have. To be able to explain the huge




Fig. 6. Estimated prototypes (20 images per digit), $\sigma_g = 0.2$ (Left), $\sigma_g = 0.3$ (Right) with images in $[-1, 1]^2$.

variability of the two anatomical shapes, more components would be interesting but at the same time more images required so that the components will not end up empty.

To emphasize the robustness of both algorithms, we run them on some binary images representing the edges of the same medical images. The exact same parameters are used and the results are shown in Figure 5, third row. Whereas the gray level mean image (f) does not represent any relevant information about the edges of the anatomical shapes, the FAM algorithm (g) tends to model the splenium and some branches of the cerebellum. Nevertheless, it does not lead to very contrasted shape boundaries as captured by the stochastic EM approach (h).

5.2 Optimization on the representation, model and algorithms

Despite the fact that many parameters (e.g. the noise variance) are self-calibrated during the estimation process, the algorithm depends on some hyper-parameters we would like to discuss briefly.

Data representation issues. The first point to be explained is the effect of the representation of the data, in particular the spline representation of both the template and the deformations (cf Section 2). We have chosen Gaussian kernels. The influence of their two scales can be seen on the template estimation. Indeed, choosing a too small geometric scale leads to very localized deformations around *fixed* control points and the resulting template is more blurry. In Figure 6, we present the results on a 20 handwritten digit images learning process. On the opposite side, a very large scale induces very smooth deformations which would no longer be relevant for the kind of deformations required to explain the database.

Concerning the photometric scale, it is straightforward that a large scale will drive to blurry template. This is particularly noticeable on digit 1 where the thickness significantly increases (cf Figure 7 two left images).

In addition, the effects of increasing scale can also be noticed on the learnt covariance matrix. Given a fatty template, the deformations required to fit the database will be forced to contract the template. This phenomena is thus important in the learnt covariance matrix. When we generate new data thanks to the estimated parameters, we can see, as in Figure 7 right images, that the template

is contracted, which is relevant, but also enlarged since the distribution on β is symmetric (this particular point is detailed in the next paragraph). Those large images are not typical from the training set.

Model distribution issues. One question is the relevance of the Gaussian distribution chosen for the deformation field. It is natural to think that the mean of the deformations around an atlas is close to zero whereas the symmetry of the distribution (the probability of a deformation field $+\beta$ equals its opposite one $-\beta$) is much more arguable. In Figure 3, we show the effects of the action of both fields on the learnt 10 digits templates. For example, digits 3 and 9 present, for some generated examples, irregular images whereas the opposite deformation leads to an image which is very similar to one or more element of the training set. Another distribution should be considered in future work.

Another issue about the model is the choice of the prior hyper-parameters. In particular, the effect of the inverse Wishart prior a_g on the geometric covariance matrix is important. Indeed, if we want to satisfy the theoretical requirements to the algorithms, we have to chose $a_g \geq 4k_g + 1$. However, the update formula is a barycenter between the expectation of the empirical covariance matrix and the prior with weights n and a_g respectively (cf: [1]). Since we are working with small sample sizes, this condition makes the update of Γ_g very constrained close to the prior Σ_g . This does not enable the geometry to be well estimated and the effects can be seen directly on the template but also on the classification rate [1]. The value of a_g used in those particular experiments is fixed to 0.5. Concerning the other weights (a_p, a_ρ) , their effects are less significant on the results and we fixed them to 200 and 2 respectively.

Stochastic algorithm issues. The FAM algorithm is deterministic and does not depend on any choice. Unfortunately, the stochastic algorithm requires several choices to optimize.

To optimize the choice of the transition kernel Π_η , we run the algorithm with different kernels and compare the evolution of the simulated hidden variables as well as the results on the estimated parameters. Some kernels, as an ordinary Hastings Metropolis algorithm using as proposal the prior or a standard random walk added to the current value, do not allow to visit well the entire support of the unobserved variable. From this point of view the hybrid Gibbs sampler we used has better properties and gives nice estimation results.



Fig. 7. Two left images: Estimated prototypes of digit 1 (20 images per class) for different hyper-parameters. Left: smaller geometry and larger photometric scales. Right: larger geometry and smaller photometric scales. Right images: Synthetic examples corresponding respectively to the two previous templates of digit 1.

To prove the convergence of the stochastic algorithms, we have to suppose that as soon as the stochastic approximation wanders outside an increasing compact set, the unobserved variable needs to be projected inside a given compact set (this is the truncation on random boundaries). In practice however, this step is never required, the results presented were obtained without this control.

Finally, the initialization of the parameters can lead to undesirable effects. For example, if the first value of the photometric parameter α is set to 0, at the first iteration of the Gibbs sampler, the proposal will be accepted with probability one. Since the candidate coordinates are simulated according to the conditional a priori, the resulting vector β leads to a variation which does not correspond to a relevant digit deformation. This implies some oscillations on the updated template. The next simulated deformation variable will try to take these oscillations into account to get closer and closer to the oscillating template, staying in its orbit. The results can be observed in Figure 6 (Right) specially for digit 1.

5.3 Results on classification rates

To get an objective way of comparing our algorithms and showing their performances, we use our model to propose a classifier which can easily be run on the USPS test set. We use the same approximations for the classification process, either a mode approximation of the posterior density or some MCMC methods to approximate the expectation required to compute the best class. Running the estimation with a FAM algorithm on all USPS database with 15 components and using a “mode” classifier gives a classification error rate of 3.5%. This is comparable to other classifiers results. The importance of the coupled photometric and geometric estimation is emphasized in [1].

Since the drawback of this method can be better proved in the presence of noise, we add an independent Gaussian noise of variance 1 on both the training set and the test set and run both estimations (with one component and 20 images per class) and both classifications. We run the parameter estimation though the “SAEM-like” algorithm presented in the previous section and test the model with these estimated parameters as a classifier. The classification error rate obtained are 22.52% when the classification uses the mode approximation and 17.07% using some MCMC methods. These results are a lot worse if the parameters are estimated with the FAM algorithm. For example, the classification error reaches 40.71% when the classification is done via the mode approximation as well.

6 Conclusion

We have presented some applications of the coherent statistical framework with BME-Template models described in [1, 11, 12]. This framework is fairly versatile and could be derived in many other important situations in CA. The possibility to work with mixture of deformable templates in a principled statistical way is also a quite enjoyable and unique feature of this setting. Reported experiments show that the deterministic FAM algorithm, despite its simplicity, performed

significantly worse especially under noisy conditions than the more sophisticated stochastic alternative. The introduction of such MCMC methods are still quite challenging in the 3D setting or for large deformation ([16] for a “FAM like” template estimation) but from an algorithmic point of view, there is a continuous interpolation from deterministic to stochastic algorithms (just increasing the number of MCMC steps) so that there is no sharp complexity gaps between to two approaches. Increasingly available computational power will make such stochastic approaches more and more appealing in the future.

References

1. Allasonnière, S., Amit, Y., Trouvé, A.: Toward a coherent statistical framework for dense deformable template estimation. *J.R.S.S.* **69** (2007) 3–29
2. Toga, W.A., Thompson, P.M.: The role of image registration in brain mapping. *Image and Vision Computing Journal* **19**(1–2) (2001) 3–24
3. Miller, M.I., Younes, L.: Group actions, homeomorphisms, and matching: A general framework. *Int. J. Comput. Vision* **41**(1–2) (2001) 61–84
4. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. *J. Math. Imaging Vision* **24** (2006) 209–228
5. Amit, Y., Grenander, U., Piccioni, M.: Structural image restoration through deformable templates. *JASA* **86** (1991) 376–387
6. Grenander, U., Miller, M.: *Pattern theory: from representation to inference*. Oxford; New York: Oxford University Press (2007)
7. Glasbey, C.A., Mardia, K.V.: A penalised likelihood approach to image warping. *Journal of the Royal Statistical Society, Series B* **63** (2001) 465–492
8. Marsland, S., Twining, C.J., Taylor, C.J.: A minimum description length objective function for groupwise non-rigid image registration. *Im. Vis. Comp.* (2007)
9. Van Leemput, K.: Probabilistic Brain Atlas Encoding Using Bayesian Inference. *MICCAI* **1** (2006) 704–711
10. Lindstrom, M., Bates, D.: Nonlinear mixed effects models for repeated measures data. *Biometrics* **46**(3) (1990) 673–687
11. Allasonnière, S., Kuhn, E., Trouvé, A.: Bayesian deformable models building via stochastic approximation algorithm: A convergence study. in revision
12. Allasonnière, S., Kuhn, E.: Stochastic algorithm for parameter estimation for dense deformable template mixture model. submitted
13. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society* **1** (1977) 1–22
14. Delyon, B., Lavielle, M., Moulines, E.: Convergence of a stochastic approximation version of the EM algorithm. *Ann. Statist.* **27**(1) (1999) 94–128
15. Kuhn, E., Lavielle, M.: Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM Probab. Stat.* **8** (2004) 115–131 (electronic)
16. Ma, J., Miller, M., Trouvé, A., Younes, L.: Bayesian template estimation in computational anatomy. To appear in *NeuroImage* (2008)

Semiparametric estimation of rigid transformations on compact Lie groups

J  r  mie Bigot¹, Jean-Michel Loubes¹ and Myriam Vimond²

¹ Institut de Math  matiques de Toulouse, Universit   de Toulouse, 31062 Toulouse Cedex 9, France,

² Universit   Haute Bretagne, Equipe de Statistique, Place du Recteur H. Le Moal, CS 24307, 35043 Rennes Cedex, France

Abstract. We study a simple model for the estimation of rigid transformations between noisy images. The transformations are supposed to belong to a compact Lie group, and a new matching criteria based on the Fourier transform is proposed. Consistency and asymptotic normality of the resulting estimators are studied. Some simulations are used to illustrate the methodology, and we describe potential applications of this approach to various image registration problems.

1 Introduction

Originating in Grenander’s pattern theory, transformation Lie groups are commonly used to model the deformations of images. The study of the properties and intrinsic geometries of such deformation groups is now an active field of research (see e.g. [3], [10], [12], [13], [14], [18]). In this setting, an important problem is the estimation of the deformations that may exist between similar images in the presence of additive noise. Rigid displacements such as translations, rotations or affine transformations can be modeled by finite dimensional Lie groups. Hence, the estimation of rigid deformation parameters can be formulated as a semi-parametric estimation problem which is an important field of research in mathematical statistics [4]. Indeed, semiparametric modeling is concerned with statistical problems where the parameters of interest are finite dimensional but where their observation is blurred by an infinite dimensional parameter. Here, the finite-dimensional parameters are the Lie group elements, and the infinite-dimensional parameter is an unknown 2D or 3D image, which is warped to obtain the different deformations. This image, that has to be recovered, is often called a template. The semiparametric framework provides optimal recovery of the warping parameters, contrary to nonparametric methods, leading to a better estimation of the template obtained by aligning the observed images.

As Lie groups are typically nonlinear spaces, an important question is the development of information geometry tools to extend classical notions such as asymptotic normality and efficiency, or the Cramer-Rao bound originally proposed for parameters lying in an Euclidean space. In the context of parametric statistics, several generalizations of these concepts to arbitrary manifolds have

been proposed (see e.g. [17]). However, in the more general situation of semi-parametric models, there is not so much work dealing with the estimation of parameters lying in a Lie group.

A first attempt in this direction has been proposed in [8], [20] for the simple problem of recovering shifts between one-dimensional noisy curves observed on an interval. The goal of this paper is extend such an approach to the more general case where the shift parameters belong to a compact Lie group. Thanks to a general shift property of the Fourier transform on compact Lie groups, a matching criterion based on the Fourier transform of the data can be defined, and we study some statistical properties of the resulting estimators.

In Section 2, a simple model for shifts on groups is introduced. Some properties on the Fourier transform are briefly reviewed (see [16] for further details), and are then used to define a general matching criterion on compact Lie groups. In Section 3, the consistency and the asymptotic normality of the estimators are established. Some numerical simulations are presented in Section 4, and extensions of the model are given in Section 5 for the registration of spherical images. Finally, we give some perspectives for future work. Note that the proofs of our theorems are based on the theory of M-estimation (see e.g. [19]), but are quite long and technical, and are thus not given in this manuscript but can be found in [2].

2 A shift model on Lie groups and a matching criteria based on the Fourier transform

2.1 The Fourier transform on compact Lie groups

Let G be a compact Lie group. Denote by e the identity element, and by hg the binary operation between two elements $h, g \in G$. Let $\mathbb{L}^2(G)$ be the Hilbert space of complex valued, square integrable functions on the group G with respect to the Haar measure dg .

To define a Fourier transform on $\mathbb{L}^2(G)$, a fundamental tool is the theory of group representations, which aims at studying the properties of groups via their representations as linear transformation of vector spaces. More precisely, a representation is an homomorphism from the group G to the automorphism group of a vector space. So let V be a finite-dimensional vector space, a *representation* of G in V is thus defined as a continuous homomorphism $\pi : G \rightarrow \text{GL}(V)$, where $\text{GL}(V)$ denotes the set of automorphisms of V . Hence it provides a linear transformation which depends on the vector space on which the group acts.

Every irreducible representation π of a compact group G in a vector space V is finite dimensional, so we denote by d_π the dimension of V . By choosing a basis for V , it is often convenient to identify $\pi(g)$ with a matrix of size $d_\pi \times d_\pi$ with complex entries. The function $g \mapsto \text{Tr} \pi(g)$ is called the *character* of π . The characters carry the essential information about the group representation. Moreover, the fundamental theorem of *Schur orthogonality* states that the characters form an orthonormal system in $\mathbb{L}^2(G)$ when π ranges over the dual set \hat{G} of irreducible

and equivalent representations of G . In the case of compact groups, the dual \hat{G} is a countable set, and the *Peter-Weyl Theorem* states that the characters are dense in $\mathbb{L}^2(G)$. Indeed, if π is a finite dimensional representation of G in the vector space V , then one can define, for every $f \in \mathbb{L}^2(G)$, the linear mapping $\pi(f) : V \rightarrow V$ by

$$\pi(f)v = \int_G f(g)\overline{\pi(g)}^T v dg, \text{ for } v \in V.$$

The matrix $\pi(f)$ is the generalization to the case of compact group of the usual notion of Fourier coefficient. Then, Peter-Weyl Theorem implies that (for simplicity the same notation is used for π and its equivalence class in \hat{G})

$$f(g) = \sum_{\pi \in \hat{G}} d_\pi \text{Tr}(\pi(g)\pi(f)) \text{ and } \|f\|_{\mathbb{L}^2(G)}^2 = \sum_{\pi \in \hat{G}} d_\pi \text{Tr}(\pi(f)\overline{\pi(f)}^T) \quad (1)$$

In the sequel, we will also denote by $\langle A, B \rangle_{HS} = \text{Tr}(\overline{A}^T B)$ the Hilbert-Schmidt inner product between two finite dimensional $d_\pi \times d_\pi$ matrices A and B .

2.2 A simple shift model on groups

Let \mathcal{X} be a subset of \mathbb{R}^d (with $d = 2, 3$ in our applications) and G be a compact Lie group acting on \mathcal{X} . A general deformation model for a set of J noisy signals would be

$$Y_j(x) = f^*(g_j^{-1} \cdot x) + W_j(x) \text{ for } x \in \mathcal{X}, \quad (2)$$

where $f^* : \mathcal{X} \mapsto \mathbb{R}$ is an unknown template, $W_j(x)$ is some additive noise, and g_j are the deformations to estimate. Our interest is to provide a very general framework for image registration under warping effect given by model (2), and thus to deal with the general problem of estimation of the elements of a Lie Group G acting on a space \mathcal{X} . In this paper, a simplest model for which $\mathcal{X} = G$ is studied to give the main ideas of our approach. An example for which $\mathcal{X} \neq G$ is given in Section 5 to show that our approach can be extended to more complex situations, and we discuss possible extensions to more complex situations in the concluding section.

Now, consider the following white noise model : for $j = 1, \dots, J$ and $g \in G$

$$dY_j(g) = f_j(g)dg + \epsilon dW_j(g), \quad (3)$$

where $f_j(g) = f^*(h_j^{*-1}g)$. The function $f^* : G \rightarrow \mathbb{R}$ is the unknown common shape of the observed images Y_j . The parameter $h_j^*, j = 1, \dots, J$ are the unknown shift parameters that we wish to estimate. $W_j, j = 1, \dots, J$ are independent standard Brownian sheets on the topological space G with measure dg , ϵ is an unknown noise level parameter which will tend to zero in our asymptotic considerations. Note that the white noise model (3) is a continuous model which is a very useful tool for the theoretical study of statistical problem in image analysis. In practice, the noisy images Y_j are typically discretely sampled on a

regular grid, but the model (3) can be shown to lead to comparable asymptotic theory in a sampled data model [5].

Obviously, without any further restriction on the set of possible shifts, the model (3) is not identifiable. Indeed, if s is an element of G with $s \neq e$, then one can replace the h_j^* 's in equation (3) by $\tilde{h}_j = h_j^* s$ and f^* by $\tilde{f}(g) = f^*(sg)$ without changing the formulation of the model. To ensure identification, we further assume that the set of parameters G^J is reduced to the subset $\mathcal{A}_0 \subset \mathcal{A}$ such that

$$\mathcal{A}_0 = \{(h_1, \dots, h_J) \in \mathcal{A}, h_1 = e\}. \quad (4)$$

This choice implies that the first image will be a reference and all the other images will be warped onto this particular choice. Now, remark that since $\pi(g)^T = \pi(g^{-1})$, one has that $\pi(f_j) = \pi(f)\pi(h_j^{*-1})$ for all $j = 1, \dots, J$. This equality is classically referred to as the *shift property* of the Fourier transform, and is at the heart of our estimation procedure to exhibit the shift parameters h_j^* .

2.3 A matching criterion based on the Fourier transform

For $h = (h_1, \dots, h_J) \in \mathcal{A}_0$, we propose to minimize the following criterion inspired by recent results in [8] and [20] for the estimation of shifts between curves:

$$M(h_1, \dots, h_J) = \frac{1}{J} \sum_{j=1}^J \left\| f_j \circ T_{h_j} - \frac{1}{J} \sum_{j'=1}^J f_{j'} \circ T_{h_{j'}} \right\|_{\mathbb{L}^2(G)}^2, \quad (5)$$

where $T_h : g \in G \rightarrow hg \in G$. Using the Parseval-Plancherel formula, the criterion may be rewritten in the Fourier domain as:

$$M(h) = M(h_1, \dots, h_J) = \frac{1}{J} \sum_{j=1}^J \sum_{\pi \in \hat{G}} d_\pi \left\| \pi(f_j) \pi(h_j) - \frac{1}{J} \sum_{j'=1}^J \pi(f_{j'}) \pi(h_{j'}) \right\|_{HS}^2, \quad (6)$$

for $h = (h_1, \dots, h_J) \in \mathcal{A}_0$. Given that $\pi(f_j) = \pi(f^*) \pi(h_j^{*-1})$, the criterion M has a minimum at $h^* = (h_1^*, \dots, h_J^*)$ such that $M(h^*) = 0$.

2.4 The empirical criterion

Our estimation method is then based on the Fourier Transform of the noisy data given by model (3). Let π an irreducible representation of G into V . We consider the following linear mappings from V to V which are defined from the model (3):

$$\pi(Y_j) = \int_G \pi(g^{-1}) dY_j(g) = \pi(f_j) + \epsilon \pi(W_j), \quad j = 1 \dots J,$$

where $\pi(W_j) = \int_G \pi(g^{-1}) dW_j(g)$, $j = 1 \dots J$. Let us denote by $(\pi_{kl}(W_j))$ the matrix coefficients of $\pi(W_j) : \pi_{kl}(W_j) = \int_G \pi_{kl}(g^{-1}) dW_j(g)$. Using the Schur

orthogonality and the fact that W_j is a standard Brownian sheet on G , one obtains that the complex variables $\pi_{kl}(W_j)$ are independent identically distributed Gaussian variables $\mathcal{N}_{\mathbb{C}}(0, d_{\pi}^{-1})$.

Let \hat{G}_{ϵ} be a finite subset of \hat{G} such that the sequence \hat{G}_{ϵ} increases when ϵ tends to 0 and $\cup_{\epsilon>0} \hat{G}_{\epsilon} = \hat{G}$. Practical choices for the set \hat{G}_{ϵ} will be discussed later on for the case of Abelian groups and the non-commutative group $\text{SO}(3)$. Then, we consider the following criterion:

$$M_{\epsilon}(h_1, \dots, h_J) = \frac{1}{J} \sum_{j=1}^J \sum_{\pi \in \hat{G}_{\epsilon}} d_{\pi} \left\| \pi(Y_j) \pi(h_j) - \frac{1}{J} \sum_{j'=1}^J \pi(Y_{j'}) \pi(h_{j'}), \right\|_{HS}^2 \quad (7)$$

and our M-estimator is given by $\hat{h}_{\epsilon} = \arg \min_{h \in \mathcal{A}_0} M_{\epsilon}(h)$.

3 Consistency and asymptotic normality of the M-estimator

Theorem 1. *Assume that $f^* \in \mathbb{L}^2(G)$ is such that there does not exist a closed subgroup H (except $H = \{e\}$ or $H = G$) such that $f(gh) = f(g)$ for all $g \in G$ and $h \in H$. Moreover, suppose that for all $\pi \in \hat{G}$ such that $\pi(f^*)$ is not identically null, then $\pi(f^*)$ is invertible. Then, M has a unique minimum at h^* , and if $\lim_{\epsilon \rightarrow 0} \epsilon^2 \sum_{\pi \in \hat{G}_{\epsilon}} d_{\pi}^2 = 0$, then \hat{h}_{ϵ} converges in probability to $h^* = (h_1^*, \dots, h_J^*)$.*

This Theorem shows the consistency of the estimators of the parameters h^* when the noise level goes to zero (recall that in the discretized model, this is equivalent to the fact that the number of observations increase). We stress that this result is obtained by minimizing a quadratic contrast without prior knowledge of the main pattern, thanks to the empirical criterion (7) which enables to get rid of f^* . Asymptotic normality of these estimators and thus rates of convergence will be provided in Theorem 2. The first condition of Theorem 1 ensures identifiability of the model since a function which would not satisfy this condition would be shift invariant, preventing any estimation. The second condition can be viewed as a smoothness condition which ensures unicity of the minimization scheme. In the following, we will denote by $\Re(z)$ the real part of any complex number z .

As the estimator \hat{h}_{ϵ} takes its values in a Lie group, it is not that obvious to define a notion of asymptotic normality as the space \mathcal{A}_0 is typically not a linear space if the group G is not Abelian. To overcome this, a classical approach is to use the exponential map to “project” \hat{h}_{ϵ} into the Lie algebra \mathcal{G} of G . If we can write $\hat{h}_{\epsilon} = \exp(\hat{u}_{\epsilon})$, then one can study the asymptotic normality of \hat{u}_{ϵ} which belongs to the linear space \mathcal{G} supposed to be of finite dimension p . For this, we first re-express the criterion M_{ϵ} defined on G^J as a function \tilde{M}_{ϵ} defined on \mathcal{G}^J . If G is a compact group, then the exponential map is surjective, but it is not necessarily injective. Hence, define \mathcal{U} to be a compact neighborhood of $0 \in \mathcal{G}^J$ onto which the exponential map is a smooth diffeomorphism, and let $\mathcal{A} = \exp(\mathcal{U})$. For $u \in \mathcal{U}^J$,

we re-express our criterion as $\tilde{M}(u_1, \dots, u_J) = M(\exp(u_1), \dots, \exp(u_J))$, and $\tilde{M}_\epsilon(u_1, \dots, u_J) = M_\epsilon(\exp(u_1), \dots, \exp(u_J))$. Then, define

$$\hat{u}_\epsilon = (\hat{u}_1, \dots, \hat{u}_J) = \arg \min_{u \in \mathcal{U}_1} \tilde{M}_\epsilon(u_1, \dots, u_J),$$

where $\mathcal{U}_1 = \{(u_1, \dots, u_J) \in \mathcal{U}^J, u_1 = 0\}$. Suppose that $h_1^* = \exp(u_1^*), \dots, h_J^* = \exp(u_J^*)$, with $u^* = (u_1, \dots, u_J) \in \mathcal{U}^J$, then the following result holds:

Theorem 2. *Assume that the conditions of Theorem 1 hold. Moreover, assume that for all $j = 2, \dots, J$ and $k = 1, \dots, p$*

$$\lim_{\epsilon \rightarrow 0} \epsilon \sum_{\pi \in \hat{G}_\epsilon} d_\pi^2 \|d_e \pi(d_{u_j^*} \exp(x^k))\|_{HS}^2 = 0, \quad (8)$$

and that

$$\lim_{\epsilon \rightarrow 0} \sup_{u_1 \in \mathcal{U}} \left\{ \sum_{\pi \in \hat{G} \setminus \hat{G}_\epsilon} d_\pi \|\pi(f^*)\|_{HS}^2 \|d_e \pi(d_{u_1} \exp(x^{k_1}))\|_{HS}^2 \right\} = 0 \quad (9)$$

$$\lim_{\epsilon \rightarrow 0} \sup_{u_1 \in \mathcal{U}} \left\{ \sum_{\pi \in \hat{G} \setminus \hat{G}_\epsilon} d_\pi \|\pi(f^*)\|_{HS}^2 \left\| \nabla_{u_1}^{x^{k_2}} d_e \pi(d_{u_1} \exp(x^{k_1})) \right\|_{HS} \right\} = 0 \quad (10)$$

$$\lim_{\epsilon \rightarrow 0} \epsilon^2 \sup_{u_1 \in \mathcal{U}} \left\{ \sum_{\pi \in \hat{G}_\epsilon} d_\pi^2 \|d_e \pi(d_{u_1} \exp(x^{k_1}))\|_{HS}^2 \right\} = 0 \quad (11)$$

$$\lim_{\epsilon \rightarrow 0} \epsilon^2 \sup_{u_1 \in \mathcal{U}} \left\{ \sum_{\pi \in \hat{G}_\epsilon} d_\pi^2 \left\| \nabla_{u_1}^{x^{k_2}} d_e \pi(d_{u_1} \exp(x^{k_1})) \right\|_{HS} \right\} = 0, \quad (12)$$

where x^1, \dots, x^p is an arbitrary basis of \mathcal{G} . Then,

$$\epsilon^{-1}(\hat{u}_\epsilon - u^*) \rightarrow N(0, H^{-1} \Sigma H^{-1}), \text{ as } \epsilon \rightarrow 0,$$

where Σ is a positive definite $(J-1)p \times (J-1)p$ matrix whose entries for $2 \leq j_1, j_2 \leq J$ and $1 \leq k_1, k_2 \leq p$ are given by

$$\begin{aligned} \Sigma_{(j_1, k_1) \times (j_1, k_1)} &= \frac{4}{J^2} \sum_{\pi \in \hat{G}} d_\pi \left(1 - \frac{1}{J}\right) \|\pi(f^*) d_e \pi(d_{u_{j_1}^*} \exp(x^{k_1}))\|_{HS}^2 \\ \Sigma_{(j_1, k_1) \times (j_2, k_2)} &= \frac{4}{J^2} \sum_{\pi \in \hat{G}} d_\pi \left(1 - \frac{1}{J}\right) \Re \left\langle \pi(f^*) d_e \pi(d_{u_{j_1}^*} \exp(x^{k_1})), \pi(f^*) d_e \pi(d_{u_{j_2}^*} \exp(x^{k_2})) \right\rangle_{HS}, \end{aligned}$$

and for $j_1 \neq j_2$ by

$$\Sigma_{(j_1, k_1) \times (j_2, k_2)} = -\frac{4}{J^2} \sum_{\pi \in \hat{G}} d_\pi \frac{1}{J} \Re \left\langle \pi(f^*) d_e \pi(d_{u_{j_1}^*} \exp(x^{k_1})), \pi(f^*) d_e \pi(d_{u_{j_2}^*} \exp(x^{k_2})) \right\rangle_{HS},$$

and where H is a positive definite $(J-1)p \times (J-1)p$ matrix whose entries for $2 \leq j_1, j_2 \leq J$ and $1 \leq k_1, k_2 \leq p$ are given by

$$\begin{aligned} H_{(j_1, k_1) \times (j_2, k_2)} = & -\frac{2}{J} \sum_{\pi \in \hat{G}} d_\pi \Re \left\{ \left\langle \pi(f^*) d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_2}) \right) d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_1}) \right) \right. \right. \\ & + \pi(f^*) \nabla_{u_{j_1}^*}^{x^{k_2}} d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_1}) \right), \pi(f^*) \rangle_{HS} \\ & \left. \left. + \frac{1}{J} \left\langle \pi(f^*) d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_1}) \right), \pi(f^*) d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_2}) \right) \right\rangle_{HS} \right\}, \end{aligned}$$

and for $j_1 \neq j_2$ by

$$H_{(j_1, k_1) \times (j_2, k_2)} = -\frac{2}{J^2} \sum_{\pi \in \hat{G}} d_\pi \Re \left\langle \pi(f^*) d_e \pi \left(d_{u_{j_1}^*} \exp(x^{k_1}) \right), \pi(f^*) d_e \pi \left(d_{u_{j_2}^*} \exp(x^{k_2}) \right) \right\rangle_{HS}.$$

The convergence result in the above theorem must be understood for the vector $\hat{u}_\epsilon = (\hat{u}_2, \dots, \hat{u}_J) \in \mathcal{G}^{J-1}$ since the first component is fixed to $\hat{u}_1 = 0$. Moreover, the notation $d_e \pi$ stands for the differential of π at e , while $d_{u_1} \exp(x^{k_1})$ corresponds to the differential of the exponential at u_1 in the direction x^{k_1} . We point out that the estimator converges at the parametric rate of convergence, and thus optimal rate of convergence, which would not have been the case if we had considered a preliminar estimate of f^* . This is one of the main achievements of the semiparametric type methodology proposed in this paper. Proving the optimality up to the constants imply studying the semiparametric efficiency of the estimators and falls beyond the scope of this paper. Some intuitions about such a result is provided in Section 6.

3.1 Abelian groups : the special case of the torus

The assumptions of Theorem 2 are rather technical and difficult to state in the very general case. However, for Abelian groups their statement is much simpler, which is due to the fact that the mapping $d \exp_u : \mathcal{G} \rightarrow \mathcal{G}$ reduces to the identity on \mathcal{G} i.e. $d_u \exp(v) = v$ for all $u \in G$ and $v \in \mathcal{G}$. The assumptions can be rewritten as

$$\lim_{\epsilon \rightarrow 0} \epsilon^2 \# \{ \hat{G}_\epsilon \} = 0, \quad \lim_{\epsilon \rightarrow 0} \epsilon \sum_{\pi \in \hat{G}_\epsilon} |d_e \pi(x^k)|^2 = 0, \quad \lim_{\epsilon \rightarrow 0} \sum_{\pi \in \hat{G} \setminus \hat{G}_\epsilon} |\pi(f^*)|^2 |d_e \pi(x^k)|^2 = 0, \quad (13)$$

where x^1, \dots, x^p is an arbitrary basis of \mathcal{G} . Thus, Condition (??) states that the common shape is differentiable and its derivatives are square integrable on \mathcal{G} . Conditions (13) give some assumptions on the choice of \hat{G}_ϵ .

As an illustrative example, let us consider the case where $G = (\mathbb{R}/\mathbb{Z})^p$ which corresponds to the classical multi-dimensional Fourier decomposition of a function $f \in \mathbb{L}^2([0, 1]^p)$

$$f(x) = \sum_{\ell \in \mathbb{Z}^p} c_\ell(f) e_\ell(x), \text{ for } x = (x_1, \dots, x_p) \in [0, 1]^d \text{ and } \ell = x = (\ell_1, \dots, \ell_p) \in \mathbb{Z}^d,$$

where $e_\ell(x) = \pi(x) = e^{-i2\pi(\sum_{k=1}^p \ell_k x_k)}$ and $c_\ell(x) = \pi(f) = \int_{[0,1]^d} f(x) e_\ell(x) dx$. Note that also that $d_e \pi(x^k) = -i2\pi \ell_k$. Now, take $\hat{G}_\epsilon = \{(\ell_1, \dots, \ell_p) \in \mathbb{Z}^p, |\ell_k| \leq \ell_\epsilon \text{ for all } k = 1, \dots, p\}$, for some positive integer ℓ_ϵ . Then, the following corollary holds:

Corollary 1. *Let $G = (\mathbb{R}/\mathbb{Z})^p$ and $f^* \in \mathbb{L}^2([0,1]^p)$ be a periodic function satisfying the conditions of Theorem 1. Assume that $h^* \in G^J$ or equivalently that $u^* \in ([0,1]^p)^J$. If*

$$\epsilon \ell_\epsilon^{p+2} = o(1) \text{ and } \sum_{(\ell_1, \dots, \ell_p) \in \mathbb{Z}^p} (|\ell_1|^2 + \dots + |\ell_p|^2) |c_\ell(f^*)|^2 < \infty,$$

then $\epsilon^{-1}(\hat{u}_\epsilon - u^*) \rightarrow N(0, \Gamma^{-1})$, as $\epsilon \rightarrow 0$, where the matrix Γ is given by

$$\begin{aligned} \Gamma_{(j_1, k_1) \times (j_1, k_2)} &= \sum_{\ell \in \mathbb{Z}} \left(1 - \frac{1}{J}\right) |c_\ell(f^*)|^2 (2\pi)^2 \ell_{k_1} \ell_{k_2}, \\ \Gamma_{(j_1, k_1) \times (j_2, k_2)} &= -\frac{1}{J} \sum_{\ell \in \mathbb{Z}} |c_\ell(f^*)|^2 (2\pi)^2 \ell_{k_1} \ell_{k_2} \text{ for } j_1 \neq j_2, \end{aligned}$$

4 Numerical simulations and some illustrative examples

4.1 A general gradient descent algorithm

To compute the estimator \hat{h}_ϵ one has to minimize the function $M_\epsilon(h)$. As this criterion is defined on a Lie group, a direct numerical optimization is generally not feasible. Finding minima of functions defined on a Lie group is generally done by reformulating the problem as an optimization problem on its Lie algebra. Since the expression of the gradient of $M_\epsilon(u)$ is available in a closed form, the following gradient descent algorithm with an adaptive step can be easily implemented:

Initialization : let $u^0 = 0 \in \mathcal{G}^J$, $\gamma_0 = \frac{1}{\|\nabla_{u^0} \tilde{M}_\epsilon\|}$, $M(0) = \tilde{M}_\epsilon(u^0)$, and set $m = 0$.

Step 2 : let $u^{new} = u^m - \gamma_m \nabla_{u^m} \tilde{M}_\epsilon$ and $M(m+1) = \tilde{M}_\epsilon(u^{new})$

While $M(m+1) > M(m)$ **do**

$\gamma_m = \gamma_m / \kappa$, and $u^{new} = u^m - \gamma_m \nabla_{u^m} \tilde{M}_\epsilon$, and $M(m+1) = \tilde{M}_\epsilon(u^{new})$

End while

Then, take $u^{m+1} = u^{new}$ and set $m = m + 1$.

Step 3 : if $M(m) - M(m+1) \geq \rho(M(1) - M(m+1))$ then return to Step 2, else stop the iterations, and take $\hat{h}_\epsilon = \exp(u^{m+1})$.

In the above algorithm, $\rho > 0$ is a small stopping parameter and $\kappa > 1$ is a parameter to control the choice of the adaptive step γ_m . Note that to satisfy the identifiability constraints the first p components of u^m are held fixed to zero at each iteration m .

4.2 Registration of translated 2D images

As an illustrative example, we consider the registration of translated 2D images (see [9] for a related work in a similar setting for 2D images). The chosen template f^* is the Shepp-Logan phantom image [1] of size $N \times N$ with $N = 100$ (see Figure 1). Noisy images can be generated by translating this image and adding Gaussian noise to each pixel value:

$$Y^j(i_1, i_2) = f^*\left(\frac{i_1}{N} - h_j^1, \frac{i_2}{N} - h_j^2\right) + \sigma z_j(i_1, i_2), \quad 1 \leq i_1, i_2 \leq N, \quad j = 1, \dots, J \quad (14)$$

where i_1, i_2 denotes a pixel position in the image, $z_j(i_1, i_2) \sim_{i.i.d.} N(0, 1)$, σ is the level of noise, and $h_j^1, h_j^2 \in [0, 1]$ are the unknown translation parameters to estimate. One could argue that the sampled data model (14) does not truly correspond to the white noise model (3). However, as previously explained there exists a correspondence between these two models in the sense that they are asymptotically equivalent if $\epsilon = \frac{\sigma}{N}$ (see [5]).

We have repeated $M = 100$ simulations with $J = 6$ noisy images simulated from the model (14). The various values taken for the translation parameters are the bold numbers given in Table 1. A typical example of a simulation run is shown in Figure 1. Here, $G = [0, 1] \times [0, 1]$ and its Lie algebra is $\mathcal{G} = \mathbb{R}^2$. The criterion $\tilde{M}_\epsilon(u)$ can be easily implemented via the use of the fast Fourier transform for 2D images:

$$\tilde{M}_\epsilon(u) = \frac{1}{J} \sum_{j=1}^J \sum_{|\ell_1| \leq \ell_\epsilon} \sum_{|\ell_2| \leq \ell_\epsilon} \left| y_{\ell_1, \ell_2}^j e^{i2\pi(\ell_1 u_1^j + \ell_2 u_2^j)} - \frac{1}{J} \sum_{j'=1}^J y_{\ell_1, \ell_2}^{j'} e^{i2\pi(\ell_1 u_1^{j'} + \ell_2 u_2^{j'})} \right|^2$$

for $u = (u_1^1, u_1^2, \dots, u_J^1, u_J^2)$, and where the y_{ℓ_1, ℓ_2}^j 's are the empirical Fourier coefficients of the image Y^j . Moreover, if $(x_1^1, x_1^2, \dots, x_J^1, x_J^2)$ denotes the canonical basis of the product space $(\mathbb{R}^2)^J$, then the components of the gradient of $\tilde{M}_\epsilon(u)$ are given by

$$\frac{\partial}{\partial x_j^k} \tilde{M}_\epsilon(u) = -\frac{2}{J} \sum_{|\ell_1| \leq \ell_\epsilon} \sum_{|\ell_2| \leq \ell_\epsilon} \Re \left((i2\pi \ell_k) y_{\ell_1, \ell_2}^j e^{i2\pi(\ell_1 u_1^j + \ell_2 u_2^j)} \overline{\left(\frac{1}{J} \sum_{j'=1}^J y_{\ell_1, \ell_2}^{j'} e^{i2\pi(\ell_1 u_1^{j'} + \ell_2 u_2^{j'})} \right)} \right).$$

According to Corollary 1, the smoothing parameter ℓ_ϵ should be chosen such that $\epsilon \ell_\epsilon^4 = o(1)$. Since the models (14) and (3) are asymptotically equivalent if $\epsilon = \frac{\sigma}{N}$, this condition becomes $\ell_\epsilon = \ell_N = o(N^{1/4})$. With $N = 100$, this suggests to take $\ell_N \leq 100^{1/4} \approx 3.16$. In Table 1 we give the empirical average of the estimated parameters over the $M = 100$ simulations, for the choice $\ell_N = 3$, together with their standard deviation. The results are extremely satisfactory as averages are very close to the true values and standard deviations are very small.

Fig. 1. A typical simulation run for $J = 6$ images generated from the model (14).**Table 1.** Average and standard deviation (in brackets) of the estimators $\hat{h}_j = (\hat{h}_j^1, \hat{h}_j^2)$ over $M = 100$ simulations. The bold numbers represent the true values of the parameters (h_j^1, h_j^2) .

	$j = 2$	$j = 3$	$j = 4$	$j = 5$	$j = 6$
h_j^1	0.07	0.1	0.05	-0.05	-0.08
\hat{h}_j^1	0.0704 (0.0031)	0.0997 (0.0031)	0.0494 (0.0028)	-0.0502 (0.0031)	-0.0801 (0.0032)
h_j^2	0.02	0.08	-0.10	-0.05	0.06
\hat{h}_j^2	0.0201 (0.0031)	0.0803 (0.0031)	-0.1002 (0.0030)	-0.0493 (0.0029)	0.0604 (0.0032)

5 Registration of spherical images

In many applications, data can be organized as spherical images i.e. as functions defined on the unit sphere \mathbb{S}^2 . For instance, spherical images are widely used in robotics since the sphere is a domain where perspective projection can be mapped, and an important question is the estimation of the camera rotation from such images (see [11]). Obviously such data do not correspond exactly to the shift model on group (3) as spherical images are defined on \mathbb{S}^2 while the shifts parameters belong the special orthogonal group $\text{SO}(3)$. However this setting corresponds to the general model (2) with $\mathcal{X} = \mathbb{S}^2$ and $G = \text{SO}(3)$, and a matching criterion similar to the one defined in equation (6) can still be defined by combining the spherical harmonics on \mathbb{S}^2 with the irreducible representations of $\text{SO}(3)$.

Indeed, let $x \in \mathbb{S}^2$ be a point on the unit sphere parameterized with spherical coordinates $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi[$. Then any $f \in \mathbb{L}^2(\mathbb{S}^2)$ can be decomposed as (see [6])

$$f(x) = \sum_{\ell=0}^{+\infty} \sum_{m=-\ell}^{\ell} c_{\ell}^m(f) \psi_{\ell}^m(x), \text{ with } c_{\ell}^m(f) = \int_0^{\pi} \int_0^{2\pi} f(\theta, \phi) \overline{\psi_{\ell}^m(\theta, \phi)} d\phi \sin(\theta) d\theta,$$

and where the functions $(\psi_{\ell}^m, \ell \in \mathbb{N}, m = -\ell, \dots, \ell)$ are the usual spherical harmonics which form an orthonormal basis of $\mathbb{L}^2(\mathbb{S}^2)$. Since these functions form a basis for the irreducible representations $(\pi_{\ell})_{\ell \in \mathbb{N}}$ of $\text{SO}(3)$ which are matrices of size $(2\ell + 1) \times (2\ell + 1)$, it follows that the action of a rotation $h \in \text{SO}(3)$ on a function $f \in \mathbb{L}^2(\mathbb{S}^2)$ is given by (see [6])

$$f(h^{-1}x) = \sum_{\ell=0}^{+\infty} c_{\ell}(f)^T \pi_{\ell}(h^{-1}) \psi_{\ell}(x) \quad \text{for all } x \in \mathbb{S}^2, \quad (15)$$

where $c_{\ell}(f) = (c_{\ell}^m(f))_{m=-\ell, \dots, \ell}$ and $\psi_{\ell}(x) = (\psi_{\ell}^m(x))_{m=-\ell, \dots, \ell}$ denotes vectors in $\mathbb{C}^{2\ell+1}$.

Now, suppose that one observes a set of noisy spherical images f_j that satisfy the following shift model: for $j = 1, \dots, J$ and $x \in \mathbb{S}^2$

$$dZ_j(x) = f_j(x)dx + \epsilon dW_j(x) \text{ with } dx = d\phi \sin(\theta)d\theta, \quad (16)$$

where $f_j(x) = f^*(h_j^{*-1}x)$, and $h_j^* \in \text{SO}(3)$, $j = 1, \dots, J$ are rotation parameters to estimate. For $h = (h_1, \dots, h_J) \in G^J$, the shift property (16) and the orthonormality of the spherical harmonics imply that the following matching criterion

$$N(h) = \frac{1}{J} \sum_{j=1}^J \left\| f_j \circ T_{h_j} - \frac{1}{J} \sum_{j'=1}^J f_{j'} \circ T_{h_{j'}} \right\|_{\mathbb{L}^2(\mathbb{S}^2)}^2, \quad (17)$$

where $T_{h_j} : x \in \mathbb{S}^2 \rightarrow h_j x \in \mathbb{S}^2$, can be written as

$$N(h) = \frac{1}{J} \sum_{j=1}^J \sum_{\ell=0}^{+\infty} \left\| c_\ell(f_j)^T \pi_\ell(h_j) - \frac{1}{J} \sum_{j'=1}^J c_\ell(f_{j'})^T \pi_\ell(h_{j'}) \right\|_{\mathbb{C}^{2\ell+1}}^2, \quad (18)$$

where $c_\ell(f_j)^T = c_\ell(f^*)^T \pi_\ell(h_j^{*-1})$ and $\|\cdot\|_{\mathbb{C}^{2\ell+1}}^2$ denotes the usual euclidean norm in $\mathbb{C}^{2\ell+1}$. Then, remark that the spherical harmonic coefficients of the noisy images Z_j are given by (in vector form) by $c_\ell(Z_j) = \int_{\mathbb{S}^2} \psi_\ell(x) dZ_j(x) = c_\ell(f_j) + \epsilon c_\ell(W_j)$, $j = 1 \dots J$, where $c_\ell(W_j) = \int_{\mathbb{S}^2} \psi_\ell(x) dW_j(x)$ is a complex random vector whose components are independent and identically distributed Gaussian variables $\mathcal{N}_{\mathbb{C}}(0, 1)$. Now, let ℓ_ϵ be an appropriate frequency cut-off parameter. The following empirical criterion can thus be proposed for registering spherical images:

$$N_\epsilon(h_1, \dots, h_J) = \frac{1}{J} \sum_{j=1}^J \sum_{\ell=0}^{\ell_\epsilon} \left\| c_\ell(Z_j)^T \pi_\ell(h_j) - \frac{1}{J} \sum_{j'=1}^J c_\ell(Z_{j'})^T \pi_\ell(h_{j'}) \right\|_{\mathbb{C}^{2\ell+1}}^2, \quad (19)$$

and an M-estimator of the rotation parameters is given by $\hat{h}_\epsilon = \arg \min_{h \in \mathcal{A}_0} N_\epsilon(h)$.

The criterion N_ϵ is very similar to the criterion M_ϵ , and the study of the consistency and the asymptotic normality of \hat{h}_ϵ can be done by following exactly the arguments as those developed in the previous sections.

6 Some perspectives and future work

The results on the asymptotic normality of the estimators show that there exists a significant difference between semi-parametric estimation on a linear Euclidean space and semi-parametric estimation on a nonlinear manifold. If the group G is non-commutative, then the asymptotic covariance matrix of the estimator \hat{u}_ϵ depends on the point u^* (and thus on h^*). Hence, this matrix can be interpreted as a Riemannian metric on G . This is a classical result in parametric

statistics for random variables whose law is indexed by parameters belonging to a finite-dimensional manifold. In such setting, the Fisher information matrix is a Riemannian metric and lower bounds analogue to the classical Cramer-Rao bound for parameters in an Euclidean space can be derived (see e.g. [17]). If G is supposed to be an Abelian group, then the asymptotic covariance matrix of the estimator does not depend on the point h^* since the parameter space for the shifts is a flat manifold.

An important issue is then the question of optimality of our estimators. We are currently studying analogs of the Cramer-Rao bound for the semi-parametric model (3), and in particular we are investigating if the covariance matrix given in Theorem 2 corresponds to the Fisher information matrix of this model. This result would provide a proof of the optimality of our reconstruction, even in a non asymptotic framework.

Another important question is the implementation of our approach for non-commutative groups. The numerical computation of our method for the registration of spherical images is more involved than the one used for the alignment of 2D images. Indeed, one has to deal with both the problem of computing the Fourier transform for images defined on a sphere, and with the problem of computing the irreducible representations of the group $SO(3)$ from its Lie algebra. We are currently working on the development of an efficient and fast numerical scheme to minimize the criterion N_ϵ , and we believe that this approach could yield good results for the registration of spherical images.

Finally, it should be mentioned that this work is rather preliminary and practical applications are restricted to estimation of shifts in 2D images and rotation on a sphere for the registration of spherical images. Another application that would be of great interest is the analysis of images of the fundus of the eye as described in [7]. However, many other interesting applications in medical images involve the study of more sophisticated Lie groups of transformations that are generally non-compact. We believe that there is a good chance to obtain satisfactory results for the estimation of deformations on non-commutative and compact groups. However, an important challenge is to investigate the extension of this work to the case of non-compact groups and at least to identify the difficulties introduced by such a generalization.

Acknowledgments

We are very much indebted to the three anonymous referees for their constructive criticism, comments and remarks that resulted in a substantial revision of the original manuscript.

References

1. Anil, K. J. (1989) Fundamentals of digital image processing. Prentice-Hall, Inc..
2. Bigot, J., Loubes, J.-M. & Vimond, M. (2008) Semiparametric Estimation of Shifts on Lie Groups for Image Registration. submitted paper.

3. Beg, M. F., Miller, M. I., Trouvé, A. & Younes, L. (2005) Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.*, **61**, 139157.
4. P.J. Bickel, Klaassen C.A.J., Ritov Y. and Wellner J.A. (1998) *Efficient and Adaptive Estimation for Semiparametric Models*. Springer Verlag.
5. Brown, L.D. & Low, M.G. (1996). Asymptotic equivalence of nonparametric regression and white noise. *Ann. Statist.*, **24**, 2384–2398.
6. Chirikjian, G. S. & Kyatkin, A. B. (2001). *Engineering Applications of Noncommutative Harmonic Analysis*. CRC Press.
7. Deguchi, K., Noami, J. and Hontani, H. (2000). 3D fundus pattern reconstruction and display from multiple fundusimages. *Proceedings 15th International Conference on Pattern Recognition*, Volume 4, 94 - 97.
8. Gamboa, F., Loubes, J.-M. and Maza, E. (2007) Semi-parametric estimation of shifts. *Electronic Journal of Statistics*, **1**, 616–640.
9. Glasbey, C. A. and Mardia, K. V. (2001) A penalized likelihood approach to image warping. *Journal of The Royal Statistical Society Series B*, **3**, 465-492.
10. Klassen, E., Srivastava, A., Mio, W. & Joshi, S. (2004) Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, **26**, 372383.
11. Makadia, A. and Daniilidis, K. (2006) Rotation Recovery from Spherical Images without Correspondences. *IEEE Trans. Pattern Anal. Mach. Intell.*, **28**(7), 1170–1175.
12. Marsland, S. and Twining, C. (2004) Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images. *IEEE Trans. Med. Imagng*, **23**.
13. Michor, P. and Mumford, D. (2006) Riemannian geometries on spaces of planar curves. *J. Eur. Math. Soc.*, **8**, 1-48.
14. Pennec, X. and Fillard, P. (2008) Statistical computing on non-linear spaces for computational anatomy. *Biomedical Image Analysis: Methodologies and Applications*, Springer, Editors: Paragios, N. and Duncan, J. and Ayache, N.
15. Rudin, W. (1990). *Fourier analysis on groups*. Wiley Classics Library, Reprint of the 1962 original edition, New York.
16. Sepanski, M.R. (2007). *Compact Lie groups*. Graduate Texts in Mathematics, Springer, New York.
17. Smith, S.T. (2005) Covariance, subspace, and intrinsic Cramér-Rao bounds. *IEEE Transactions on Signal Processing*, **53**, 1610-1630.
18. Trouvé, A. and Younes, L. (2005) Metamorphoses through lie group action. *Found. Computnl Math.*, **5**, 173198.
19. van der Vaart, A. W. (1998). *Asymptotic statistics*. Cambridge University Press.
20. Vimond, M. (2008) Efficient estimation in homothetic shifted in regression models. To be published in the *Annals of Statistics*.

Geodesic Shape Spaces of Surfaces of Genus Zero^{*}

Xiuwen Liu¹, Washington Mio², Yonggang Shi³, and Ivo Dinov³

¹ Department of Computer Science, Florida State University, Tallahassee, FL 32306

² Department of Mathematics, Florida State University, Tallahassee, FL 32306

³ Laboratory of Neuro Imaging, UCLA School of Medicine, Los Angeles, CA 90095

Abstract. We construct shape spaces of elastic spherical surfaces immersed in Euclidean space \mathbb{R}^k . The spaces are equipped with geodesic metrics that depend on the tension and rigidity of the surfaces. We develop algorithms to calculate geodesics and geodesic distances, as well as tools to quantify local shape similarities and contrasts, thus obtaining a local-global formulation. We give examples of geodesic interpolations and illustrations of the use of the model in brain mapping.

1 Introduction

The development of computational models of shapes of surfaces in Euclidean space is motivated by core problems in computational anatomy such as mapping the human brain, characterizing normal variation in anatomy and identifying pathological changes. In recent years, significant progress has been made in the study of Riemannian shape spaces of curves (see e.g. [1–5]). However, a computationally feasible approach to higher-dimensional shapes still poses many challenges. In this paper, we construct a geodesic shape space of elastic surfaces of genus zero in \mathbb{R}^k , with shapes realized as piecewise linear immersions of a triangle mesh K homeomorphic to the 2-sphere. We represent a parametric shape by its discrete exterior derivative because first-order representations provide a good balance between computational tractability and geometrical accuracy. Representations beyond order zero also are more effective in capturing deformations such as elastic bends and folds. The space of immersions is equipped with a family of Riemannian metrics that reflect the elasticity properties of the surfaces. The general model is anisotropic and inhomogeneous, as it allows the tension and rigidity of a surface to vary throughout its extension. The representation is invariant under translations and a shape space of spherical surfaces is obtained by normalizing scale and accounting for the action of orthogonal transformations.

The paper is devoted to the development of the elastic shape models and algorithms to calculate geodesic paths and geodesic distances in shape space. We use these basic tools to estimate mean shapes and to construct an anatomical

^{*} This work was supported in part by NSF grants DMS-0713012 and CCF-0514743, and NIH Roadmap for Medical Research grant U54 RR021813.

atlas of the right hippocampus. Our shape model is used in conjunction with existing surface parametrization and registration techniques [6–8], but research in these areas is still ongoing and is complementary to the study of shape spaces. Although shape distance is a global measurement of shape dissimilarity, the formulation used also enables us to characterize and quantify local shape differences. This is important in applications in order to identify the regions where the most significant morphological contrasts occur. Shapes of surfaces can be studied from several other viewpoints. Alternative approaches include: (i) models based on landmark representations [9]; (ii) models that rely on diffeomorphisms acting on a volume containing a surface to describe shape deformations [10]; (iii) models based on medial-axis representations [11].

The paper is organized as follows. In Section 2, we discuss the shape representation and examine invariance under shape-preserving transformations. In Section 3, we define the elastic shape metrics. The algorithm to calculate geodesics via energy minimization is presented in Sections 4 and 5, where illustrations are also provided. Energy density functions and localization techniques are discussed in Section 6. We conclude with a summary and some discussion in Section 7.

2 Shape Representation

Let K be a finite simplicial complex whose underlying polyhedron $|K|$ is homeomorphic to the unit 2-sphere \mathbb{S}^2 . K will be fixed throughout the discussion. A parametric spherical shape will be realized as a mapping $\alpha: |K| \rightarrow \mathbb{R}^k$, which is linear on each simplex of K . Thus, if $V = \{z_1, \dots, z_p\}$ is the vertex set of K , α is completely determined by its restriction $\alpha: V \rightarrow \mathbb{R}^k$ to V . Next, we define the discrete exterior derivative $d\alpha$, which is a measure of the variation of α along the edges of K . Fix an orientation for each edge and each triangle of K and let $E = \{e_1, \dots, e_m\}$ and $F = \{T_1, \dots, T_n\}$ be the oriented edge and face sets of K . For an oriented edge e , let e^- and e^+ denote its initial and terminal vertices, respectively. Then, $d\alpha: E \rightarrow \mathbb{R}^k$ is given by $d\alpha(e_i) = \alpha(e_i^+) - \alpha(e_i^-)$. Note that it suffices to define $d\alpha$ over the oriented edges in E because if we reverse the orientation of an edge, the variation of α gets multiplied by -1 . We only consider mappings such that $d\alpha(e_i) \neq 0$, for every $e_i \in E$, which we refer to as *immersions*. This just means that no edge gets crushed to a point under α . For each i , $1 \leq i \leq m$, we write the modular component of $d\alpha(e_i)$ in logarithmic scale as $r_i = \log \|d\alpha(e_i)\|$ and the directional component as $v_i = d\alpha(e_i) / \|d\alpha(e_i)\|$. Thus, the (oriented) i th edge of α is represented by the vector $e^{T_i} v_i$. Writing each v_i as a row vector $v_i = [v_{i1} \dots v_{ik}]$, the parametric surface α will be represented by the pair $(r, v) \in \mathbb{R}^m \times \mathbb{R}^{m \times k}$, where $r = [r_1 \dots r_m]^T$ and v is the $m \times k$ matrix whose i th row is v_i . Since the Riemannian metrics on $\mathbb{R}^m \times \mathbb{R}^{m \times k}$ to be used in the development of our shape models differ from the standard Euclidean metric, we will use the notation \mathbb{E} for this space to emphasize this fact. Thus, \mathbb{E} is the space of dimension $m(k+1)$ formed by all pairs (r, v) , with the topology induced by the Euclidean metric. Thus far, the only constraint on (r, v) is that

each row of v must be a unit vector. We denote by M the subspace of \mathbb{E} formed by all pairs (r, v) with this property.

2.1 Pre-Shape Space

The representation of α via (r, v) is clearly invariant under translations. To fix scale, we set the total edge length to be unitary; that is, we normalize the pair (r, v) to satisfy

$$G_1(r, v) = \sum_{i=1}^m e^{r_i} = 1. \quad (1)$$

Remark. Alternatively, one may normalize scale by fixing the total area, or simply drop this condition if a scale-sensitive model is desired.

What pairs (r, v) represent the exterior derivative of a mapping $\alpha: |K| \rightarrow \mathbb{R}^k$? The integrability conditions will reveal the further constraints to be imposed on (r, v) . If p is an oriented path in K formed by a sequence of oriented edges, let the integer n_i be the net number of times that the oriented edge $e_i \in E$ is traversed by p , where a negative sign indicates reversal of orientation. Given an immersion α , its variation along p may be expressed as $\sum_{i=1}^m n_i d\alpha(e_i) = \sum_{i=1}^m n_i e^{r_i} v_i$. The variation of α along any (oriented) cycle c in K clearly must vanish. Conversely, given (r, v) , define the integral of (r, v) along the path p to be $I(p) = \sum_i n_i e^{r_i} v_i$. If $I(c)$ vanishes along every cycle c , then (r, v) represents the exterior derivative of some α , uniquely determined up to translations. This can be seen as follows: fix a vertex v_0 of K and a point $x_0 \in \mathbb{R}^k$ and define $\alpha(v_0) = x_0$. For any other vertex v , choose a path p from v_0 to v and let $\alpha(v) = x_0 + I(p)$. The vanishing condition over cycles ensures that $\alpha(v)$ is independent of the path chosen.

Verifying that $I(c) = 0$, for every cycle c , is not computationally feasible. However, it suffices to check this vanishing condition on cycles that are boundaries of oriented triangles (this uses the fact that the 2-sphere is simply connected). This is because the integral along any cycle in K can be accounted for by combining the integrals along the boundaries of appropriately chosen oriented triangles. For each oriented triangle $T_\ell \in F$, let $\varepsilon_{\ell_1} e_{\ell_1}, \varepsilon_{\ell_2} e_{\ell_2}, \varepsilon_{\ell_3} e_{\ell_3}$ be the oriented edges of T_ℓ , where $\varepsilon_{\ell_i} = \pm 1$, $e_{\ell_i} \in E$, and the indexes satisfy $\ell_1 < \ell_2 < \ell_3$. Then, the argument above shows that a pair (r, v) represents the exterior derivative of an immersion α if and only if it satisfies the following kn integrability conditions:

$$G_{\ell,j}(r, v) = \sum_{i=1}^3 \varepsilon_{\ell_i} e^{r_{\ell_i}} v_{\ell_i,j} = 0. \quad (2)$$

where $1 \leq \ell \leq n$ and $1 \leq j \leq k$. It is not difficult to see that one may drop condition (2) over one of the triangles, say T_n , since the variation along the boundary of any of the triangles can be expressed as a combination of the variations along the boundaries of the remaining ones. Thus, we obtain $k(n-1)$ independent integrability conditions.

Pairs $(r, v) \in M$ satisfying (1) and (2) will be called *pre-shapes* and the space of all pre-shapes will be denoted P . A pre-shape gives a representation of immersions that is invariant to scale and translations. Thus, we have introduced

a nested sequence $P \subset M \subset \mathbb{E}$ of spaces, which will be useful at various stages of our constructions.

Remark. A more formal explanation of the integrability of (r, v) can be given in terms of simplicial cohomology. The vanishing of the integral of (r, v) along the boundaries of oriented triangles means that the simplicial 1-cochain (with coefficients in \mathbb{R}^k) determined by $e_i \mapsto e^{r_i} v_i$, $1 \leq i \leq m$, is a 1-cocycle. This cocycle represents the trivial cohomology class because $|K|$ is simply connected. Thus, the cocycle is integrable. This interpretation also allows us to identify the further integrability conditions needed for a complex K of different topology. The details will be discussed in future work.

2.2 Shape Space

Let $O(k)$ be the group of $k \times k$ orthogonal matrices. If $U \in O(k)$ and α is an immersion, the result of applying the rigid transformation U to α is the composition $U \circ \alpha: V \rightarrow \mathbb{R}^k$. The induced action on (r, v) is $(r, v) \mapsto (r, vU^T)$. The action on the modular component is trivial because edge lengths do not change under U . Pre-shapes clearly get transformed into pre-shapes and the orbit of (r, v) is $\mathcal{O}(r, v) = \{(r, vU^T) : U \in O(k)\}$. Since pre-shapes that differ by a rigid transformation have the same shape, we define the shape space S as the orbit space of P under the action of $O(k)$. In other words, as the quotient $S = P/O(k)$.

3 Riemannian Metrics

We introduce Riemannian structures on \mathbb{E} that will lead to a family of elastic pre-shape and shape metrics. The general model is anisotropic and inhomogeneous, but the family contains a special 1-parameter homogeneous class. The idea is that each edge of a shape may offer different resistance to stretching/compression and bending, which will be quantified by its tension and rigidity coefficients. The geodesic distance will be related to the minimum energy required to deform a shape into another under these conditions.

Let $a, b: E \rightarrow \mathbb{R}^+$ be positive functions defined on the edge set, where $a_i = a(e_i)$ and $b_i = b(e_i)$ represent the tension and rigidity coefficients of the i th edge. Consider the Riemannian structure on \mathbb{E} given at (r, v) by

$$\langle (h, w), (\bar{h}, \bar{w}) \rangle_{(r, v)} = \sum_{i=1}^m a_i h_i \bar{h}_i e^{r_i} + \sum_{i=1}^m b_i (w_i \cdot \bar{w}_i) e^{r_i}. \quad (3)$$

If we think of (h, w) as an infinitesimal deformation of (r, v) , then $\|(h, w)\|_{(r, v)}^2 = \langle (h, w), (h, w) \rangle_{(r, v)}$ is the energy cost of the deformation. Henceforth, the submanifolds P and M will have the Riemannian structure induced from (3). If (r, v) and (\bar{r}, \bar{v}) are pre-shapes, we let $\delta((r, v), (\bar{r}, \bar{v}))$ be the geodesic distance

between them on the pre-shape manifold P . If s and \bar{s} are the shapes associated with the orbits of (r, v) and (\bar{r}, \bar{v}) , the shape distance is defined as

$$d(s, \bar{s}) = \inf_{U, V \in O(k)} \delta((r, vV^T), (\bar{r}, \bar{v}U^T)) = \inf_{U \in O(k)} \delta((r, v), (\bar{r}, \bar{v}U^T)). \quad (4)$$

The last equality follows from the fact that the Riemannian metric is compatible with the action of $O(k)$; that is, $O(k)$ acts by isometries. Most of the remaining work will be devoted to the calculation of geodesics in P and the geodesic distance δ . The minimization over $O(k)$ is a relatively simpler calculation, as indicated below.

4 Path Spaces and the Energy Functional

Our next goal is to construct a pre-shape geodesic between $(r, v), (\bar{r}, \bar{v}) \in P$. The idea is to begin with a path in P connecting the pre-shapes and gradually deform it to a geodesic following the negative gradient flow of the path energy. Implicit in this statement is that a path space equipped with a Riemannian structure has been constructed.

4.1 Path Spaces

Let $I = [0, 1]$. If $p: I \rightarrow M$ is a path in M , we write its modular and directional components as $p(t) = (r^t, v^t)$. Let Y be the manifold of all absolutely continuous paths in M (with square integrable derivative). A tangent vector to Y at p represents an infinitesimal deformation of the path and is given by a vector field (h^t, w^t) along p with the property that $w_i^t \cdot v_i^t = 0$, for every $1 \leq i \leq m$ and $\forall t \in I$. This follows from the fact that $\|v_i^t\|^2 = 1$ so that the deformation of each v_i^t has to be tangential to the unit sphere. Consider the Riemannian metric on the path space Y given by

$$\langle (h, w), (\bar{h}, \bar{w}) \rangle_p = \langle (h^0, w^0), (\bar{h}^0, \bar{w}^0) \rangle_{p(0)} + \int_0^1 \langle D_t(h^t, w^t), D_t(\bar{h}^t, \bar{w}^t) \rangle_{p(t)} dt, \quad (5)$$

where D_t denotes covariant differentiation in M along p . This type of inner product in function space was, to our knowledge, introduced by Palais [12]. Motivated by the inclusion $P \subset M$, we consider the following submanifolds of Y :

- (i) The space $Z_M \subset Y$ of paths in M satisfying the boundary conditions $p(0) = (r, v)$ and $p(1) = (\bar{r}, \bar{v})$.
- (ii) The space $Z_P \subset Z_M$ of paths in the pre-shape space P satisfying the boundary conditions described in (i). A path in Z_P has the property that each (r^t, v^t) also satisfies (1) and (2).

We thus have three nested path spaces: $Z_P \subset Z_M \subset Y$.

4.2 The Energy Functional

On the path space Y , define the energy functional $E: Y \rightarrow \mathbb{R}$ by

$$E(p) = \frac{1}{2} \int_0^1 \langle (\partial_t r^t, \partial_t v^t), (\partial_t r^t, \partial_t v^t) \rangle_{p(t)} dt. \quad (6)$$

A pre-shape geodesic between $(r, v), (\bar{r}, \bar{v}) \in P$ is a path $p \in Z_P$, which is a critical point of the energy E restricted to Z_P . We are particularly interested in minimal energy paths since they represent minimal-length geodesics. We exploit the inclusions $Z_P \subset Z_M \subset Y$ in our approach to the minimization problem, the point being that it is easier to calculate the gradient of E as a functional on Z_M . The use of (5) leads to a particularly elegant and computationally robust expression for the gradient.

5 Geodesics

In this section, we develop an algorithm to calculate a geodesic in P between two pre-shapes (r, v) and (\bar{r}, \bar{v}) using a gradient search to minimize the energy.

5.1 Initialization

To initialize the process, let $\alpha, \bar{\alpha}: V \rightarrow \mathbb{R}^k$ be immersions of K associated with (r, v) and (\bar{r}, \bar{v}) , respectively. We linearly interpolate α and $\bar{\alpha}$ to obtain a 1-parameter family of mappings $\alpha_t: K \rightarrow \mathbb{R}^k$. If $d\alpha_t(e_i) = 0$, we gently deform α_t to make it non-singular and then scale each α_t to turn $d\alpha_t$ into a pre-shape. The path $p(t) = d\alpha_t$ is used to initialize the search.

5.2 Covariant Integration and Parallel Transport

To calculate the gradient of E at p , we first discuss covariant integration in M of a vector field (f^t, x^t) along a path $p(t) = (r^t, v^t)$, where f^t , where x^t denote the modular and directional components of the field. The tangentiality to M of the field (f^t, x^t) just means that $x_i^t \cdot v_i^t = 0, \forall t \in I$ and every $i, 1 \leq i \leq m$. As shown in Appendix B, a vector field (F^t, X^t) along p is tangential to M and is a *covariant integral* of (f^t, x^t) if and only if it satisfies the system of differential equations

$$\begin{cases} \partial_t F_i^t = f_i^t - \frac{1}{2}(\partial_t r_i^t)F_i^t + \frac{1}{2}\frac{b_i}{a_i}(X_i^t \cdot \partial_t v_i^t) \\ \partial_t X_i^t = x_i^t - \frac{1}{2}(X_i^t \partial_t r_i^t + F_i^t \partial_t v_i^t) - (X_i^t \cdot \partial_t v_i^t)v_i^t. \end{cases} \quad (7)$$

Numerical integration of (7) with initial conditions (F^0, X^0) will be used in the calculation of geodesics. In the special case where the field (f^t, x^t) is identically zero, the integral field (F^t, X^t) is the *parallel transport* of (F^0, X^0) along p .

5.3 The Gradient of the Energy

Given a path $p \in Z_P$, we first calculate the gradient of E at p as a functional on the path space Y . For this purpose, we consider a variation $(r^t(\mu), v^t(\mu))$ of $p(t) = (r^t, v^t)$ in Y along a direction (h^t, w^t) , where $\mu \in (-\epsilon, \epsilon)$ is the variation parameter. This means that $(r^t(0), v^t(0)) = (r^t, v^t)$, the path $(r^t(\mu), v^t(\mu)) \in Y$ for each fixed μ ,

$$h^t = \frac{\partial}{\partial \mu} r^t(\mu) \Big|_{\mu=0} \quad \text{and} \quad w^t = \frac{\partial}{\partial \mu} v^t(\mu) \Big|_{\mu=0}. \quad (8)$$

Differentiating (6) at $\mu = 0$, we obtain

$$dE_p(h^t, w^t) = \int_0^1 \langle D_t(h_t, w_t), (\partial_t r^t, \partial_t v^t) \rangle_{p(t)} dt,$$

where D_t denotes covariant derivative. If we set $(f^t, x^t) = (\partial_t r^t, \partial_t v^t)$ in (7) and integrate the system with initial condition $(F^0, X^0) = (0, 0)$, we get a vector field (F^t, X^t) along the path p . Then, we may rewrite the variation $dE_p(h^t, w^t)$ as

$$dE_p(h^t, w^t) = \langle (h^t, w^t), (F^t, X^t) \rangle_p. \quad (9)$$

Thus, the gradient of E at p as a functional on Y is $\nabla_Y E(p) = (F^t, X^t)$. To obtain $\nabla_{Z_M} E(p)$, we project $\nabla_Y E(p)$ orthogonally onto the tangent space of Z_M at p with respect to the inner product (5). Since the space Z_M is obtained from Y by imposing the boundary conditions $(r^0, v^0) = (r, v)$ and $(r^1, v^1) = (\bar{r}, \bar{v})$ on paths, a tangent vector to Y at p is tangent to Z_M if and only if it vanishes at $t = 0$ and $t = 1$. A simple (covariant) integration by parts argument shows that the orthogonal complement of the tangent space of Z_M in the tangent space of Y at p is formed by the covariantly linear fields (that is, smooth fields with trivial second covariant derivative) in M along p . By construction, the field (F^t, X^t) vanishes at $t = 0$. Thus, to orthogonally project (F^t, X^t) onto the tangent space of Z_M , we simply need to subtract from (F^t, X^t) the covariantly linear field that vanishes at $t = 0$ and coincides with (F^t, X^t) at $t = 1$. Again, we resort to covariant integration, this time applied to the reverse of the path p . First, construct a covariant field by integrating the everywhere zero field with initial condition (F^1, X^1) along the reverse of the path p . Reversing the path again, after integration, we obtain a parallel field (G^t, Y^t) along p whose value at $t = 1$ is (F^1, X^1) . The field (tG_t, tY_t) is covariantly linear with the desired properties. Therefore, the gradient is given by

$$\nabla_{Z_M} E(p) = \nabla_Y E(p) - (tG_t, tY_t). \quad (10)$$

Our goal is to minimize E on the path space Z_P . One possible approach is to calculate $\nabla_{Z_P} E(p)$, which would allow us to implement a gradient descent directly in Z_P . However, in this case, the calculation is quadratic in the number $m = |E|$ of edges. This is undesirable as spherical meshes used in neuroimaging often have a large number of edges. Thus, we employ an alternative numerical strategy that scales linearly with m and is based on the replacement of gradient descent in Z_P by its counterpart in Z_M followed by a closest-shape projection onto Z_P . The projection algorithm is discussed in Appendix A.

5.4 Pre-shape Geodesics

We now present an algorithm to estimate a pre-shape geodesic from (r, v) to (\bar{r}, \bar{v}) . Let $\epsilon, \delta > 0$ be small real numbers:

- (i) Initialize the search with a path $p(t) = (r^t, v^t)$ in Z_P , e.g., as described in Section 5.
- (ii) Let $(f^t, x^t) = (\partial_t r^t, \partial_t v^t)$. Using (7), integrate this field covariantly along the path p with zero initial condition. The integral field (F^t, X^t) gives the gradient $\nabla_Y E(p)$.
- (iii) Using (7), calculate the parallel transport of (F^1, X^1) along the reverse of the path p . Reverse the path and the parallel field again to obtain a parallel field (G^t, Y^t) along p . By (10), the Z_M -gradient of E is $\nabla_{Z_M} E(p) = \nabla_Y E(p) - (tG_t, tY_t)$.
- (iv) Write the modular and the directional components of the gradient $\nabla_{Z_M} E(p)$ as (h^t, w^t) , $0 \leq t \leq 1$. Update $p(t) = (r^t, v^t)$ as a path in Z_M according to (a) $\tilde{r}^t = r^t - \epsilon h^t$; (b) $\tilde{v}_i^t = v_i^t$, if $w_i^t = 0$; (c) $\tilde{v}_i^t = \cos(\epsilon \|w_i^t\|) v_i^t - \frac{\sin(\epsilon \|w_i^t\|)}{\|w_i^t\|} w_i^t$, otherwise. Note that the update of v_i^t takes place along great circles of the unit sphere in \mathbb{R}^k ensuring that each \tilde{v}_i^t is a unit vector.
- (v) Project each $(\tilde{r}^t, \tilde{v}^t)$ onto the pre-shape space P (see Appendix A) to obtain a path $(r^t, v^t)_{\text{new}}$ in P .
- (vi) Iterate the process until $\|(r^t, v^t)_{\text{new}} - (r^t, v^t)\|_p < \delta$, where $\|\cdot\|_p$ denotes the Palais norm defined in (5).

5.5 Shape Geodesic and Shape Distance

The geodesic shape metric defined in (4) involves a minimization of the pre-shape distance over the orthogonal group $O(k)$. We briefly indicate how the minimization problem can be treated. We obtain an initial estimate of

$$\hat{U} = \operatorname{argmin}_{U \in O(k)} \delta((r, v), (\bar{r}, \bar{v}U^T)), \quad (11)$$

for example, by minimizing $\|(r, v) - (\bar{r}, \bar{v}U^T)\|_{(r, v)}^2$. This is the same as minimizing $\sum_{i=1}^m b_i \|v_i - \bar{v}_i U^T\|^2 e^{r_i}$, which is equivalent to maximizing the expression $\sum_{i=1}^m (b_i e^{r_i/2} v_i) \cdot (\bar{v}_i U^T)$. This problem is analogous to the one that arises in Procrustes alignment of shapes [9] and admits a closed form solution. Once this initial estimate is obtained, a gradient search can be used to locally refine the estimation. Using the algorithm to calculate pre-shape distances, we compute the $O(k)$ -gradient numerically via finite differences using the Lie group structure of $O(k)$.

Figures 1(a) and (b) show geodesics between cortical surfaces with 122,880 edges and hippocampal surfaces with 30,720 edges. The geodesic interpolations were calculated with homogeneous elasticity coefficients $a_i = 0.3$, $b_i = 0.7$ for the cortices and $a_i = 0.15$, $b_i = 0.85$ for the hippocampal surfaces. Correspondences between the shapes were established with the direct mapping techniques of [7, 8] and the shapes were parameterized with the methods of [6]. Figure 2 further illustrates the method with the calculation of the sample Fréchet mean shape of 8 hippocampal surfaces.




Fig. 1. (a) Two cortical surfaces extracted from MR images and three intermediate stages of a geodesic interpolation with homogeneous coefficients $a_i = 0.3$ and $b_i = 0.7$; (b) a similar calculation with hippocampal surfaces with $a_i = 0.15$ and $b_i = 0.85$.



Fig. 2. The mean of 8 right hippocampal surfaces calculated with homogeneous coefficients $a_i = 0.15$ and $b_i = 0.85$. The mean shape is shown on the last panel.

6 Energy and Localization

A pre-shape geodesic $p(t) = (r^t, v^t)$, $0 \leq t \leq 1$, has parallel velocity field. In particular, it is traversed with constant speed ω , where ω is the length of p . Thus, the energy of p is $E(p) = \int_0^1 \|\partial_t p(t)\|_{p(t)}^2 dt = \omega^2$. On the other hand, we can write the energy as

$$E(p) = \sum_{i=1}^m \int_0^1 (a_i |\partial_t r_i^t|^2 e^{r_i} + b_i \|\partial_t v_i^t\|^2 e^{r_i}) dt, \quad (12)$$

which expresses $E(p)$ as a sum of the local contributions of the edges. Thus, the energy density function

$$\rho(e_i) = \frac{1}{\omega^2} \int_0^1 (a_i |\partial_t r_i^t|^2 e^{r_i} + b_i \|\partial_t v_i^t\|^2 e^{r_i}) dt \quad (13)$$

quantifies the fraction of the total deformation energy associated with the i th edge. Although geodesic distance is a global quantifier of shape difference, the energy density provides a means to measure local shape differences and identify the regions where shape similarity and divergence are most pronounced. One can further decompose the local energy into its tension and rigidity components to separately quantify the local shape differences due to stretching and bending. One may also modify ρ to a function defined on the vertex set by letting the value on a vertex be the average value of ρ on the edges incident with that vertex. Figures 3 displays the initial shapes of the geodesics in Figure 1 overlaid with color maps of the respective energy density functions.




Fig. 3. Plots of the energy density functions for the geodesics shown in Figure 1.

7 Summary and Discussion

We constructed a shape space of elastic spherical surfaces immersed in Euclidean space equipped with a family of geodesic metrics that depend on the elasticity of the surfaces. The metrics reflect the resistance offered by the shapes to deformations by stretching and bending. Although the general model is anisotropic, a special homogeneous sub-collection associated with constant tension and rigidity coefficients should be of particular interest in applications. The selection of elasticity parameters for a particular problem based on shape discrimination or other criteria requires further study. We developed an algorithm to calculate geodesics between spherical shapes and introduced energy density functions that provide a localization tool that allows us to identify the regions where shape dissimilarity is most pronounced. Energy density functions also let us characterize the nature of the local elastic deformations as due to stretching or bending.

We illustrated the applicability of the methods with the calculation of geodesic deformations between pairs of cortical and hippocampal surfaces segmented from MR images. Geodesic distances, geodesic interpolations, and models of shape variation developed in shape space enable us to quantify and visualize anatomical resemblance and divergence between individuals and across populations. We applied the methods developed in the paper to the construction of an atlas of the right hippocampus as the sample Fréchet mean shape of a group of 8 segmented hippocampi. Further statistical modeling of the shape of surfaces, extensions of the model to non-spherical shapes, and applications of the computational models to problems in neuroimaging are some of the topics to be investigated in future work.

References

1. Younes, L.: Computable elastic distance between shapes. *SIAM Journal of Applied Mathematics* **58** (1998) 565–586
2. Michor, P., Mumford, D.: Riemannian geometries on spaces of plane curves. *J. Eur. Math. Soc.* **8** (2006) 1–48
3. Klassen, E., Srivastava, A., Mio, W., Joshi, S.: Analysis of planar shapes using geodesic paths on shape manifolds. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **26** (2004) 372–383
4. Mio, W., Srivastava, A., Joshi, S.: On shape of plane elastic curves. *Int. Journal of Computer Vision* **73**(3) (2007) 307–324
5. Michor, P., Mumford, D., Shah, J., Younes, L.: A metric on shape space with explicit geodesics. *arXiv:0706.4299v1* (2007)

6. Praun, E., Hoppe, H.: Spherical parametrization and remeshing. In: ACM SIGGRAPH 2003. (2003) 340–349
7. Shi, Y., Thompson, P., Dinov, I., Osher, S., Toga, A.: Direct cortical mapping via solving partial differential equations on implicit surfaces. *Medical Image Analysis* **11**(3) (2007) 207–223
8. Shi, Y., Thompson, P., Zubizaray, G., Ross, S., Tu, Z., Dinov, I., Toga, A.: Direct mapping of hippocampal surfaces with intrinsic shape context. *NeuroImage* **36**(3) (2007) 792–807
9. Kendall, D.G.: Shape manifolds, Procrustean metrics and complex projective spaces. *Bulletin of London Mathematical Society* **16** (1984) 81–121
10. Grenander, U., Miller, M.I.: Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics* **LVI**(4) (1998) 617–694
11. Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions on Medical Imaging* **23**(8) (2004) 995–1005
12. Palais, R.S.: Morse theory on Hilbert manifolds. *Topology* **2** (1963) 299–340
13. do Carmo, M.P.: *Riemannian Geometry*. Birkhauser (1994)

A Closest-Shape Projection

Given $(r, v) \in M$ near the pre-shape space P , we need an efficient estimation of the pre-shape closest to (r, v) . Consider the $kn + 1$ residual functions $\rho_1(r, v) = 1 - G_1(r, v)$ and $\rho_{\ell,j}(r, v) = -G_{\ell,j}(r, v)$, whose vanishing is equivalent to $(r, v) \in P$. We estimate the closest pre-shape using Newton’s method to find the zeros of $H(r, v) = \frac{1}{2}\rho_1^2(r, v) + \frac{1}{2}\sum_{\ell,j}\rho_{\ell,j}^2(r, v)$. If we relax the condition that each row of v must satisfy $\|v_i\| = 1$ and just treat each v_i as an arbitrary vector in \mathbb{R}^k , then the negative gradient of H is

$$-\nabla H(r, v) = \rho_1(r, v)\nabla G_1(r, v) + \sum_{\ell,j}\rho_{\ell,j}(r, v)\nabla G_{\ell,j}(r, v), \quad (14)$$

where $\nabla G_1^r(r, v) = [1/a_1 \dots 1/a_m]^T$ is the modular component of $\nabla G_1(r, v)$ and the directional component is zero. The modular part of $\nabla G_{\ell,j}(r, v)$ and the j th column of its directional component, which is its only nonzero column, are

$$\begin{bmatrix} 0 & \frac{\varepsilon_{\ell_1}}{a_{\ell_1}}v_{\ell_1,j} & 0 & \frac{\varepsilon_{\ell_2}}{a_{\ell_2}}v_{\ell_2,j} & 0 & \frac{\varepsilon_{\ell_3}}{a_{\ell_3}}v_{\ell_3,j} & 0 \end{bmatrix}^T \quad \text{and} \quad \begin{bmatrix} 0 & \frac{\varepsilon_{\ell_1}}{b_{\ell_1}} & 0 & \frac{\varepsilon_{\ell_2}}{b_{\ell_2}} & 0 & \frac{\varepsilon_{\ell_3}}{b_{\ell_3}} & 0 \end{bmatrix}^T. \quad (15)$$

Here, the nonzero entries occur in rows ℓ_1 , ℓ_2 and ℓ_3 . Write the modular and directional parts of the gradient as $-\nabla H(r, v) = (\delta^r, \bar{\delta}^v)$. To make $\bar{\delta}_i^v$ tangential to \mathbb{S}^2 at v^i , replace it with $\delta_i^v = \bar{\delta}_i^v - (\bar{\delta}_i^v \cdot v_i)v_i$. Letting $\epsilon(r, v) = H(r, v)/\|(\delta_1, \delta_2)\|_{(r,v)}^2$, update (r, v) as follows:

$$\begin{cases} r = r + \epsilon\delta_1; \\ v_i = \cos(\epsilon\|\delta_i^v\|)v_i + \frac{\sin(\epsilon\|\delta_i^v\|)}{\|\delta_i^v\|}\delta_i^v, & \text{if } \delta_i^v \neq 0; \end{cases} \quad (16)$$

and v_i stays unchanged, otherwise. Iterate until $H(r, v)$ becomes small.

B Covariant Integration

For $(z, y) = (z, y_1, \dots, y_k) \in \mathbb{R} \times \mathbb{R}^k$, we use the subscript 0 to identify the z -coordinate and the subscript j , $1 \leq j \leq k$, for the coordinate y_j . Given $a_i, b_i > 0$, define a Riemannian metric on $\mathbb{R} \times \mathbb{R}^k$ whose metric tensor at (z, y) is $g_{00}(z, y) = a_i e^z$, $g_{jj}(z, y) = b_i e^z$, and 0, otherwise. Then, the Riemannian structure on \mathbb{E} , defined in Section 3, is isometric to the Cartesian product of these $(k+1)$ -dimensional models over $1 \leq i \leq m$. Thus, to derive the differential equation that governs covariant integration along a path in M , it suffices to derive the corresponding differential equation for covariant integration in $\mathbb{R} \times \mathbb{S}^{k-1}$ with respect to the induced metric. The Christoffel symbols of the Levi-Civita connection on $\mathbb{R} \times \mathbb{R}^k$ are $\Gamma_{00}^0 = \Gamma_{0j}^j = \Gamma_{j0}^j = 1/2$, $\Gamma_{jj}^0 = -b_i/2a_i$, $1 \leq j \leq k$, and zero otherwise. Therefore (cf. [13]), the covariant derivative of a vector field (F^t, X^t) along a path (z^t, y^t) in $\mathbb{R} \times \mathbb{R}^k$ is given by

$$\begin{cases} D_t F^t = \partial_t F^t + \frac{1}{2}(\partial_t z^t)F^t - \frac{1}{2}\frac{b_i}{a_i}[X^t \cdot \partial_t y^t] \\ D_t X^t = \partial_t X^t + \frac{1}{2}(X^t \partial_t z^t + F^t \partial_t y^t) . \end{cases} \quad (17)$$

If (z^t, y^t) is a path in $\mathbb{R} \times \mathbb{S}^{k-1}$, then fields that are tangential to $\mathbb{R} \times \mathbb{S}^{k-1}$ are those that satisfy the additional orthogonality condition $F^t \cdot y^t = 0$. Thus, if (f^t, x^t) and (F^t, X^t) are both tangential to $\mathbb{R} \times \mathbb{S}^{k-1}$, the fact that the covariant derivative of (F^t, X^t) in the submanifold $\mathbb{R} \times \mathbb{S}^{k-1}$ is (f^t, x^t) may be rephrased as $D_t F^t = f^t$ and $D_t X^t = x^t + \tau^t y^t$, where τ^t is a scalar field to be determined. Substituting in (17), we obtain

$$\begin{cases} \partial_t F^t = f^t - \frac{1}{2}\partial_t z^t F^t + \frac{1}{2}\frac{b_i}{a_i}[X^t \cdot \partial_t y^t] \\ \partial_t X^t = x^t - \frac{1}{2}(X^t \partial_t z^t + F^t \partial_t y^t) + \tau^t y^t . \end{cases} \quad (18)$$

Differentiating $X^t \cdot y^t = 0$, we get $\partial_t X^t \cdot y^t = -X^t \cdot \partial_t y^t$. From (18), it follows that $\tau^t = -X^t \cdot \partial_t y^t$, where we used the facts that $X^t \cdot y^t = 0$ and $\partial_t y^t \cdot y^t = 0$. Substituting this value of τ^t in (18), we obtain (7).

Characterization of Anatomical Shape Based on Random Walk Hitting Times

Grace Vesom¹, Nathan D. Cahill^{1,2}, Lena Gorelick³, and J. Alison Noble¹

¹ Institute of Biomedical Engineering, University of Oxford, Oxford OX3 HDQ, UK

² Research & Innovation, Carestream Health, Inc., Rochester, NY 14608, USA

³ Computer Vision Group, The Weizmann Institute of Science, Rehovot 76100, Israel

Abstract. This paper presents an implicit shape representation for describing anatomical shapes with high inter-patient variability based on the expected boundary hitting time of a random walk, which happens to be the solution to the Poisson equation. The main contribution of this paper is to test the validity of the Poisson-based mapping for learning anatomical shape variability, comparing its compactness and completeness with the commonly used Signed Distance Transform and using the liver and the caudate nucleus as examples. Based on these findings, we discuss its use as a shape prior for image segmentation.

1 Introduction

An integral part of modern disease management is treatment planning, which involves several aspects of medical image analysis, from segmenting objects of interest for volume measurement to intra-patient registration for monitoring size change or morphology. Because of the difficulty of these tasks, many segmentation or registration frameworks require a shape representation for the object in question. Shape representations can be abstractly divided into two categories: explicit and implicit. Explicit shape representations are parameterized to create ordered collections of components. Examples include Active Shape Models [1] and spherical harmonics [2]. Mostly, they define a $(N - 1)$ -dimensional boundary in a N -dimensional image space, but can include additional information such as a medial representation [3] or, like Active Appearance Models [4], intensity and texture.

An implicit representation of an N -dimensional shape is a function $f(\mathbf{x})$ in an N -dimensional space that takes on the value zero if and only if \mathbf{x} is a position on the $(N - 1)$ -dimensional shape boundary. The function is defined over the inside domain of the closed object, and sometimes, depending on the application, over the entire image domain.

Implicit shape descriptions are free of topological constraints and have local support. In recent literature, the Signed Distance Transform (SDT) has proved to be a popular implicit shape representation – attractive for ease of computation, definition over the entire image space, and ability to fit into a level-set segmentation framework. In addition to the SDT, Hong et al. [5] developed an

implicit shape representation based on an approximation to the solution of the heat equation over the shape. Their shape representation is attractive because it can be quickly computed by convolution with a Gaussian kernel. However, it does not adjust the Gaussian kernel to account for arbitrary shape boundaries.

To be able to summarize anatomical object variability, we seek a shape descriptor which can extract characteristic global object structure while preserving individual variations. In our work, we consider another implicit shape representation called the Poisson Transform (PT) and look at its application in learning and modeling anatomical variation. The PT was first used by Gorelick et al. [6] to analyze 2D shape properties on silhouettes in order to perform shape classification and retrieval. We extended that idea in [7] to include the shape exterior and used this in the context of deformable image registration. In the next section, we revisit the mathematical representation of the SDT, then provide a summary of and computational method for the PT. We use this new mapping to produce an implicit shape representation, and then analyze it for completeness and compactness in comparison to the SDT. We use Principal Component Analysis on the SDT for our training data to summarize shape variation in the same manner as [8–10]. We investigate the validity of the SDT representation for anatomical objects with highly variable shape against the PT representation. Besides the technical report by Lamecker et al. [11], no studies have assessed the compactness or completeness of liver shape description. We conclude by discussing its application to image segmentation.

2 Implicit Shape Representations

Let us define S to be the interior of the shape and assume that $S \subset \Omega$, where Ω is the entire image space. Then $\partial S = \bar{S} - S$ is the boundary of S , where \bar{S} represents the closure of the open set S .

2.1 Signed Distance Transform

The SDT, also known as the signed Euclidean distance transform, is represented here by W_{SDT} . It yields two pieces of information; the magnitude provides the Euclidean distance between a point and the closest point on ∂S , and the sign indicates whether the current point is inside S (negative) or outside S (positive). Mathematically speaking, W_{SDT} is a solution to the Eikonal equation:

$$\begin{aligned} |\nabla W_{SDT}(\mathbf{x})| &= 1 \quad \forall \mathbf{x} \in \Omega, \\ W_{SDT}(\mathbf{x}) &= 0 \quad \forall \mathbf{x} \in \partial S. \end{aligned} \tag{1}$$

Examples of the SDT for the liver and caudate nucleus can be seen in Figs. 1(a) and 1(c).

2.2 Random Walk Hitting Time

Gorelick et al. [6] presented an implicit shape representation based on the expected time for a symmetric random walk to reach the shape boundary. Based on the solution to Poisson's equation, this representation has several advantages as it is differentiable everywhere and has the interior regularity property.

In order to derive the Poisson Transform, consider a shape S defined on a 3-dimensional discrete lattice with isotropic spacing h . Now consider a particle initially at position $(x, y, z) \in S$. If the particle undergoes a symmetric random walk, then we define $U(x, y, z)$ to be the expected time (in number of steps) for the particle to reach any point on the boundary of S , given that the particle started at (x, y, z) . This expected time is also known as the "hitting time" because it represents the average time it will take the particle to "hit" the boundary. Therefore, if $(x, y, z) \in \partial S$, then $U(x, y, z) = 0$. If $(x, y, z) \in S$, then the hitting time at (x, y, z) can be related to the hitting time at each of its six neighbors by a conditioning argument:

$$U(x, y, z) = 1 + \frac{1}{6} \left(U(x+h, y, z) + U(x-h, y, z) + U(x, y+h, z) + U(x, y-h, z) + U(x, y, z+h) + U(x, y, z-h) \right). \quad (2)$$

Note that (2) is a discretized version of the Poisson equation

$$\Delta U(x, y, z) = -\frac{6}{h^2}, \quad (3)$$

with $\Delta U = U_{xx} + U_{yy} + U_{zz}$ denoting the Laplacian of U . For simplicity, we chose $h = 1$.

To use comparable notation with the SDT definition, let W_{PT} replace U to represent the symmetric random walk hitting time to the shape boundary in an open domain, and let $\mathbf{x} = (x, y, z)$ for the remainder of this paper. Like the SDT, we seek to use the sign of the metric values to define points inside and outside the boundary shape, which results in changing the sign of (3), while the magnitude of those values still provides the symmetric random walk time to a point on the boundary. Thus, for the shape interior, the *Poisson Transform* W_{PT} satisfies

$$\begin{aligned} \Delta W_{PT}(\mathbf{x}) &= 6 \quad \forall \mathbf{x} \in S, \\ W_{PT}(\mathbf{x}) &= 0 \quad \forall \mathbf{x} \in \partial S. \end{aligned} \quad (4)$$

On the exterior of the shape, if no further boundary conditions are prescribed, the random walk analogy fails and the Poisson equation has infinitely many solutions. Therefore, in order to provide an extension of this shape representation to the entirety of $\hat{\Omega}$, we need to define some sort of external boundary condition. One option is to enforce Neumann boundary conditions on $\partial\Omega$; however, this choice gives undesirable behavior of the gradient of the Poisson Transform at the

image boundaries. Instead, we choose to define an open sphere $T \supset \bar{\Omega}$ centered on the centroid of Ω , and solve the following boundary value problem:

$$\begin{aligned} \Delta W_{\text{PT}}(\mathbf{x}) &= -6 \quad \forall \mathbf{x} \in T - \bar{S}, \\ \langle \nabla W_{\text{PT}}(\mathbf{x}), \mathbf{n}(\mathbf{x}) \rangle &= 0 \quad \forall \mathbf{x} \in \partial T, \\ W_{\text{PT}}(\mathbf{x}) &= 0 \quad \forall \mathbf{x} \in \partial S, \end{aligned} \tag{5}$$

where $\mathbf{n}(\mathbf{x})$ is the outward pointing normal vector to the surface ∂T at \mathbf{x} . The Neumann boundary condition on ∂T effectively reflects the symmetric random walk in the direction normal to the sphere. Once W_{PT} has been found on T , it can be cropped to $\bar{\Omega}$.

Computing a discrete approximation of W_{PT} inside S can be done in linear time with a multigrid algorithm [6]. This algorithm can be extended to approximate W_{PT} inside $T - \bar{S}$ by appropriately discretizing the Neumann boundary conditions, as established by Cahill et al. [7]. We have found that applying two w-cycles that employ naive boundary conditions at the coarse levels (placing the boundary at the nearest coarse grid points, instead of modifying the nearby coarse equations to account for the fine, pixel-level location of the boundary) provides a good approximation. For further details on the multigrid method, see [12].

Examples of the SDT for the liver and caudate nucleus can be seen in Figs. 1(b) and 1(d).




Fig. 1. Slices from example 3D liver and caudate nucleus images of both implicit functions, inside and outside the shape

3 Application to Segmentation

In medical image segmentation, image data information is typically balanced by integrating a priori knowledge into the algorithm through a shape prior, providing a global constraint in a segmentation framework.

3.1 Shape Representation

Using a shape prior for image segmentation requires an application-specific choice of shape representation. Selecting a shape representation or descriptor is a function of the information to be extracted from an image, the signal-to-noise of the image, the anatomical object in question, and the ultimate goal of the task. For the purpose of segmentation, the shape descriptor also subsequently affects the dissimilarity measure for comparing shapes and method of integration for a model into a segmentation framework. Golland et al. [13] found that the choice of shape representation changed their classifier function results by about 15%, independent of shape analysis and alignment.

3.2 Shape Analysis Methods

With highly variable anatomical shapes, large numbers of samples, or training shapes, are required to correctly portray a shape population. Shape models are used to efficiently summarize the training shape space.

Much research looks for the best shape analysis method to summarize the shape space, while some methods never question the validity of its use, such as with Principal Component Analysis (PCA). That method is utilized to reduce the dimensionality of the training shape space, while preserving the most variation. PCA is often challenged because of its sensitivity to low sample size, as well as the assumption of a Gaussian uni-modal distribution for the features measured in the resulting model. In the case of the SDT and PT, the shape representation space is not closed under linear operations, but PCA is still often employed. Shape analysis alternatives include independent component analysis [14], principal factor analysis [15], and support vector machines [13]. Additionally, methods to work around obstacles in PCA have shown some success. They transform training data into a feature space via Mercer kernels [16], or into a vector space using Fourier approximation [17] or the logarithm of odds ratio [18], where linear operations of shape representation transformations are closed.

With the SDT, PCA has been consistently employed for shape analysis for the purpose of learning shape variability in creating a statistical shape model. Notwithstanding, Hong et al. [5] used their integral kernel shape representation as a template with a deformation field for segmentation, removing the need for any shape analysis.

3.3 Implicit Anatomical Shape Priors

In image segmentation of highly variable shapes, it is essential to integrate a priori information as a global constraint with data-driven energy functions. We concentrate here on statistical shape models (SSMs), as opposed to deformable shape models or classifiers.

Leventon et al. [8] employed SSMs from implicit shape representations, using the SDT to represent shapes and PCA to summarize the shape space to create a shape prior in image segmentation. Thereafter, literature involving implicit

shape models concentrates on shape alignment prior to PCA and shape prior integration into a region-based level-set segmentation framework. Rousson et al. [10] used the sum of squared distances as a dissimilarity measure between the prior and target, while minimizing the rotation and translation pose parameters during the level set evolution. Tsai et al. [9] similarly proposed an alignment method, but also included scale as a pose parameter and explored three energy functionals, optimizing the functionals via gradient descent. Yang et al. [19] took another step by introducing neighbor priors, developing a hierarchical multi-object segmentation. Accounting for neighboring structures in the shape prior, however, gave rise to concerns about closedness of linear combination of signed distance functions. Neighboring structures were modeled by the difference of their representative SDTs, however any linear combination of these differences may not represent neighboring structures as mutually exclusive.

3.4 Analysis of Implicit Shape Representations SDT and PT

While any dimensionality reduction method can arguably be more suitable for shape analysis, we wish to show that a shape descriptor that suppresses redundancy in variability, yet elicits distinguishing global features, will result in more meaningful factors. PCA is the most widely used form of shape analysis in medical imaging because of its ease of computation and interpretation. For these reasons, we have thus chosen to use PCA for a comparative shape analysis study on the SDT and PT.

The nature of the PT allows us to combine global and local shape extraction – we can see a common shape for each structure, while local variations are preserved closer to the boundary (Figs. 3(b) and 2(b)). Consequently, we believe PT is best suited to represent objects with high shape variability, for which the liver and caudate nucleus are well known and well studied.




Fig. 2. Inside level sets of the liver from SDT and PT




Fig. 3. Inside level sets of the caudate nucleus from SDT and PT

Consider the case of representing the shape of the liver. Although liver shape has consistent convexities, concavities, and ridges, it is still very highly variable across the population. It is also adjacent to structures of similar density, making it extremely difficult to segment in Computed Tomography (CT) images. Liver image segmentation has been a well-established problem in medical imaging analysis. A variety of published methods range from deformable models using the evolutionary algorithm [20] and data-driven methods [21] to graph-cuts [22], demonstrating the difficulty of the task.

The shape of the caudate nucleus is often investigated in order to identify brain abnormalities. An irregular shape in longitudinal studies and intra-patient asymmetric shape can reveal developmental disorders, such as schizophrenia [23].

In 3D, we computed the SDT and PT for 33 liver images and 20 caudate nucleus images. Slices through the mapping for each transform and class are shown in Fig. 1, and 3D renderings of the inner level sets are shown in Figs. 3 and 2.

3.5 The Data

Our data has been obtained from the MICCAI 3D Segmentation Challenge [24] and the Churchill Hospital, part of the Oxford Radcliffe Hospitals NHS Trust in Oxford, UK. We used the segmented training data sets from MICCAI, and we manually segmented the additional liver data sets subject to the approval of a clinician. All binary images were rescaled to physical Cartesian coordinates and smoothed to remove aliasing effects. The liver surfaces were aligned as binary images using the Principal Axis Transformation [25], while the caudate nucleus shapes were aligned using first-order moments, with a few slight manual rotations.

3.6 Analysis

For each case, we ran PCA over the entire data set, Figs. 4(a) and 5(a), and then for randomly chosen sets in multiples of 5, to show the number of modes needed to summarize a cumulative variance threshold of 95% given the changing number of training shapes, Figs. 4(b) and 5(b).




Fig. 4. Compactness and completeness of the SDT and PT for the shape of the liver

We defined compactness of a shape description as requiring as few parameters as possible and show relative compactness here in comparison to the SDT. From Figs. 4(a) and 5(a), we see that to summarize 95% of shape variance for the liver and the caudate nucleus, the SDT requires 15 and 8 modes, while the PT requires only 5 and 3, respectively – approximately a 3-fold improvement in both cases.

3.7 Discussion

With a large number of training data, we would expect that the number of modes required should approach a constant value, as we would expect in a complete shape model. If the number of modes required grows with the number of training data, then it can be inferred that every randomly chosen set of 5 training data added were somehow orthogonal to the existing set, which is improbable. Then PCA might be incapable of accurately reducing the dimensionality of a shape space spanned by the SDT on liver and caudate nucleus shapes. As another possibility, the SDT might be an inadequate representation of highly variable




Fig. 5. Compactness and completeness of the SDT and PT for the shape of the caudate nucleus

anatomical objects and create a shape space where common factors are difficult to resolve.

For a SSM based on the liver for image segmentation, Lamecker et al. [11] used a subsampled geometric mesh representation. Mesh point values, like all values in a SDT-map, are directly affected by any noise on the shape surface. Lamecker report compactness and completeness curves consistent with that of the signed distance function in their liver SSM analysis, and required 18 modes to summarize 95% variance for 33 training shapes and 21 modes for 42 training shapes.

In contrast, the PT completeness curve flattens quickly as the number of training shapes grows, converging to a fixed number of modes, as we should expect in a complete model (Fig. 4(b)). Through our comparative PCA study we can see the PT is a relatively more compact and complete shape descriptor than SDT.

For high dimensional data and low sample size data sets, PCA is argued to be incomplete. That is to say some actual shapes may be unrealizable from the basis, which is chosen for low-approximation while summarizing the most variance. Yet one can argue that the shape analysis results are dependent on what has been extracted from the original data into the shape representation. Boundary representations give information about the object delineation, and SDT-maps propagate boundary variations beyond the neighborhood of occurrence, creating superfluous variations throughout the image domain. However, the PT-maps

reveal global shape information from a data set yet adapt to individual object boundaries (see Figs. 1–2).

4 Conclusion and Future Work

In this paper, we introduced using a new implicit shape representation called the Poisson Transform, an extrapolation of Gorelick’s innovative approach to shape characterization. In our comparative dimensionality reduction study, we showed that the PT gives approximately a 3-fold improvement over the SDT in the number of principal components required to represent 95% of the liver and caudate nucleus data set variance. The PT is capable of local support on the boundary for individual instances but substantially stable against those variations to give rise to a characteristic object shape across a data set. This balance in description enables optimized learning for highly variable anatomical objects. We are currently working on methods of integrating a shape model based on the PT into an image segmentation framework and overcoming obstacles mentioned in [13, 18].

Acknowledgments

The authors would like to thank Professor J. M. Brady for his insightful discussions; Tobias Heimann, Martin Styner, and Bram van Ginneken for allowing us to use their data; and Dr. Joanne Brady from the Churchill Hospital in Oxford for selecting patients and overseeing manual segmentations.

References

1. Cootes, T., Taylor, C., Cooper, D., Graham, J., et al.: Active shape models-their training and application. *Computer Vision and Image Understanding* **61**(1) (1995) 38–59
2. Brechbuhler, C., Gerig, G., Kubler, O.: Parametrization of closed surfaces for 3-D shape description. *Computer Vision and Image Understanding* **61**(2) (1995) 154–170
3. Styner, M., Gerig, G.: Medial models incorporating object variability for 3D shape analysis. *Information Processing in Medical Imaging* (2001) 502–516
4. Cootes, T., Edwards, G., Taylor, C.: Active appearance models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **23**(6) (2001) 681–685
5. Hong, B.W., Prados, E., Soatto, S., Vese, L.: Shape Representation based on Integral Kernels: Application to Image Matching and Segmentation. In: *Proc. CVPR*. (2006) I: 833–840
6. Gorelick, L., Galun, M., Sharon, E., Basri, R., Brandt, A.: Shape Representation and Classification Using the Poisson Equation. *Transactions on Pattern Analysis and Machine Intelligence* **28**(12) (2006) 1991–2005
7. Cahill, N.D., Vesom, G., Gorelick, L., Brady, J., Noble, J.A., Brady, J.M.: Investigating Implicit Shape Representations for Alignment of Livers from Serial CT Examinations. In: *Proc. International Symposium on Biomedical Imaging*. (May 2008)

8. Leventon, M.E., Grimson, W.E.L., Faugeras, O.: Statistical shape influence in geodesic active contours. In: Computer Vision and Pattern Recognition, 2000. Proceedings. IEEE Conference on. Volume 1. (2000) 316–323
9. Tsai, A., Yezzi Jr, A., Wells, W., Tempany, C., Tucker, D., Fan, A., Grimson, W., Willsky, A.: A shape-based approach to the segmentation of medical imagery using level sets. Medical Imaging, IEEE Transactions on **22**(2) (2003) 137–154
10. Rousson, M., Paragios, N., Deriche, R.: Implicit Active Shape Models for 3D Segmentation in MR Imaging. In: MICCAI 2004. Volume 3216., Springer Berlin / Heidelberg (September 2004) 209–216
11. Lamecker, H., Lange, T., Seebass, M.: Segmentation of the Liver using a 3D Statistical Shape Model. Technical report, Zuse Institute Berlin (April 2004)
12. Trottenberg, U., Oosterlee, C.W., Schüller: Multigrid. Academic Press (2001)
13. Golland, P., Grimson, W., Shenton, M., Kikinis, R.: Detection and analysis of statistical differences in anatomical shape. Medical Image Analysis **9**(1) (February 2005) 69–86
14. Hyvarinen, A., Oja, E.: Independent component analysis: algorithms and applications. Neural Networks **13**(4-5) (2000) 411–430
15. Ballester, M., Linguraru, M., Aguirre, M., Ayache, N.: On the adequacy of principal factor analysis for the study of shape variability. Proc. SPIE **5747** (2005) 1392–1399
16. Cremers, D., Kohlberger, T., Schnörr, C.: Shape statistics in kernel space for variational image segmentation. Pattern Recognition **36**(9) (September 2003) 1929–1943
17. Srivastava, A., Joshi, S.H., Mio, W., Liu, X.: Statistical shape analysis: clustering, learning, and testing. Pattern Analysis and Machine Intelligence, IEEE Transactions on **27**(4) (2005) 590–602
18. Pohl, K.M., Fisher, J., Shenton, M., McCarley, R., Grimson, W., Kikinis, R., Wells, W.: Using the logarithm of odds to define a vector space on probabilistic atlases. Medical Image Analysis **11**(5) (October 2007) 465–477
19. Yang, J., Staib, L.H., Duncan, J.S.: Neighbor-constrained segmentation with level set based 3-D deformable models. Medical Imaging, IEEE Transactions on **23**(8) (2004) 940–948
20. Heimann, T., Munzing, S., Meinzer, H., Wolf, I.: A Shape-Guided Deformable Model with Evolutionary Algorithm Initialization for 3D Soft Tissue Segmentation. In: Information Processing in Medical Imaging. Volume 4584. Springer Berlin/Heidelberg (2007) 1–12
21. Soler, L., Delingette, H., Malandain, G., Montagnat, J., Ayache, N., Koehl, C., Dourthe, O., Malassagne, B., Smith, M., Mutter, D., et al.: Fully automatic anatomical, pathological, and functional segmentation from CT scans for hepatic surgery. Computer Aided Surgery **6**(3) (2001) 131–142
22. Boykov, Y., Funka-Lea, G.: Graph Cuts and Efficient N-D Image Segmentation. International Journal of Computer Vision **70**(2) (November 2006) 109–131
23. Levitt, J., Westin, C., Nestor, P., Estepar, R., Dickey, C., Voglmaier, M., Seidman, L., Kikinis, R., Jolesz, F., McCarley, R., et al.: Shape of caudate nucleus and its cognitive correlates in neuroleptic-naïve schizotypal personality disorder. Biological Psychiatry **55**(2) (2004) 177–184
24. van Ginneken, B., Heimann, T., Styner, M.: 3D Segmentation in the Clinic: A Grand Challenge. In van Ginneken, B., Heimann, T., Styner, M., eds.: 3D Segmentation in the Clinic: A Grand Challenge. (2007) 7–15
25. Alpert, N.M.: The Principal Axes Transformation—A Method for Image Registration. Journal of Nuclear Medicine **31**(10) (October 1990) 1717–1722

Tiling Manifolds with Orthonormal Basis

Moo K. Chung¹², Anqi Qiu⁴, Brendon, M. Nacewicz²,
Seth Pollak³, Richard J. Davidson²³

¹Department of Biostatistics and Medical Informatics,

²Waisman Laboratory for Brain Imaging and Behavior,

³ Department of Psychology and Psychiatry

University of Wisconsin, Madison, USA

⁴Division of Bioengineering, Faculty of Engineering

National University of Singapore, Singapore

`mkchung@wisc.edu`

Abstract. One main obstacle in building a sophisticated parametric model along an arbitrary anatomical manifold is the lack of an easily available orthonormal basis. Although there are at least two numerical techniques available for constructing an orthonormal basis such as the Laplacian eigenfunction approach and the Gram-Smidt orthogonalization, they are computationally not so trivial and costly. We present a relatively simpler method for constructing an orthonormal basis for an arbitrary anatomical manifold. On a unit sphere, a natural orthonormal basis is the spherical harmonics which can be easily computed. Assuming the manifold is topologically equivalent to the sphere, we can establish a smooth mapping ζ from the manifold to the sphere. Such mapping can be obtained from various surface flattening techniques. If we project the spherical harmonics to the manifold, they are no longer orthonormal. However, we claim that there exists an orthonormal basis that is the function of spherical harmonics and the spherical mapping ζ .

The detailed step by step procedures for the construction is given along with the numerical validation using amygdala surfaces as an illustration. As an application, we propose the pullback representation that reconstructs surfaces using the orthonormal basis obtained from an average template. The pullback representation introduces less inter-subject variability and thus requires far less number of coefficients than the traditional spherical harmonic representation. The source code used in the study is freely available at

<http://www.stat.wisc.edu/~mchung/research/amygdala>.

1 Introduction

We present a novel orthonormal basis construction method for an arbitrary anatomical surface that is topologically equivalent to a sphere. The method avoids the well known Gram-Smidt orthogonalization procedure [7], which is inefficient for high resolution polygonal meshes. In order to perform the Gram-Smidt orthogonalization as described in [7], for a surface mesh with n vertices,

we need to perform the Choleski decomposition as well as the inversion of matrix of size $n \times n$. For a cortical mesh generated with FreeSurfer [6], n can easily reach up to 200000.

On the other hand, Qiu et al. [15] constructed an orthonormal basis as the eigenfunctions of the Laplace-Beltrami operator in a bounded regions of interest (ROI) on a cortical surface (Figure 3). The finite element method (FEM) is used to numerically construct the orthonormal basis by solving a system of large linear equations. The weakness of the FEM approach is the computational burden of inverting a matrix of size $n \times n$. One advantage of the eigenfunction approach is that since the eigenfunctions and eigenvalues are directly related to the Laplace-Beltrami operator, it is trivial and geometrically intuitive to construct the heat kernel analytically and perform a various heat kernel smoothing based modeling [4].

We propose a completely different method that avoids the computational bottleneck by using a conceptually different machinery. We assume an arbitrary anatomical surface to be topologically equivalent to a sphere. Then using a smooth mapping ζ obtained from a surface flattening technique, we project the spherical harmonics to the anatomical surface. Obviously the projected spherical harmonics will no longer be orthonormal. However, if we correct the metric distortion introduced from the surface flattening, we may able to make the projected spherical harmonics orthonormal somehow. This is the basic idea behind our new proposed method. For the surface flattening, we present a new method that treats the mapping ζ as the geodesic path of the heat equilibrium state.

As an application of the proposed technique, we present the novel *pullback representation* for parameterizing anatomical boundaries that outperforms the traditional spherical harmonic (SPHARM) representation [2] [3] [8] [16] [17]. We claim that our proposed representation has far less intersubject variability in the estimated parameters than SPHARM and converges faster to the true boundary with less number of basis.

2 Methods

It is assumed that the anatomical boundary \mathcal{M} is a smooth 2-dimensional Riemannian manifold parameterized by two parameters. The one-to-one mapping ζ from point $\mathbf{p} = (p_1, p_2, p_3)' \in \mathcal{M}$ to $\mathbf{u} = (u_1, u_2, u_3)' \in S^2$, a unit sphere, can be obtained from various surface flattening techniques such as conformal mapping [1] [8] [9], quasi-isometric mapping [18], area preserving mapping [2] [16] [17] and the deformable surface algorithm [13]. Since the conformal mapping tend to introduce huge area distortion, most spherical harmonic literature tend to use area preserving mapping [2] [16] [17].

In this paper, we present a new flattening technique via the geodesic trajectory of the equilibrium state of heat diffusion. The proposed flattening technique is numerically simpler than any other available methods and does not require optimizing a cost function. The methodology is illustrated using the 47 amyg-




Fig. 1. The diffusion equation with a heat source (amygdala) and a heat sink (enclosing sphere) corresponds. After sufficient amount of diffusion, the heat equilibrium state is reached. By tracing the geodesic path from the heat source to the heat sink using the geodesic contours, we obtain a smooth mapping ζ .




Fig. 2. Amygdala surface flattening is done by tracing the geodesic path of the heat equilibrium state. The numbers corresponds to the different the geodesic contours. For simple shapes like amygdala, 5 to 10 contours are sufficient for tracing the geodesic path.

dala binary segmentation obtained from the 3-Tesla magnetic resonance images (MRI).

High resolution anatomical MRI were obtained using a 3-Tesla GE SIGNA scanner with a quadrature head coil. Details on image acquisition parameters are given in [14]. MRIs are reoriented to the pathological plane for optimal segmentation and comparison with an atlas. This global alignment guarantee that amygdala are approximately aligned in the same orientation.

Manual amygdala segmentation was done by a trained expert and the reliability of the manual segmentation was validated by two raters on 10 amygdala resulting in interclass correlation of 0.95 and the intersection over the union of 0.84 [14]. Afterwards a marching cubes algorithm was used to extract the boundary of the binary segmentation as a triangle mesh with approximately 2000-3000 vertices. The amygdala surface is then mapped onto a sphere using the new flattening algorithm.

2.1 Diffusion-Based Surface Flattening

Given an amygdala binary segmentation \mathcal{M}_a , we put a larger sphere \mathcal{M}_s that encloses the amygdala (Figure 1 left). The amygdala and sphere serve as Dirichlet boundary conditions for solving the Laplace equation. The amygdala is assigned the value 1 while the enclosing sphere is assigned the value -1, i.e.

$$f(\mathcal{M}_a, \sigma) = 1, f(\mathcal{M}_s, \sigma) = -1 \quad (1)$$

for all σ . The amygdala and the sphere serve as a heat source and a heat sink respectively. Then we solve an isotropic diffusion

$$\frac{\partial f}{\partial \sigma} = \Delta f \quad (2)$$

within the empty space bounded by the amygdala and the sphere. Δ is the 3D Laplacian. After enough diffusion, the system reaches the heat equilibrium state where the additional diffusion does not make any difference in the heat distribution (Figure 1 middle). Once we obtained the equilibrium state, we trace the geodesic path from the heat source to the heat sink for every mesh vertices. The trajectory of the geodesic path provides a smooth mapping from the amygdala surface to the sphere. The geodesic path can be easily traced by constructing geodesic contours that correspond to the level set of the equilibrium state (Figure 1 right). Then the geodesic path is constructed by finding the shortest distance from one contour to the next and iteratively connecting the path together. Figure 2 shows the process of flattening using five contours corresponding to the temperature 0.6, 0.2, -0.2, -0.6, -1.0.

Although we did not apply our flattening technique to other anatomical objects, the proposed method can be applied to more complex object than the amygdala. At the equilibrium state, we no longer has change in heat change over time, i.e. $\frac{\partial f}{\partial \sigma} = 0$, so we have the Laplace equation

$$\Delta f = 0$$

with the same boundary condition. The Laplace equation has been previously used to trace the distance between outer and inner cortical surfaces and to compute cortical thickness [10] [12] [19]. Since the solution to the Laplace equation with the boundary condition (1) is unique even for highly convoluted and folded structures, the geodesic path will be uniquely defined.

2.2 Orthonormal basis in two sphere S^2

Suppose a unit sphere S^2 is represented as a high resolution triangle mesh consisting of the vertex set $\mathcal{V}(S^2)$. We have used an almost uniformly sampled mesh with 2562 vertices and 5120 faces. Let us parameterize coordinates $\mathbf{u} \in S^2$ with parameters θ, φ :

$$(u_1, u_2, u_3) = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta),$$

where $(\theta, \varphi) \in \mathcal{N} = [0, \pi] \otimes [0, 2\pi)$. The polar angle θ is the angle from the north pole and φ is the azimuthal angle. The orthonormal basis on the unit sphere is given by the eigenfunctions of

$$\Delta f + \lambda f = 0,$$

where Δ is the spherical Laplacian. The eigenfunction Y_{lm} corresponding to the eigenvalue $l(l+1)$ is called the spherical harmonic of degree l and order m [5]. With respect to the inner product

$$\langle f, g \rangle_{S^2} = \int_{S^2} f(\mathbf{u})g(\mathbf{u}) d\mu(\mathbf{u}), \quad (3)$$

with measure $d\mu(\mathbf{u}) = \sin\theta d\theta d\varphi$, Y_{lm} form the orthonormal basis in $L^2(S^2)$, the space of square integrable functions on S^2 , i.e.

$$\langle Y_{lm}, Y_{l'm'} \rangle_{S^2} = \delta_{ll'} \delta_{mm'}. \quad (4)$$

The inner product can be numerically computed as the Riemann sum over mesh vertices as

$$\langle Y_{lm}, Y_{l'm'} \rangle_{S^2} \approx \sum_{\mathbf{u}_j \in \mathcal{V}(S^2)} Y_{lm}(\mathbf{u}_j) Y_{l'm'}(\mathbf{u}_j) D_{S^2}(\mathbf{u}_j), \quad (5)$$

where $D_{S^2}(\mathbf{u}_j)$ is the discrete approximation of $d\mu(\mathbf{u})$. Let $T_{\mathbf{u}_j}^1, T_{\mathbf{u}_j}^2, \dots, T_{\mathbf{u}_j}^{j_m}$ be the area of triangles containing the vertex \mathbf{u}_j . Then we estimate $D_{S^2}(\mathbf{u}_j)$ as

$$D_{S^2}(\mathbf{u}_j) = \frac{1}{3} \sum_{k=1}^{j_m} T_{\mathbf{u}_j}^k. \quad (6)$$

The discrete approximation (6) defines the area of triangles at a mesh vertex. The factor $1/3$ is chosen in such a way that

$$\sum_{\mathbf{u}_j \in \mathcal{V}(S^2)} D_{S^2}(\mathbf{u}_j) = 12.5514 = 4 \cdot 3.1378,$$

analogous to the relationship

$$\int_{S^2} d\mu(\mathbf{p}) = 4\pi.$$

The discrepancy between the integral and its discrete counter part is due to the mesh resolution and it should become smaller as the mesh resolution increases.

Based on the proposed discretization scheme, we have computed the inner product (5) for all degrees $0 \leq l, l' \leq 20$. Figure 3 (left) shows the inner products for every possible pairs. Since for up to the k -th degree, there are total $(k+1)^2$ basis functions, we have total 441^2 possible inner product pairs, which is displayed as a matrix. For the diagonal terms, we obtained 0.9988 ± 0.0017 while for the off-diagonal terms, we have obtained 0.0000 ± 0.0005 indicating our basis and the discretization scheme is orthonormal with two decimal accuracy.

2.3 Orthonormal basis on manifold \mathcal{M}

For $f, g \in L^2(\mathcal{M})$, the orthonormality is defined with respect to the inner product

$$\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} f(\mathbf{p})g(\mathbf{p}) d\mu(\mathbf{p}).$$

Using the spherical harmonics in S^2 , it is possible to construct an orthonormal basis in \mathcal{M} numerically without the computational burden of solving the large matrix inversion associated with the eigenfunction method or the Gram-Smidt orthogonalization. Since the spherical harmonics are orthonormal in S^2 and, the manifolds S^2 and \mathcal{M} can be deformed to each other by the mapping ζ , one would guess that the orthonormal basis in \mathcal{M} can be obtained somehow using the spherical harmonics. Surprisingly this guess is not wrong as we will show in this section.

For $f \in L^2(S^2)$, let us define the *pullback operation* $*$ as

$$\zeta^* f = f \circ \zeta.$$

While f is defined on S^2 , the pullbacked function $\zeta^* f$ is defined on \mathcal{M} . The schematic of the pull back operation is given in Figure 6 (a). Then even though we do not have orthonormality on the pullbacked spherical harmonics, i.e.,

$$\langle \zeta^* Y_{lm}, \zeta^* Y_{l'm'} \rangle_{\mathcal{M}} \neq \delta_{ll'} \delta_{mm'},$$

we can make them orthonormal by using the Jacobian determinant of the mapping ζ somehow.

Consider the Jacobian J_{ζ} of the mapping $\zeta : \mathbf{p} \in \mathcal{M} \rightarrow \mathbf{u} \in S^2$ defined as

$$J_{\zeta} = \frac{\partial \mathbf{u}(\theta, \varphi)}{\partial \mathbf{p}(\theta, \varphi)}.$$

For functions $f, g \in L^2(S^2)$, we have the following change of variable relationship:

$$\langle f, g \rangle_{S^2} = \int_{\mathcal{M}} \zeta^* f(\mathbf{p}) \zeta^* g(\mathbf{p}) |\det J_{\zeta}| d\mu(\mathbf{p}). \quad (7)$$

Similarly we have the inverse relationship given as

$$\langle \zeta^* f, \zeta^* g \rangle_{\mathcal{M}} = \int_{S^2} f(\mathbf{u})g(\mathbf{u}) |\det J_{\zeta}^{-1}| d\mu(\mathbf{u}). \quad (8)$$

By letting $f = Y_{lm}$ and $g = Y_{l'm'}$ in (7), we obtain

$$\delta_{ll'} \delta_{mm'} = \int_{\mathcal{M}} \zeta^* Y_{lm} \zeta^* Y_{l'm'} |\det J_{\zeta}| d\mu(\mathbf{p}) \quad (9)$$

Equation (9) demonstrates that functions

$$Z_{lm} = |\det J_{\zeta}|^{1/2} \zeta^* Y_{lm} \quad (10)$$

are orthonormal in \mathcal{M} . We will refer l as degree and m as order of the basis function. Using the Riesz-Fischer theorem [11], it is not hard to show that Z_{lm} form a complete basis in $L^2(\mathcal{M})$.




Fig. 3. Left: inner products of eigenfunctions of the Laplacian for every pairs [15]. The pairs are rearranged from low to high degree. Right: representative eigenfunctions Ψ_j on the left amygdala template surface obtained by solving $\Delta\Psi_j + \lambda_j\Psi_j = 0$.

2.4 Numerical Implementation

Although the expression (10) provides a nice analytical form for an orthonormal basis for an arbitrary manifold \mathcal{M} , it is not practical. If one want to use the basis (10), the Jacobian determinant needs to be numerically estimated somehow. We present a new discrete estimation technique for the surface Jacobian determinant that avoids estimating unstable spatial derivative estimation.

The Jacobian determinant J_ζ of the mapping ζ can be expressed in terms of the Riemannian metric tensors associated with the manifolds S^2 and \mathcal{M} . Consider determinants $\det g_{S^2}$ and $\det g_{\mathcal{M}}$ of the Riemannian metric tensors associated with the parameterizations $\mathbf{u}(\theta, \varphi)$ and $\mathbf{p}(\theta, \varphi)$ respectively. Note that the integral of the area elements $\sqrt{\det g_{S^2}}$ and $\sqrt{\det g_{\mathcal{M}}}$ with respect to the parameter space \mathcal{N} gives the total area of the manifolds, i.e.

$$\int_{\mathcal{N}} \sqrt{\det g_{S^2}} d\mu(\theta, \varphi) = 4\pi, \quad \int_{\mathcal{N}} \sqrt{\det g_{\mathcal{M}}} d\mu(\theta, \varphi) = \mu(\mathcal{M}).$$

Then we have the relationship

$$|\det J_{\zeta^{-1}}| = \frac{\sqrt{\det g_{\mathcal{M}}}}{\sqrt{\det g_{S^2}}}, \quad |\det J_\zeta| = \frac{\sqrt{\det g_{S^2}}}{\sqrt{\det g_{\mathcal{M}}}}.$$

Note that the Jacobian determinant $\det J_\zeta$ measures the amount of contraction or expansion in the mapping ζ from \mathcal{M} to S^2 . So it is intuitive to have this quantity to be expressed as the ratio of the area elements. Consequently the discrete estimation of the Jacobian determinant at mesh vertex $\mathbf{u}_j = \zeta(\mathbf{p}_j)$ is obtained as

$$|\det J_\zeta| \approx \frac{D_{S^2}(\mathbf{u}_j)}{D_{\mathcal{M}}(\mathbf{p}_j)}.$$




Fig. 4. Left: inner products of spherical harmonics computed using formula (3) for every pairs. The pairs are rearranged from low to high degree and order. There are total $(20 + 1)^2 = 441$ possible pairs for up to degree 20. Right: representative orthonormal basis Z_{lm} on the left amygdala template surface.

Then our orthonormal basis is given by

$$Z_{lm}(\mathbf{p}_j) = \sqrt{\frac{D_{S^2}(\zeta(\mathbf{p}_j))}{D_{\mathcal{M}}(\mathbf{p}_j)}} \zeta^* Y_{lm}(\mathbf{p}_j). \quad (11)$$

The numerical accuracy can be determined by computing the inner product

$$\begin{aligned} \langle Z_{lm}, Z_{l'm'} \rangle_{\mathcal{M}} &\approx \sum_{\mathbf{p}_j \in \mathcal{V}(\mathcal{M})} Z_{lm}(\mathbf{p}_j) Z_{l'm'}(\mathbf{p}_j) D_{\mathcal{M}}(\mathbf{p}_j). \\ &= \sum_{\mathbf{p}_j \in \mathcal{V}(\mathcal{M})} \zeta^* Y_{lm}(\mathbf{p}_j) \zeta^* Y_{l'm'}(\mathbf{p}_j) D_{S^2}(\zeta(\mathbf{p}_j)) \\ &= \sum_{\mathbf{u}_j \in \mathcal{V}(S^2)} Y_{lm}(\mathbf{u}_j) Y_{l'm'}(\mathbf{u}_j) D_{S^2}(\mathbf{u}_j) \\ &= \langle Y_{lm}, Y_{l'm'} \rangle_{S^2} \end{aligned}$$

Since this is tautology, the order of the numerical accuracy in Z_{lm} is identical to that of spherical harmonics given in the previous section. There is no need for additional validation other than given in the previous section. Hence we conclude that our basis is in fact orthonormal within two decimal accuracy. Figure 4 shows the result of our numerical procedure applied to the average amygdala surface template. The template surface is constructed by averaging the surface using the spherical harmonic correspondence given in [3].

We have also constructed an orthonormal basis on a cortical surface with more than 40000 mesh vertices (Figure 5). The diagonal elements in the inner product matrix are 0.9999 ± 0.0001 indicating that our basis is orthonormal within three decimal accuracy. As the mesh resolution increases, we expect to have increased accuracy. The proposed orthonormal basis construction methods




Fig. 5. Orthonormal basis Z_{lm} on a cortical surface. The basis is projected on a sphere to show how the nonuniformity of the Jacobian determinant is affecting the spherical harmonics Y_{lm} . The color scale is thresholded at ± 0.003 for better visualization.

avoids inverting matrix of size larger than 40000×40000 associated with the eigenfunction approach and the Gram-Smidt orthogonalization process.

Although the pattern of tiling in the eigenfunction approach (Figure 3) and the pullback based method (Figure 4) looks different, it can be shown that they are actually linearly dependent.

3 Application: Pullback Representation

As an application of the proposed orthonormal basis construction, we present a new variance reducing Fourier Series representation that outperforms the traditional spherical harmonic representation [2] [3] [8] [16] [17]. We will call this method as the *pullback representation*.

The spherical harmonic (SPHARM) representation models the surface coordinates with respect to a unit sphere as

$$\mathbf{p}(\theta, \varphi) = \sum_{l=0}^k \sum_{m=-l}^l \mathbf{p}_{lm}^0 Y_{lm}(\theta, \varphi) \quad (12)$$

where $\mathbf{p}_{lm}^0 = \langle \mathbf{p}, Y_{lm} \rangle_{S^2}$ are spherical harmonic coefficients, which can be viewed as random variables. The coefficients are estimated using the *iterative residual fitting algorithm* [3] that breaks a larger least squares problem into smaller ones in an iterative fashion. The MATLAB code for performing the iterative residual fitting algorithm for arbitrary surface mesh is given in <http://www.stat.wisc.edu/~mchung/softwares/weighted-SPHARM/weighted-SPHARM.html>. Note that all MRIs were reoriented to the pathological plane guaranteeing an approximate global alignment before the surface flattening to increase the robustness of the coefficient estimation.




Fig. 6. Left: schematic showing how the pullback operation $*$ is working. Point $\mathbf{p} \in \mathcal{M}$ is mapped to $\mathbf{u} \in S^2$ via our new flattening technique. As an illustration $f = Y_{3,2} + 0.6Y_{2,1}$ is plotted on S^2 . The function f is pulled back onto \mathcal{M} by ζ . Right: sample standard deviation of Fourier coefficients of for 47 subjects plotted over the index of basis. In average, the traditional SPHARM representation (black) has 88% more variability than the pull back method (red).

The shortcoming of the spherical harmonic representation is that the reconstruction is respect to a unit sphere that is not geometrically related to the original anatomical surface. On the other hand, the pullback representation will reconstruct the surface with respect to the average template surface reducing substantial amount of variability compared to SPHARM.

In the pullback representation, we represent the surface coordinates with respect to the template surface \mathcal{M} as

$$\mathbf{p}(\theta, \varphi) = \sum_{l=0}^k \sum_{m=-l}^l \mathbf{p}_{lm}^1 Z_{lm}(\theta, \varphi) \quad (13)$$

with $\mathbf{p}_{lm}^1 = \langle \mathbf{p}, Z_{lm} \rangle_{\mathcal{M}}$. Then we claim that the pullback representation has smaller variance in the estimated coefficients so that

$$\text{Var}(\mathbf{p}_{lm}^1) \leq \text{Var}(\mathbf{p}_{lm}^0). \quad (14)$$

The equality in (14) is obtained when the template \mathcal{M} becomes the unit sphere, in which case the spherical mapping ζ collapses to the identity, and the inner products coincide. We have computed the sample standard deviation of Fourier coefficients for 47 subjects using the both representations. In average, the SPHARM contains 88% more intersubject variability compared to the pullback representation (Figure 6 right). This implies that SPHARM is an inefficient representation and requires more number of basis to represent surfaces compared to the pullback method.

Although the pullback method is more efficient, the both representations (12) and (13) converge to each other as k goes to infinity. We have computed the




Fig. 7. Comparison of SPHARM and the pullback representations for degree 5 to 25. Red colored numbers are the average Euclidean distance between two representations in mm.

squared Euclidean distance between two representations numerically (Figure 7). In average, the difference is 0.0569 mm for 20 degree representation negligible for 1mm resolution MR. Figure 7 also visually demonstrate that the pullback representation converges to the true manifold faster than SPHARM again showing the inefficiency of the SPHARM representation.

4 Conclusion

We have introduced a computationally efficient way to construct an approximate orthonormal basis on an arbitrary manifold by pulling back the spherical harmonics to the manifold and accounting for the metric distortion using the Jacobian determinant. The proposed technique is very general so that it can be applicable to other types of anatomical manifolds. The constructed basis on an amygdala is used to show the new pullback representation that reconstruct the manifold as linear combination of the basis functions.

References

1. S. Angenent, S. Hacker, A. Tannenbaum, and R. Kikinis. On the laplace-beltrami operator and brain surface flattening. *IEEE Transactions on Medical Imaging*, 18:700–711, 1999.
2. C. Brechbuhler, G. Gerig, and O. Kubler. Parametrization of closed surfaces for 3d shape description. *Computer Vision and Image Understanding*, 61:154–170, 1995.
3. M.K. Chung, L. Shen Dalton, K.M., A.C. Evans, and R.J. Davidson. Weighted fourier representation and its application to quantifying the amount of gray matter. *IEEE Transactions on Medical Imaging*, 26:566–581, 2007.

4. M.K. Chung, S. Robbins, Davidson R.J. Alexander A.L. Dalton, K.M., and A.C. Evans. Cortical thickness analysis in autism with heat kernel smoothing. *NeuroImage*, 25:1256–1265, 2005.
5. R. Courant and D. Hilbert. *Methods of Mathematical Physics: Volume II*. Interscience, New York, english edition, 1953.
6. B. Fischl and A.M. Dale. Measuring the thickness of the human cerebral cortex from magnetic resonance images. *PNAS*, 97:11050–11055, 2000.
7. K.M. Gorski. On determining the spectrum of primordial inhomogeneity from the coBE dmr sky maps: I. method. *Astrophysical Journal*, 430, 430:L85, 1994.
8. X. Gu, Y.L. Wang, T.F. Chan, T.M. Thompson, and S.T. Yau. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transactions on Medical Imaging*, 23:1–10, 2004.
9. M. K. Hurdal and K. Stephenson. Cortical cartography using the discrete conformal approach of circle packings. *NeuroImage*, 23:S119S128, 2004.
10. S.E. Jones, B.R. Buchbinder, and I. Aharon. Three-dimensional mapping of cortical thickness using laplace’s equation. *Human Brain Mapping*, 11:12–32, 2000.
11. A.N. Kolmogorov and S.V. Fomin. Introductory real analysis. Dover Publications, Inc, New York, 1970.
12. J. P. Lerch and A.C. Evans. Cortical thickness analysis examined through power analysis and a population simulation. *NeuroImage*, 24:163–173, 2005.
13. J.D. MacDonald, N. Kabani, D. Avis, and A.C. Evans. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from mri. *NeuroImage*, 12:340–356, 2000.
14. B.M. Nacewicz, K.M. Dalton, T. Johnstone, M.T. Long, E.M. McAuliff, T.R. Oakes, A.L. Alexander, and R.J. Davidson. Amygdala volume and nonverbal social impairment in adolescent and adult males with autism. *Arch. Gen. Psychiatry*, 63:1417–1428, 2006.
15. A. Qiu, D. Bitouk, and M.I. Miller. Smooth functional and structural maps on the neocortex via orthonormal bases of the laplace-beltrami operator. *IEEE Transactions on Medical Imaging*, 25:1296–1396, 2006.
16. L. Shen, J. Ford, F. Makedon, and A. Saykin. surface-based approach for classification of 3d neuroanatomical structures. *Intelligent Data Analysis*, 8:519–542, 2004.
17. M. Styner, I. Oguz, S. Xu, C. Brechbuhler, D. Pantazis, J. Levitt, M. Shenton, and G. Gerig. Framework for the statistical shape analysis of brain structures using spharm-pdm. In *Insight Journal, Special Edition on the Open Science Workshop at MICCAI*, 2006.
18. B. Timsari and R. Leahy. An optimization method for creating semi-isometric flat maps of the cerebral cortex. In *The Proceedings of SPIE, Medical Imaging*, 2000.
19. A. Yezzi and J.L. Prince. An eulerian pde approach for computing tissue thickness. *IEEE Transactions on Medical Imaging*, 22:1332–1339, 2003.

Diffusion Tensor Imaging and Deconvolution on Spaces of Positive Definite Symmetric Matrices

Peter T. Kim ^{*} and Donald St. P. Richards ^{**}

University of Guelph and Penn State University

Abstract. Diffusion tensor imaging can be studied as a deconvolution density estimation problem on the space of positive definite symmetric matrices. We develop a nonparametric estimator for the common density function of a random sample of positive definite matrices. Our estimator is based on the Helgason-Fourier transform and its inversion, the natural tools for analysis of compositions of random positive definite matrices. Under smoothness conditions on the density of the intrinsic error in the random sample, we derive bounds on the rates of convergence of our nonparametric estimator to the true density.

Key words: Deconvolution; diffusion tensor imaging; Harish-Chandra c -function; Helgason-Fourier transform; inverse problems; Laplace-Beltrami operator; magnetic resonance imaging; Sobolev spaces; Wishart distribution.

1 Introduction

The appearance in medical imaging of sequences of random positive definite matrices has become commonplace due to developments in *diffusion tensor imaging* (DTI), a method of imaging based upon the observation that water molecules in biological tissue are always in motion. For the purposes of mathematical modeling, it is generally assumed that the diffusion of water molecules at any given location in biological tissue follows a Brownian motion. A diffusion tensor image then is represented by the 3×3 positive-definite covariance matrix of the local diffusion process at the given location (Fletcher and Joshi [3, 4]). DTI seeks to detect the diffusion of water protons between and within distinct tissue cells, and to derive estimates of the dominant orientation and direction of the Brownian motion (Le Bihan [9], Hasan, *et al.* [5]).

In DTI brain imaging, the diffusion of water molecules within and between *voxels*, the three-dimensional volume elements that constitute an image, reveal

^{*} Department of Mathematics and Statistics, University of Guelph, Guelph, Ontario N1G 2W1, Canada. Research supported in part by the Natural Sciences and Engineering Research Council of Canada, grant DG-46204. Email: pkim@uoguelph.ca.

^{**} Department of Statistics, The Pennsylvania State University, University Park, PA 16802. Research supported in part by the U.S. National Science Foundation, grants DMS-0112069 and DMS-0705210. Email: richards@stat.psu.edu.

both the orientation of fibers comprising white-matter tracts in the brain and the coherence of fibers, the extent to which fibers are aligned together. DTI may be the only non-invasive, *in vivo* procedure which enables the study of deep brain white-matter fibers. Consequently, DTI has been found to be highly promising for comparing the human brain in normal states with abnormal states caused by strokes, epileptic seizures, tumors, white-matter abnormalities, multiple sclerosis lesions, HIV-infection, traumatic brain injuries, aging, Alzheimer's disease, alcoholism, and developmental disorders; and there are potential applications to psychiatric conditions including schizophrenia, autism, cognitive and learning disabilities (Neumann-Haefelin, *et al.* [12], Rosenbloom, *et al.* [14], Pomara, *et al.* [13], Matthews and Arnold [10]). In addition, DTI has been applied in research on the pathology of organ and tissue types such as the human breast, kidney, lingual, cardiac, skeletal muscles, and spinal cord (Damon, *et al.* [2]).

It is well-known that magnetic resonance imaging, from which diffusion tensors are derived, is endowed inherently with random noise. Hence, DTI data also contain noise (Basu, *et al.* [1]), and it is natural that statistical inferential issues arise in the analysis of DTI data (Koltchinskii, *et al.* [8]; Schwartzman, *et al.* [15, 16]; Zhu, *et al.* [18]).

In this paper, we study the problem of estimating the probability density function of a population of positive definite matrices based on a random sample from that population. An instance in which this problem arises may be obtained from Schwartzman [15] who studied the two-sample comparison of twelve children divided into two groups according to reading ability, where the issue is to compare physical characteristics of brain tissue of the two groups on the basis of DTI images. In addition to comparing the population parameters, it is natural to seek estimators of the underlying density functions, and then it will be important to estimate the rates of convergence of the density estimators.

The deconvolution density estimation problem has been widely studied on Euclidean spaces. This classical setting, the commutative nature of the underlying mathematical operations renders the problem amenable to classical mathematical methods. The deconvolution problem has also been studied on certain compact manifolds. In that setting, the problem is solvable using well-known generalizations of classical Fourier analysis. By contrast, \mathcal{P}_m , the space of $m \times m$ positive definite matrices, is a noncompact Riemannian manifold and has an intrinsic non-commutative nature, and it is natural to expect that the deconvolution problem will be more difficult in that setting. To the best of our knowledge, no results yet are available for the general deconvolution problem on \mathcal{P}_m .

To date, the primary statistical emphases regarding mathematical methods on \mathcal{P}_m are motivated by properties of the Wishart distribution, see for example Muirhead [11]. To solve the general deconvolution problem on \mathcal{P}_m , more advanced mathematical methods are required. While much is known in the mathematical literature about the necessary methods (Helgason [6], Terras [17]), virtually nothing about those methods has appeared in the statistical literature. Using these new methods, we develop nonparametric methods for solving the deconvolution problem on \mathcal{P}_m .

In summary, section 2 provides notation, and introduces the *Helgason-Fourier transform* and its inversion formula, involving the *Harish-Chandra c -function*; we have provided the necessary details of these concepts so as to make the paper fully accessible to readers who are new to this area. We formalize in section 3 the statistical procedure in terms of measurement errors on \mathcal{P}_m and present the main results on general deconvolution density estimation on \mathcal{P}_m and we provide the explicit details in the case of the Wishart distribution. All proofs can be found in Kim and Richards [7].

2 Preliminaries

Throughout the paper, we denote by G the general linear group $\text{GL}(m, \mathbb{R})$, of $m \times m$ nonsingular real matrices and by K the group, $\text{O}(m)$, of $m \times m$ orthogonal matrices. The group G acts transitively on \mathcal{P}_m , the space of $m \times m$ positive definite matrices, by the action

$$G \times \mathcal{P}_m \rightarrow \mathcal{P}_m, \quad (g, x) \mapsto g'xg,$$

$g \in G$, $x \in \mathcal{P}_m$, where g' denotes the transpose of g . Under this action, the isotropy group of the identity in G is K , hence the homogeneous space $K \backslash G$ can be identified with \mathcal{P}_m by the “natural” mapping

$$K \backslash G \rightarrow \mathcal{P}_m, \quad Kg \mapsto g'g. \quad (1)$$

In distinguishing between left and right cosets, we place the quotient operation on the left and right of the group, respectively.

A random matrix $X \in \mathcal{P}_m$ is said to be K -invariant if $X \stackrel{\mathcal{L}}{=} k'Xk$ for all $k \in K$, where “ $\stackrel{\mathcal{L}}{=}$ ” denotes equality in distribution. A function f on \mathcal{P}_m is called K -invariant if $f(k'xk) = f(x)$ for all $k \in K$, $x \in \mathcal{P}_m$; we will indicate that f is K -invariant by writing its domain as \mathcal{P}_m/K , with a similar notation for K -invariant positive definite random matrices.

By means of the relationship (1) between $K \backslash G$ and \mathcal{P}_m , we identify K -invariant functions on \mathcal{P}_m with K -biinvariant functions on G , i.e., functions $\tilde{f} : G \rightarrow \mathbb{C}$ which satisfy $\tilde{f}(g) = \tilde{f}(k_1 g k_2)$ for all $k_1, k_2 \in K$ and $g \in G$. In particular, $\mathcal{P}_m/K \simeq K \backslash G / K$ where “ \simeq ” denotes diffeomorphic equivalence.

Consider random matrices $X, \varepsilon \in \mathcal{P}_m$ with corresponding group elements $\tilde{X}, \tilde{\varepsilon} \in G$, respectively. By the natural map (1), $K\tilde{X} \mapsto X$, equivalently $\tilde{X}'\tilde{X} = X$ and, similarly, $K\tilde{\varepsilon} \mapsto \varepsilon$. Then $\tilde{X}\tilde{\varepsilon} \in G$ is mapped *via* (1) to

$$\tilde{X}\tilde{\varepsilon} \mapsto (\tilde{X}\tilde{\varepsilon})'(\tilde{X}\tilde{\varepsilon}) \equiv \tilde{\varepsilon}'\tilde{X}'\tilde{X}\tilde{\varepsilon} = \tilde{\varepsilon}'X\tilde{\varepsilon} \in \mathcal{P}_m. \quad (2)$$

If $\tilde{\varepsilon}$ is K -biinvariant, i.e., $\tilde{\varepsilon} \stackrel{\mathcal{L}}{=} k_1 \tilde{\varepsilon} k_2$ for all $k_1, k_2 \in K$ then $\tilde{\varepsilon}'\tilde{\varepsilon} = \varepsilon = k' \Lambda k$ where $k \in K$ and $\Lambda \in \mathcal{P}_m$ is the diagonal matrix of eigenvalues of ε . Consequently, we define $\varepsilon^{1/2} = k' \Lambda^{1/2} k$ and note that, by K -biinvariance of $\tilde{\varepsilon}$, the relationship (2) on the group G corresponds in distribution to $\varepsilon^{1/2} X \varepsilon^{1/2} \in \mathcal{P}_m$, where $X \in \mathcal{P}_m$ and $\varepsilon \in \mathcal{P}_m/K$. Bearing this in mind, we formally make the following

definition. Suppose that $X \in \mathcal{P}_m$ and $\varepsilon \in \mathcal{P}_m/K$ are random matrices. Then the composition of X and ε is

$$X \circ \varepsilon = \varepsilon^{1/2} X \varepsilon^{1/2} \quad (3)$$

where $\varepsilon^{1/2}$ is the positive definite square root of ε .

There is an alternative approach leading to compositions (3). We begin by noting that \mathcal{P}_m is a Riemannian manifold, hence each pair $X, Y \in \mathcal{P}_m$ defines a unique geodesic path. Since \mathcal{P}_m is also a homogeneous space, *viz.*, $\mathcal{P}_m = K \backslash G$, then there exists a unique $V \in G$ such that

$$Y = V X V'. \quad (4)$$

This model can be viewed as a multiplicative analog of classical regression analysis, with Y serving as the dependent variable, X as the independent variable, and V as the error variable. We can also study this relationship between X and Y using the logarithm map on \mathcal{P}_m . To that end, each observation on Y may be transformed into an observation on y where $\exp(y) = Y$ with $y \in T(\mathcal{P}_m)$, the tangent space of \mathcal{P}_m . Similarly, each measurement on X may be transformed into a measurement on $x \in T(\mathcal{P}_m)$ where $\exp(x) = X$. By postulating the existence of a “small” error $v \in T(\mathcal{P}_m)$ such that $y = x + v$ and, by exponentiating this linear “regression” relationship between y and x , we are led naturally to (4); see Terras [17], section 4.1-4.2 for details on the exponentiation of such linear relationships on $T(\mathcal{P}_m)$.

If we assume that the measurement error V is *isotropic* or has no preferred orientation, i.e., V has a K -invariant distribution, then the conclusion is the model (4) in which $V \in G/K$. As noted in [14], in the context of DTI, the assumption that water molecules diffuse isotropically is appropriate for regions such as the ventricles, which are large fluid-filled spaces deep in the brain. On the other hand, water molecules located in white-matter fiber are constrained by the axon sheath; this forces greater movement along the longitudinal axes of fibers than across the axes, and then diffusion may be isotropic only at sufficiently small scales.

Returning to (4), we apply polar coordinates on G , *viz.*, $V = k\varepsilon^{1/2}$ where $k \in K$ and $\varepsilon \in \mathcal{P}_m$. Since V is K -invariant then $V \stackrel{\mathcal{L}}{=} k'V = k'k\varepsilon^{1/2} = \varepsilon^{1/2}$; hence (4) reduces to $Y \stackrel{\mathcal{L}}{=} \varepsilon^{1/2} X \varepsilon^{1/2} \equiv X \circ \varepsilon$, which agrees with (3) since $\varepsilon = V'V$ is K -invariant. Thus, the problem is to estimate nonparametrically the density of X based on a random sample from $Y = X \circ \varepsilon$ where the error matrix ε is isotropic.

Let $C_c^\infty(\mathcal{P}_m)$ denote the collection of complex-valued, infinitely differentiable, compactly supported functions f on \mathcal{P}_m . For $w \in \mathcal{P}_m$ and $j = 1, \dots, m$, denote by $|w_j|$ the principal minor of order j of w . For $s = (s_1, \dots, s_m) \in \mathbb{C}^m$, the *power function* $p_s : \mathcal{P}_m \rightarrow \mathbb{C}$ is defined by

$$p_s(w) = \prod_{j=1}^m |w_j|^{s_j},$$

$w \in \mathcal{P}_m$.

Let dk denote the Haar measure on K , normalized to have total volume equal to one. Then

$$h_s(w) = \int_K p_s(k'wk) dk \quad (5)$$

$w \in \mathcal{P}_m$, $s \in \mathbb{C}^m$, is a *zonal spherical function* on \mathcal{P}_m . It is well-known that the spherical functions play a fundamental role in harmonic analysis on symmetric spaces; see Helgason [6]. In particular, (5) is a special case of Harish-Chandra's formula for the general spherical functions. If s_1, \dots, s_m are nonnegative integers then, except for a constant factor, (8) is an integral formula for the zonal polynomials which arise often in aspects of multivariate distribution theory; see Muirhead [11], pp. 231–232.

Let $w = (w_{ij}) \in \mathcal{P}_m$. Then, up to a constant factor, the unique G -invariant measure on \mathcal{P}_m is

$$d_*w = |w|^{-(m+1)/2} \prod_{1 \leq i \leq j \leq m} dw_{ij}, \quad (6)$$

where $|w|$ is the determinant of w .

For $s = (s_1, \dots, s_m) \in \mathbb{C}^m$ and $k \in K$, the *Helgason-Fourier transform* ([17], p. 87) of a function $f \in C_c^\infty(\mathcal{P}_m)$ is

$$\mathcal{H}f(s, k) = \int_{\mathcal{P}_m} f(w) \overline{p_s(k'wk)} d_*w, \quad (7)$$

where $\overline{p_s(k'wk)}$ denotes complex conjugation and d_*w is the G -invariant measure (6).

For the case in which $f \in C_c^\infty(\mathcal{P}_m/K)$, we make the change of variables $w \rightarrow k'_1wk_1$ in (7), $k_1 \in K$, and integrate with respect to the Haar measure dk_1 . Applying the invariance of f and the formula (5) for the zonal spherical function, we deduce that $\mathcal{H}f(s, k)$ does not depend on k ; specifically, $\mathcal{H}f(s, k) = \widehat{f}(s)$ where

$$\widehat{f}(s) = \int_{\mathcal{P}_m} f(w) \overline{h_s(w)} d_*w, \quad (8)$$

is the *zonal spherical transform* of f .

Let $A = \{\text{diag}(a_1, \dots, a_m) : a_j > 0, j = 1, \dots, m\}$, denote the group of diagonal positive definite matrices in G , and $N = \{n = (n_{ij}) \in G : n_{ij} = 0, 1 \leq j < i \leq m; n_{jj} = 1, j = 1, \dots, m\}$ be the subgroup of G consisting of upper-triangular matrices with all diagonal entries equal to 1. It is well-known (see Terras [17], p. 20) that each $g \in G$ can be decomposed uniquely as $g = kan$ where $k \in K$, $a \in A$, and $n \in N$; this result is the *Iwasawa decomposition*, and (k, a, n) are called the *Iwasawa coordinates* of g .

For $a, b \in \mathbb{C}$ with $\text{Re}(a), \text{Re}(b) > 0$, let $B(a, b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ denote the well-known beta function, where $\Gamma(\cdot)$ is the classical gamma function. For $s = (s_1, \dots, s_m) \in \mathbb{C}^m$, the *Harish-Chandra c -function* is

$$c_m(s) = \prod_{1 \leq i < j \leq m-1} \frac{B(\frac{1}{2}, s_i + \dots + s_j + \frac{1}{2}(j-i+1))}{B(\frac{1}{2}, \frac{1}{2}(j-i+1))}.$$

Let $\rho \equiv (\frac{1}{2}, \dots, \frac{1}{2}, \frac{1}{4}(1-m))$ and set

$$\omega_m = \frac{\prod_{j=1}^m \Gamma(j/2)}{(2\pi i)^m \pi^{m(m+1)/4} m!} . \quad (9)$$

We shall use the notation

$$\mathbb{C}^m(\rho) = \{s \in \mathbb{C}^m : \operatorname{Re}(s) = -\rho\} \quad (10)$$

because this subset of \mathbb{C}^m arises frequently in the sequel, and we also define

$$d_* s = \omega_m |c_m(s)|^{-2} ds_1 \cdots ds_m \quad (11)$$

since this measure is ubiquitous in our development.

Let $M = \{\operatorname{diag}(\pm 1, \dots, \pm 1)\}$ be the collection of $m \times m$ diagonal matrices with entries ± 1 on the diagonal; then M is a subgroup of K and is of order 2^m . By factorizing the Haar measure dk on K , it may be shown ([17], p. 88) that there exists an invariant measure $d\bar{k}$ on the coset space K/M such that

$$\int_{\bar{k} \in K/M} d\bar{k} = 1.$$

In stating the inversion formula for the Helgason-Fourier transform, we make particular use of the notation (9)-(11). The inversion formula then is the following result.

For $f \in C_c^\infty(\mathcal{P}_m)$ and $w \in \mathcal{P}_m$,

$$f(w) = \int_{\mathbb{C}^m(\rho)} \int_{\bar{k} \in K/M} \mathcal{H}f(s, k) p_s(k'wk) d\bar{k} d_* s. \quad (12)$$

In particular, if $f \in C_c^\infty(\mathcal{P}_m/K)$ then

$$f(w) = \int_{\mathbb{C}^m(\rho)} \widehat{f}(s) h_s(w) d_* s. \quad (13)$$

We refer to Terras [17], p. 87 ff. for a detailed treatment of this inversion formula and many references to the literature.

For $f \in L^1(\mathcal{P}_m)$ and $h \in L^1(\mathcal{P}_m/K)$ we define $f * h$, the convolution of f and h , by

$$(f * h)(w) = \int_{\mathcal{P}_m} f(z) h(z^{-1/2} w z^{-1/2}) d_* z,$$

$w \in \mathcal{P}_m$. Thus, if f and h are the density functions of independent random matrices $X \in \mathcal{P}_m$ and $\varepsilon \in \mathcal{P}_m/K$, respectively, then $f * h$ is the density function of the composition $X \circ \varepsilon$.

For $f \in C_c^\infty(\mathcal{P}_m)$ and $h \in C_c^\infty(\mathcal{P}_m/K)$, the *convolution property of the Helgason-Fourier transform* is that

$$\mathcal{H}(f * h)(s, k) = \mathcal{H}f(s, k) \widehat{h}(s), \quad (14)$$

$s \in \mathbb{C}^m$, $k \in K$; see Terras [17], Theorem 1, p. 88.

3 Deconvolution density estimation on \mathcal{P}_m

On the group G , the deconvolution problem arises from the statistical model

$$\tilde{Y} \stackrel{\mathcal{L}}{=} \tilde{X}\tilde{\varepsilon} \quad (15)$$

where \tilde{X} is a random unobservable, $\tilde{\varepsilon}$ is an independent random error, and \tilde{Y} is the observed random measurement. We assume that $f_{\tilde{\varepsilon}}$, the density of $\tilde{\varepsilon}$, is known and K -biinvariant and that the unknown densities $f_{\tilde{X}}$ and $f_{\tilde{Y}}$ of \tilde{X} and \tilde{Y} , respectively, are K -invariant. Under the equivalence $K \backslash G \simeq \mathcal{P}_m$, we have $\tilde{\varepsilon} \mapsto \varepsilon$, $\tilde{X} \mapsto X$, and $\tilde{Y} \mapsto Y$ together with the identification $f_{\varepsilon}(\tilde{g}) = f_{\varepsilon}(\tilde{g}'\tilde{g})$, $f_X(\tilde{g}) = f_X(\tilde{g}'\tilde{g})$ and $f_Y(\tilde{g}) = f_Y(\tilde{g}'\tilde{g})$, $g \in G$. Since \tilde{X} and $\tilde{\varepsilon}$ are independent then X and ε also are independent, and (15) implies that

$$Y \stackrel{\mathcal{L}}{=} \varepsilon^{1/2} X \varepsilon^{1/2}; \quad (16)$$

hence $f_Y = f_X * f_{\varepsilon}$.

Applying to (16) the convolution property (14) of the Helgason-Fourier transform, we obtain

$$\mathcal{H}f_Y(s, k) = \mathcal{H}f_X(s, k) \hat{f}_{\varepsilon}(s), \quad (17)$$

$s \in \mathbb{C}^m$, $k \in K$. Given a random sample Y_1, \dots, Y_n from Y , we estimate the density function f_X as follows. We form $\mathcal{H}_n f_Y$, the *empirical Helgason-Fourier transform*,

$$\mathcal{H}_n f_Y(s, k) = \frac{1}{n} \sum_{\ell=1}^n \overline{p_s(k' Y_{\ell} k)}. \quad (18)$$

Substituting (18) in (17), together with the assumption that $\hat{f}_{\varepsilon}(s) \neq 0$, $s \in \mathbb{C}^m$, we obtain

$$\mathcal{H}_n f_X(s, k) = \frac{\mathcal{H}_n f_Y(s, k)}{\hat{f}_{\varepsilon}(s)},$$

$s \in \mathbb{C}^m$, $k \in K$.

In analogy with classical Euclidean deconvolution, we introduce a smoothing parameter $T = T(n)$ where $T(n) \rightarrow \infty$ as $n \rightarrow \infty$, and then we apply the Helgason-Fourier inversion formula (12) using a spectral cut-off.

We introduce the notation

$$\mathbb{C}^m(\rho, T) = \{s \in \mathbb{C}^m(\rho) : \lambda_s < T\}$$

where $\mathbb{C}^m(\rho)$ is defined in (10). As an estimator of the population density f_X , we define the density estimator f_X^n given by

$$f_X^n(w) = \int_{\mathbb{C}^m(\rho, T)} \int_{\bar{k} \in K/M} \frac{\mathcal{H}_n f_Y(s, \bar{k})}{\hat{f}_{\varepsilon}(s)} p_s(\bar{k}' w \bar{k}) d\bar{k} d_* s,$$

$w \in \mathcal{P}_m$. The estimator f_X^n will serve as our nonparametric deconvolution estimator of the density f_X .

Let $\{a_n\}$ and $\{b_n\}$ be sequences of real numbers. We write $a_n \ll b_n$ to mean $a_n \leq Cb_n$ for some constant $C > 0$, as $n \rightarrow \infty$ (the Vinogradov notation). We use the notation $a_n = o(b_n)$ to mean $a_n/b_n \rightarrow 0$, as $n \rightarrow \infty$. We also write $a_n \asymp b_n$ if both $a_n \ll b_n$ and $b_n \ll a_n$; and we write $a_n \sim b_n$ if $a_n/b_n \rightarrow 1$ as $n \rightarrow \infty$.

For technical reasons, we will also assume the moment condition

$$\int_{\mathcal{P}_m} |w_1|^{-1} \cdots |w_{m-1}|^{-1} |w|^{(m-1)/2} f_Y(w) d_* w < \infty, \quad (19)$$

on the principal minors $|y_1|, \dots, |y_m|$ of $y \in \mathcal{P}_m$. This assumption will be maintained throughout the rest of the paper.

Estimation will proceed in the *Sobolev class* of functions,

$$H_\sigma(\mathcal{P}_m) = \{f \in C^\infty(\mathcal{P}_m) : \|\Delta^{\sigma/2} f\|^2 < \infty\},$$

where $2\sigma > \dim \mathcal{P}_m = m(m+1)/2$, Δ is the Laplacian on \mathcal{P}_m and, for $f \in C^\infty(\mathcal{P}_m)$,

$$\|f\| = \left(\int_{\mathcal{P}_m} |f(w)|^2 d_* w \right)^{1/2}$$

denotes the $L^2(\mathcal{P}_m)$ -norm with respect to the invariant measure $d_* w$. For $Q > 0$, we also define the bounded Sobolev class,

$$H_\sigma(\mathcal{P}_m, Q) = \{f \in C^\infty(\mathcal{P}_m) : \|\Delta^{\sigma/2} f\|^2 < Q\},$$

where $2\sigma > \dim \mathcal{P}_m = m(m+1)/2$.

Theorem 1. *Suppose there exists $\beta \geq 0$ such that*

$$|\widehat{f}_\varepsilon(s)|^{-2} \ll T^\beta$$

as $T \rightarrow \infty$, for all $s \in \mathbb{C}^m(\rho, T)$. If $f_X \in H_\sigma(\mathcal{P}_m, Q)$ and $\sigma > \frac{1}{2} \dim \mathcal{P}_m \equiv m(m+1)/4$ then, as $n \rightarrow \infty$,

$$\mathbb{E} \|f_X^n - f_X\|^2 \ll n^{-2\sigma/(2\sigma+2\beta+\dim \mathcal{P}_m)}.$$

Corollary 1. *Suppose the distribution of ε is concentrated at \mathbf{I}_m . If $f_X \in H_\sigma(\mathcal{P}_m, Q)$ where $\sigma > \frac{1}{2} \dim \mathcal{P}_m$ then as $n \rightarrow \infty$,*

$$\mathbb{E} \|f_X^n - f_X\|^2 \ll n^{-2\sigma/(2\sigma+\dim \mathcal{P}_m)}.$$

We shall also obtain a result for the situation in which the hypothesis in Theorem 1 is replaced by an exponential bound. In such a situation, we have the following result.

Theorem 2. *Suppose there exists $\beta, \gamma > 0$ such that*

$$|\widehat{f}_\varepsilon(s)|^{-2} \ll \exp(T^\beta/\gamma),$$

as $T \rightarrow \infty$, for all $s \in \mathbb{C}^m(\rho, T)$. If $f_X \in H_\sigma(\mathcal{P}_m, Q)$ with $\sigma > \frac{1}{2} \dim \mathcal{P}_m$ then, as $n \rightarrow \infty$,

$$\mathbb{E} \|f_X^n - f_X\|^2 \ll (\log n)^{-\sigma/\beta}.$$

In this situation, we consider the special case in which

$$\widehat{f}_\varepsilon(s) = \exp(-\gamma^{-1} \lambda_s^\beta),$$

$s \in \mathbb{C}^m$, where $\gamma > 0$ is a scale parameter. Again by the inversion formula (13), the underlying density function is

$$f_\varepsilon(w) = \int_{\mathbb{C}^m(\rho)} \exp(-\gamma^{-1} \lambda_s^\beta) h_s(w) d_* s,$$

$w \in \mathcal{P}_m$. The case in which $\beta = 1$ is particularly important and is called the *heat* or *Gaussian kernel*, since the latter is the fundamental solution to the heat equation on \mathcal{P}_m , see Terras [17], pp. 106-107. As a consequence, we obtain the following result.

Corollary 2. *Suppose that f_ε is Gaussian. If $f_X \in H_\sigma(\mathcal{P}_m, Q)$ where $\sigma > \frac{1}{2} \dim \mathcal{P}_m$ then, as $n \rightarrow \infty$,*

$$\mathbb{E} \|f_X^n - f_X\|^2 \ll (\log n)^{-\sigma}.$$

A case which is familiar in multivariate statistics is that of the Wishart distribution, $W_m(N, \Sigma)$, where $\Sigma \in \mathcal{P}_m$ and $N > m - 1$. For $s = (s_1, \dots, s_m) \in \mathbb{C}^m$ define the multivariate gamma function,

$$\Gamma_m(s_1, \dots, s_m) = \pi^{m(m-1)/4} \prod_{j=1}^m \Gamma(s_j + \dots + s_m - \frac{1}{2}(j-1)),$$

where $\operatorname{Re}(s_j + \dots + s_m) > (j-1)/2$, $j = 1, \dots, m$. Relative to the invariant measure $d_* w$ in (6), the probability density function of the standard Wishart distribution $W_m(N, \mathbf{I}_m)$ is

$$f_\varepsilon(w) = \frac{1}{\Gamma_m(0, \dots, 0, N/2)} \left| \frac{1}{2} w \right|^{N/2} \exp\left(-\frac{1}{2} \operatorname{tr} w\right), \quad (20)$$

$w \in \mathcal{P}_m$. We note that (20) is K -invariant and its Helgason-Fourier transform is well-known (Muirhead [11], p. 248; Terras [17], pp. 85-86),

$$\widehat{f}_\varepsilon(s) = \frac{\Gamma_m((0, \dots, 0, N/2) + s^*)}{\Gamma_m(0, \dots, 0, N/2)} h_s\left(\frac{1}{2} \mathbf{I}_m\right)$$

where $s^* = (s_{m-1}, s_{m-2}, \dots, s_2, s_1, -(s_1 + \dots + s_m))$.

For $N > m - 1$, the Wishart distribution $W_m(N, \mathbf{I}_m)$ satisfies

$$|\widehat{f}_\varepsilon(s)|^{-2} \ll \exp(\pi T^{1/2}),$$

as $T \rightarrow \infty$, where $s \in \mathbb{C}^m(\rho, T)$.

Consequently we deduce the following result.

Theorem 3. *Suppose that ε follows the Wishart distribution (20) with $N > m - 1$. If $f_X \in H_\sigma(\mathcal{P}_m, Q)$ with $\sigma > \dim \mathcal{P}_m/2$ then, as $n \rightarrow \infty$,*

$$\mathbb{E} \|f_X^n - f_X\|^2 \ll (\log n)^{-2\sigma}.$$

References

1. Basu, S., Fletcher, T., Whitaker, S.: Rician noise removal in diffusion tensor MRI. In: Larsen, R., Nielsen, M., Sporring, J. (eds.) MICCAI 2006. LNCS, vol. 4190, pp. 117–125. Springer, Heidelberg, (2006)
2. Damon, B. M., Ding, Z., Anderson, A. W., Freyer, A. S., Gore, J. C.: Validation of diffusion tensor MRI-based muscle fiber tracking. *Magn. Reson. Med.* 48(1), 97–104 (2002)
3. Fletcher, P. T., Joshi, S.: Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In: Šonka, M., Kakadiaris, I. A., Kybic, J. (eds.) ECCV 2004 Workshop CVAMIA. LNCS, vol. 3117, pp. 87–98. Springer, Heidelberg, (2004)
4. Fletcher, P. T. and Joshi, S.: Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing* 87(2) 250–262 (2007)
5. Hasan, K. M., Basser, P. J., Parker, D. L., Alexander, A. L.: Analytical computation of the eigenvalues and eigenvectors in DT-MRI. *J. Magn. Res.* 152(1), 41–47 (2001)
6. Helgason, S.: *Groups and Geometric Analysis*. Academic Press, Orlando, (1984)
7. Kim, P. T. and Richards, D. St. P.: Deconvolution density estimation on spaces of positive definite symmetric matrices. Preprint (2007)
8. Koltchinskii, V., Sakhanenko, L., Cao, S.: Integral curves of noisy vector fields and statistical problems in diffusion tensor imaging: Nonparametric kernel estimation and hypotheses testing. *Ann. Statist.* 35(4), 1576–1607 (2007)
9. Le Bihan, D.: *Diffusion and Perfusion Magnetic Resonance Imaging: Applications to Functional MRI*. Raven Press, New York, (1995)
10. Matthews, P. M., Arnold, D. L.: Magnetic resonance imaging of multiple sclerosis: new insights linking pathology to clinical evolution. *Curr. Opin. Neurol.* 14(3), 279–287 (2001)
11. Muirhead, R. J.: *Aspects of Multivariate Statistics*. Wiley, New York, (1978)
12. Neumann-Haefelin T., Moseley M. E., Albers, G. W.: New magnetic resonance imaging methods for cerebrovascular disease. *Ann. Neurol.* 47(5), 559–570 (2000)
13. Pomara, N., Crandalla, D. T., Choia, S. J., Johnson, G. Lim, K. O.: White-matter abnormalities in HIV-1 infection: A diffusion tensor imaging study. *Psychiatry Res.: Neuroimaging* 106(1), 15–24 (2001)
14. Rosenbloom, M., Sullivan, J. V., Pfefferbaum, A.: Using magnetic resonance imaging and diffusion tensor imaging to assess brain damage in alcoholics. *Alcohol Research & Health* 27(2), 146–152 (2003)
15. Schwartzman, A.: *Random Ellipsoids and False Discovery Rates: Statistics for Diffusion Tensor Imaging Data*. Ph.D. dissertation, Stanford University (2006)
16. Schwartzman, A., Dougherty, R. F., Taylor, J. E.: Cross-subject comparison of principal diffusion direction maps. *Magn. Reson. Med.* 53(6), 1423–1431 (2005)
17. Terras, A.: *Harmonic Analysis on Symmetric Spaces and Applications, II*. Springer, New York, (1988).
18. Zhu, H., Zhang, H., Ibrahim, J. G., Peterson, B. G.: Statistical analysis of diffusion tensors in diffusion-weighted magnetic image resonance data (with discussion). *J. Amer. Statist. Assoc.* 102, 1085–1102 (2007)

Tubular Surface Evolution for Segmentation of the Cingulum Bundle From DW-MRI

Vandana Mohan¹, Ganesh Sundaramoorthi², John Melonakos¹, Marc Niethammer³,
Marek Kubicki⁴, and Allen Tannenbaum¹

¹ Georgia Institute of Technology, Atlanta, GA, USA

² University of California, Los Angeles, CA, USA

³ University of North Carolina, Chapel Hill, NC, USA

⁴ Brigham Women's Hospital, MA, USA

gth115a@mail.gatech.edu

Abstract. This work provides a framework for modeling and extracting the Cingulum Bundle (CB) from Diffusion-Weighted Imagery (DW-MRI) of the brain. The CB is a tube-like structure in the brain that is of potentially of tremendous importance to clinicians since it may be helpful in diagnosing Schizophrenia. This structure consists of a collection of fibers in the brain that have locally similar diffusion patterns, but vary globally. Standard region-based segmentation techniques adapted to DW-MRI are not suitable here because the diffusion pattern of the CB cannot be described by a *global* set of simple statistics. Active surface models extended to DW-MRI are not suitable since they allow for arbitrary deformations that give rise to unlikely *shapes*, which do not respect the tubular geometry of the CB. In this work, we explicitly model the CB as a tube-like surface and construct a general class of energies defined on tube-like surfaces. An example energy of our framework is optimized by a tube that encloses a region that has *locally* similar diffusion patterns, which differ from the diffusion patterns immediately outside. Modeling the CB as a tube-like surface is a *natural shape prior*. Since a tube is characterized by a center-line and a radius function, the method is reduced to a 4D (center-line plus radius) curve evolution that is computationally much less costly than an arbitrary surface evolution. The method also provides the center-line of CB, which is potentially of clinical significance.

1 Introduction

In this work, we are interested in extracting a structure in the brain called the *cingulum bundle* (CB) from *diffusion-weighted magnetic resonance imagery* (DW-MRI) of the brain. DW-MRI is imagery that at each voxel indicates the diffusion of water molecules at each particular sampling *direction* in 3D space. Adding an extra dimension, directionality, to the data is necessary to discriminate our structure of interest - the cingulum bundle. The CB has recently become the subject of interest as an anatomical structure which may display quantifiable differences between schizophrenic and normal control populations, and studying it may aid in the diagnosis of schizophrenia [1, 2].

The Cingulum Bundle is a *thin, highly curved* structure that consists of a collection of neural fibers, which are mostly disjoint possibly intersecting, roughly aligned

and centered around a fiber. The collection of fibers approximately form a tube-like structure. The diffusion pattern in the CB varies in orientation and anisotropy smoothly along the structure, and it has a distinct diffusion pattern from surrounding areas of the brain (see Fig. 2 for a sagittal slice of the CB). The CB possesses a challenge to segment because of inhomogeneity of its diffusion pattern *globally* and the noisy nature of DW-MRI makes it difficult to detect edges separating the CB from the rest of the image.

There has been much research in detecting and characterizing neural connections between brain structures in DW-MRI. Early methods for detecting fibers, i.e., *tractography*, are based on streamlines where the fiber path is constructed by following the direction of the principal eigenvector of diffusion tensors from an initial seed point, e.g., [3, 4]. These methods have been shown to perform poorly in noisy situations and they often terminate prematurely before the fiber ending. To alleviate these problems, there has been a number of works, e.g., [5–11], where an optimal path, in some sense, is constructed from seed region(s). The procedure is repeated to detect all fibers of a bundle. These methods, however, do not explicitly provide a segmentation of the entire fiber bundle.

We are interested in segmenting the entire fiber bundle as a surface. Standard active surface techniques, e.g. [12, 13], adapted to DW-MRI are typically difficult to segment the CB since the DW-MRI of the brain are extremely noisy and contain many local features that trap the active surface in unlikely configurations that are not representative of the CB. Indeed, the CB is difficult to segment without a *shape prior* favoring its thin tube-like geometry. Standard region-based techniques adapted to DW-MRI or DT-MRI, e.g. [14], are generally not applicable to the segmentation of the CB since the statistics of the DW-MRI inside the CB cannot be described by a few *global* parameters (e.g. mean). The Mumford-Shah energy extended to DT-MRI, [15], which assumes piecewise smooth image data inside the surface, is applicable to the CB, but the technique needs a shape prior for the CB and is computationally costly since smooth functions must be determined at each update of the evolving surface. In [16], the authors model the probability distribution of the CB and design an algorithm to classify voxels of the image, and the method could benefit greatly modeling the CB geometry. Noticing that standard region-based techniques are not applicable to the CB, an edge-based active surface method for segmenting the CB is considered by [17]. However, the method is sensitive to the noise in the DW-MRI, and the method does not incorporate the tube-like geometry into the segmentation. The work [18] designs an energy on volumetric regions that incorporates “local region-based” information and a prior favoring regions that are close to an initially detected center-line curve. However, the energy is highly dependent on the correct placement of the detected center-line, which is often not exactly in the center of the CB. Moreover, the method does not enforce the tube-like geometry of the CB.

In this work, we *explicitly model the CB as a tubular surface*, and *construct a general class of energies defined on these tubular surfaces*. This enforces a tubular geometry during the segmentation process. Since the tubular surfaces we consider are determined by a center-line in 3D space and a radius function defined at each point of the center-line (see Fig. 1), the problem is reduced to optimizing an energy defined on 4D curves. *This significantly reduces the computational cost of the optimization proce-*

ture when compared to an unconstrained surface optimization. Further, we show how to construct energies that are tailored to the varying nature of the diffusion pattern in the CB.

Our method is inspired by the work of [19] in which the authors model vessels as tubular regions formed by the union of spheres along a center-line. Energies are constructed on 4D curves that represent tubes, and these energies are globally minimized using the minimal path technique [20]. The energies we construct cannot be optimized using the minimal path technique since our energies are directionally dependent - they depend on the position of the 4D curve and its *tangent*. Moreover, for the energies we consider, we are not interested in a global minimum but rather certain local minima. As we shall see (in Section 4), the optimization of our energy of interest using gradient descent requires special consideration, and this interestingly ties to the metric structure on the space of 4D curves.

2 The Cingulum Bundle

The cingulum bundle is a 5-7 mm in diameter fiber bundle that interconnects all parts of the limbic system. It originates within the white matter of the temporal pole, and runs posterior and superior into the parietal lobe, then turns, forming a “ring-like belt” around the corpus callosum, into the frontal lobe, terminating anterior and inferior to the genu of the corpus callosum in the orbital-frontal cortex [21]. Moreover, the CB consists of long, association fibers that directly connect temporal and frontal lobes, as well as shorter fibers radiating into their own gyri. The CB also includes most afferent and efferent cortical connections of cingulate cortex, including those of prefrontal, parietal and temporal areas, and the thalamostriatae bundle. In addition, lesion studies document a variety of neurobehavioral deficits resulting from a lesion located in this area, including akinetic mutism, apathy, transient motor aphasia, emotional disturbances, attentional deficits, motor activation, and memory deficits. Because of its involvement in executive control and emotional processing, the CB has been investigated in several clinical populations, including depression and schizophrenia. Previous studies, using DTI, in schizophrenia, demonstrated decrease of FA in anterior part of the cingulum bundle [1, 2], at the same time pointing to the technical limitations restricting these investigations from following the entire fiber tract.

3 Proposed Framework

In this section, we model the cingulum bundle (CB) as a tubular surface in \mathbb{R}^3 . We show that the tubular surface is completely determined by its center-line and the radius function of the discs along the center-line, and therefore, the tubular surface in \mathbb{R}^3 can be effectively reduced to a curve in \mathbb{R}^4 . We formulate a general class of energies directly on curves living in \mathbb{R}^4 , and then observe that special consideration of the metric structure on curves is needed to optimize the energy.

3.1 Modeling the Cingulum Bundle (CB) as a Tubular Surface

We are interested in tubular surfaces since these surfaces naturally model the CB. These surfaces have the additional advantage that they may be represented as space curves thus significantly reducing the computational complexity of our algorithm.




Fig. 1. Illustration of Tubular Surface model

The tubular surfaces we consider are determined by a center line, which is an open curve lying in \mathbb{R}^3 , and a radius function defined at each point of the center-line. Given an open curve $c : [0, 1] \rightarrow \mathbb{R}^3$, the center line, and a function $r : [0, 1] \rightarrow \mathbb{R}^+$, the radius function, we can define the tubular surface, $S : \mathbb{S}^1 \times [0, 1] \rightarrow \mathbb{R}^3$ (\mathbb{S}^1 is $[0, 2\pi]$ with endpoints identified) as follows:

$$S(\theta, u) = c(u) + r(u)[n_1(u) \cos \theta + n_2(u) \sin \theta] \quad (1)$$

where $n_1, n_2 : [0, 1] \rightarrow \mathbb{R}^3$ are normals to the curve c defined to be orthonormal, smooth, and such that the dot products $c'(u) \cdot n_i(u)$ vanish. See Fig. 1 for an illustration of a tubular surface. The idea is simply that the tubular surface is represented as a collection of circles each of which lie in the plane perpendicular to the center line. Note that the surface in (1) may thus be identified with a 4D space curve, $\tilde{c} : [0, 1] \rightarrow \mathbb{R}^4$, defined as a cross-product:

$$\tilde{c}(u) = (c(u), r(u))^T. \quad (2)$$

3.2 Variational Approach for Detecting the Cingulum Bundle

We now define a general class of energy functionals defined directly on 4D curves (2) that when optimized result in the 4D curve that represents the CB from DW-MRI of the brain.

Let $\mathbb{S}^2 \subset \mathbb{R}^3$ denote the 2D sphere, which is to represent the set of all possible angular acquisition directions of the scanning device for DW-MRI. Let $I : \mathbb{R}^3 \times \mathbb{S}^2 \rightarrow \mathbb{R}^+$ be the diffusion image. We are interested in weighted length functionals on 4D curves as energy functionals of interest. Indeed, let $\Psi : \mathbb{R}^4 \times \mathbb{S}^2 \rightarrow \mathbb{R}^+$ ($\Psi(x, r, v) \in \mathbb{R}^+$) be a weighting function, which we call the *potential* to be chosen, and define the energy as

$$E(\tilde{c}) = \int_{\tilde{c}} \Psi(\tilde{c}(\tilde{s}), \frac{c'(\tilde{s})}{|c'(\tilde{s})|}) d\tilde{s}, \quad \tilde{c} = (c, r) \quad (3)$$

where $d\tilde{s} = |\tilde{c}'(u)| du = \sqrt{(r'(u))^2 + |c'(u)|^2} du$ is the arclength measure of the 4D curve, and $c'(\tilde{s})/|c'(\tilde{s})|$ is the unit tangent to c , the center line. When (3) is minimized, the term $d\tilde{s}$ penalizes the non-smoothness of the center line and the radius function. The energy (3) is related to the length of a curve in a Finsler manifold [22].

The goal is to choose Ψ so that the energy is optimized by a \tilde{c} which determines a surface enclosing the diffusion pattern of the CB in the DW-MRI of the brain. The diffusion pattern in the cingulum varies in orientation and anisotropy across the length of the bundle, although *locally* similar (see Fig. 2), and that pattern differs from the pattern immediately outside the CB. This fact precludes the use of traditional region-based techniques adapted to DT-MRI since these techniques assume homogeneous statistics within the *entire* region enclosed by the surface, whereas we will assume homogeneity within local regions. In the next section, we show how Ψ may be chosen so that the energy can capture the varying diffusion pattern of the CB. The idea is to choose Ψ at a particular coordinate (x, r, p) to incorporate statistics of the DT-MRI *local* to the disc determined by (x, r, p) rather than using statistics global to the entire structure as in traditional region based methods.

3.3 Example Potentials, Ψ

In this section, we give two choices of Ψ that are meaningful for extracting the CB from DW-MRI, both based on *local* region-based statistics.

The first potential Ψ_1 at a coordinate $(x, r, v) \in \mathbb{R}^3 \times \mathbb{R}^+ \times \mathbb{S}^2$ is constructed so as to be small when the *mean diffusion profile* inside the disc, $D(x, r, v)$, differs greatly from the mean diffusion profile inside the annular region, $D(x, \alpha r, v) \setminus D(x, r, v)$ where $\alpha > 1$, outside $D(x, r, p)$. This is given by the following expressions:

$$\Psi_1(\tilde{p}, v) = \frac{1}{1 + \|\mu_{D(\tilde{p}, v)} - \mu_{D((p, \alpha r), v) \setminus D(\tilde{p}, v)}\|^2} \quad (4)$$

where the μ 's are means:

$$\mu_{D(\tilde{p}, v)}(\hat{v}) = \frac{1}{r^2} \int_{D(\tilde{p}, v)} I(x, \hat{v}) dA(x) \quad (5)$$

$$\mu_{D((p, \alpha r), v) \setminus D(\tilde{p}, v)}(\hat{v}) = \frac{1}{(\alpha^2 - 1)r^2} \int_{D((p, \alpha r), v) \setminus D(\tilde{p}, v)} I(x, \hat{v}) dA(x), \quad (6)$$

where dA is the area element and $\|\cdot\|$ is a suitable norm on functions of the form $f : \mathbb{S}^2 \rightarrow \mathbb{R}^+$, e.g.,

$$\|f_1 - f_2\|^2 = \int_{\mathbb{S}^2} |f_1(v) - f_2(v)|^2 dS(v), \quad (7)$$

where dS is the surface area element. The energy corresponding to Ψ_1 is *minimized*.

Another example potential is chosen such that the corresponding energy is related to a weighted surface area:

$$\Psi_2(x, r, p) = r \int_0^{2\pi} \phi(x + rp^\perp(\theta)) d\theta, \text{ and } p^\perp(\theta) = n_1 \cos \theta + n_2 \sin \theta \quad (8)$$

where n_1, n_2 are orthonormal vectors perpendicular to p , and $\phi : \mathbb{R}^3 \rightarrow \mathbb{R}^+$ is large near the boundary of differing diffusion regions, e.g.,

$$\phi(x) = \frac{1}{|B(x, R)|} \int_{B(x, R)} \|I(y, \cdot) - \mu_{B(x, R)}(\cdot)\|^2 dy \quad (9)$$

where $B(x, R)$ is the ball centered at x of chosen radius R , $|B(x, R)|$ denotes the volume, and the norm is defined as in (7). For this choice of potential, we are interested in *maximizing* the corresponding energy. The objective is to initialize the tubular surface inside the CB, and then increase surface area until the surface reaches the boundary of differing diffusion patterns.

4 Energy Optimization

In this section, we construct a steepest descent flow to minimize the energy of interest (3). A steepest descent is considered since we are not necessarily interested in the global maximizer or minimizer; indeed, the energy corresponding to (8) does not have a global maximizer. We begin with a tubular surface initialization (see Section 5 for the procedure), i.e., an initial 4D curve, and follow the gradient or its opposite depending on whether we want to maximize or minimize the energy.

4.1 Gradient Descent: Fixed Endpoints

The standard technique for calculating the gradient of an energy defined on curves, which is based on a geometrized \mathbb{L}^2 metric on the space of curves, cannot be applied to our energy of interest. This is because of the fact that when minimizing (3) using \mathbb{L}^2 , Ψ must satisfy a certain positivity condition (see [22]) that we cannot guarantee for our choices of Ψ otherwise the gradient descent is ill-posed. Moreover, when maximizing (3), we are indeed maximizing a weighted length, which with respect to the standard geometrized \mathbb{L}^2 curve metric, leads to an *unstable reverse diffusion*. As shown in [23], such weighted length functionals may be optimized in a *stable* manner by moving in the gradient direction of the energy (3) with respect to a *geometrized Sobolev metric*:

Definition 1. Let $\tilde{c} : [0, 1] \rightarrow \mathbb{R}^4$ be such that $\tilde{c}(0), \tilde{c}(1)$ are fixed. Let $h, k : [0, 1] \rightarrow \mathbb{R}^4$ be perturbations of \tilde{c} then

$$\begin{aligned}\langle h, k \rangle_{\mathbb{L}^2} &:= \frac{1}{L} \int_{\tilde{c}} h(\tilde{s}) \cdot k(\tilde{s}) \, d\tilde{s}, \\ \langle h, k \rangle_{\text{Sob}} &:= L \int_{\tilde{c}} h'(\tilde{s}) \cdot k'(\tilde{s}) \, d\tilde{s},\end{aligned}$$

where L is the length of the curve \tilde{c} , $d\tilde{s}$ is arclength element of \tilde{c} , and the derivatives are with respect to the arclength parameter \tilde{s} .

It can be shown that the gradient of (3) with respect to the Sobolev metric above is

$$\frac{1}{L} \nabla_{\text{Sob}} E(\tilde{c}) = K(\Psi_{\tilde{p}}) + \partial_{\tilde{s}} K(\hat{\Psi}_v \sqrt{1 + (r_{\tilde{s}}/|c_{\tilde{s}}|)^2} + \Psi \tilde{c}_{\tilde{s}}), \quad (10)$$

where

$$K(f) := \int_0^L K(\cdot, \tilde{s}) f(\tilde{s}) \, d\tilde{s}, \quad K(\tilde{s}_1, \tilde{s}_2) = \frac{1}{L} \begin{cases} \frac{\tilde{s}_2}{L} (1 - \frac{\tilde{s}_1}{L}) & 0 \leq \tilde{s}_2 \leq \tilde{s}_1 \\ \frac{\tilde{s}_1}{L} (1 - \frac{\tilde{s}_2}{L}) & \tilde{s}_1 \leq \tilde{s}_2 \leq L \end{cases}. \quad (11)$$

The expression has the additional numerical advantage that only first order derivatives are required in comparison to the standard \mathbb{L}^2 gradient, which needs second order information, and simply cannot be used anyway since it results in an unstable flow. Note that as stated in [23], the expression (10) may be computed efficiently in order N complexity, where N is the number of sample points of the curve.

4.2 Evolving Endpoints

In the previous subsection, we derived a gradient descent flow for (3) provided the endpoints of the 4D curve (i.e., the end cross sections of the tube) are fixed. We now describe how to evolve the endpoints so as to reduce the energy. This is useful for some choices of Ψ in (3), for example, Ψ_2 defined in (8). To determine the evolution of the endpoints, we compute the variation with respect to the endpoints. This results in

$$\tilde{c}_t(0) = \mp \hat{\Psi}_v \sqrt{1 + \left(\frac{r_{\tilde{s}}}{|c_{\tilde{s}}|}\right)^2} \mp \Psi \tilde{c}_{\tilde{s}}, \quad \tilde{c}_t(1) = \pm \hat{\Psi}_v \sqrt{1 + \left(\frac{r_{\tilde{s}}}{|c_{\tilde{s}}|}\right)^2} \pm \Psi \tilde{c}_{\tilde{s}}, \quad (12)$$

which will minimize/maximize the energy (depending on the sign chosen above). Therefore, the algorithm to reduce the energy is to alternatively evolve the endpoints by (12) and then evolve the 4D curve by (10).




Fig. 2. Selected slice-wise views of CB Segmentation results from proposed framework. The top row shows the DWI data and the bottom row shows the DWI data with the extracted surface rendered in 3D.

5 Experiments and Results

The algorithm was applied to DWI data of the brain from a data set that included schizophrenic and normal control subjects, with the DWI being acquired for 54 sampling directions. Results are included for the data from 2 subjects, and show the CB extracted for both the right and left bundles in each case.

In this paper, a perturbation of the anchor tract is used as the initial centerline curve and the smallest possible radius of 0.5 is used, with the surface essentially growing out from this initial radius. There are also other options to perform this initialization. Since we are given seed regions determined by an expert, an alternative initialization would be to connect the two seed regions with a streamline that passed above the Corpus Callosum (which is easy to segment). This initialization is being explored for future work.

The results included in this paper show the application of the proposed framework to the data sets, using the energy (3) using the potential Ψ_1 (4). Figure 2 shows slice-wise

views of the CB segmentation results obtained from the proposed framework indicating the homogeneity of the discs within the captured volume. The Figures, 3, and 4 show the tubular surface extracted by the proposed algorithm. It is to be noted that the surfaces are accurate while the boundaries shown in Figure 2 are the boundary locations rounded off to the grid points by the visualization process.




Fig. 3. CB Segmentation Results for Brain data set 1. Yellow shows the left CB and magenta shows the right CB.

6 Conclusions and Future work

We have proposed a novel technique to extract the cingulum bundle, which is of interest in the medical community because of its ties to schizophrenia, from DW-MRI of the brain. Unlike other standard techniques for extracting fiber bundles in the brain, we are able to *extract the entire bundle as a region at once rather* than detecting individual fibers and then combining them to form the bundle, which is laborious and prone to errors, while also performing this as a curve evolution rather than a surface evolution thus avoiding the computational disadvantages of a levelset implementation. We have *modeled the cingulum bundle as a tubular surface* and constructed a variational approach to detect the optimum tubular surface from DW-MRI, which represents the CB. Tubular surfaces provide a natural and accurate *shape prior* for the cingulum bundle, and such a shape prior is necessary due to the noisy nature of the imagery and the fact that data is not very visible or highly corrupted in certain slices. As we have shown, the tubular surface can be represented as a 4D curve, and thus, we were able to *significantly reduce the computational cost* of the algorithm compared to extracting an arbitrary surface. The proposed model was shown to yield good segmentations of the Cingulum Bundle upon visual inspection; unfortunately, there is no expert ground truth data available since it is laborious to hand segment an entire volume and certain slices do not even display the CB diffusion pattern accurately.

In future work, the authors plan to explore different choices of Ψ in the energy functional(3), and explore smoothness terms for the tubes in the energy. We will also




Fig. 4. CB Segmentation Results for Brain data set 2. Yellow shows the left CB and magenta shows the right CB.

implement the evolution of endpoints, which will be useful since the initialization will have to only be a single seed point. Further, the use of the extracted Cingulum Bundles will be explored in population studies for the discrimination of Schizophrenia. We are also interested in applying the framework other tubular structures such as the Uncinate Fasciculus in the brain.

This work was supported in part by grants from NSF, AFOSR, ARO, MURI, as well as by a grant from NIH (NAC P41 RR-13218) through Brigham and Women's Hospital. This work 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. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>.

References

1. Kubicki, M., Westin, C., Nestor, P., Wible, C., Frumin, M., Maier, S., Kikinis, R., Jolesz, F., McCarley, R., Shenton, M.: Cingulate fasciculus integrity disruption in schizophrenia: a magnetic resonance diffusion tensor imaging study. *Biological Psychiatry* **54**(11) (2003) 1171–1180
2. Wang, F., Sun, Z., Cui, L., Du, X., Wang, X., Zhang, H., Cong, Z., Hong, N., Zhang, D.: Anterior Cingulum Abnormalities in Male Patients With Schizophrenia Determined Through Diffusion Tensor Imaging (2004)
3. Mori, S., Crain, B., Chacko, V., van Zijl, P.: Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann. of Neurol.* **45**(2) (1999) 265–269
4. Conturo, T., Lori, N., Cull, T., Akbudak, E., Snyder, A., Shimony, J., McKinstry, R., Burton, H., Raichle, M.: Tracking neuronal fiber pathways in the living human brain. *Proc. Natl. Acad. Sci USA* **96**(18) (1999) 10422–10427

5. Perrin, M., Poupon, C., Cointepas, Y., Rieul, B., Golestani, N., Pallier, C., Rivière, D., Constantinesco, A., Bihan, D.L., Mangin, J.F.: Fiber tracking in q-ball fields using regularized particle trajectories. [24] 52–63
6. Parker, G.J.M., Wheeler-Kingshott, C.A.M., Barker, G.J.: Estimating distributed anatomical brain connectivity using fast marching methods and diffusion tensor imaging. *IEEE Trans. Med. Imaging* **21**(5) (2002) 505–512
7. Friman, O., Farneback, G., Westin, C.F.: A bayesian approach for stochastic white matter tractography. *IEEE Trans. Med. Imaging* **25**(8) (2006) 965–978
8. Lenglet, C., Rousson, M., Deriche, R., Faugeras, O.D., Lehericy, S., Ugurbil, K.: A riemannian approach to diffusion tensor images segmentation. [24] 591–602
9. Control Theory and Fast Marching Techniques for Brain Connectivity Mapping. In: *CVPR* (1). (2006)
10. Zhang, F., Goodlett, C., Hancock, E.R., Gerig, G.: Probabilistic fiber tracking using particle filtering. In Ayache, N., Ourselin, S., Maeder, A.J., eds.: *MICCAI* (2). Volume 4792 of *Lecture Notes in Computer Science.*, Springer (2007) 144–152
11. Cohen-Adad, J., Descoteaux, M., Rossignol, S., Hoge, R.D., Deriche, R., Benali, H.: Detection of multiple pathways in the spinal cord white matter using q-ball imaging. In: *IEEE International Symposium on Biomedical Imaging : From Nano to Macro.* (2008)
12. Xu, C., Pham, D.L., Prince, J.L.: Finding the brain cortex using fuzzy segmentation, isosurfaces, and deformable surface models. In Duncan, J.S., Gindi, G., eds.: *IPMI*. Volume 1230 of *Lecture Notes in Computer Science.*, Springer (1997) 399–404
13. Yezzi, A.J., Kichenassamy, S., Kumar, A., Olver, P.J., Tannenbaum, A.: A geometric snake model for segmentation of medical imagery. *IEEE Trans. Med. Imaging* **16**(2) (1997) 199–209
14. Lenglet, C., Rousson, M., Deriche, R.: Dti segmentation by statistical surface evolution. *IEEE Trans. Med. Imaging* **25**(6) (2006) 685–700
15. Wang, Z., Vemuri, B.C.: Dti segmentation using an information theoretic tensor dissimilarity measure. *IEEE Trans. Med. Imaging* **24**(10) (2005) 1267–1277
16. Awate, S.P., 0005, H.Z., Gee, J.C.: Fuzzy nonparametric dti segmentation for robust cingulum-tract extraction. In Ayache, N., Ourselin, S., Maeder, A.J., eds.: *MICCAI* (1). Volume 4791 of *Lecture Notes in Computer Science.*, Springer (2007) 294–301
17. Melonakos, J., Mohan, V., Niethammer, M., Smith, K., Kubicki, M., Tannenbaum, A.: Finsler tractography for white matter connectivity analysis of the cingulum bundle. In: *MICCAI* (1). (2007) 36–43
18. Melonakos, J., Niethammer, M., Mohan, V., Smith, K., Kubicki, M., Tannenbaum, A.: Locally-constrained region-based methods for dw-mri segmentation. In: *MMBIA*. (2007)
19. Li, H., Yezzi, A.: Vessels as 4-d curves: Global minimal 4-d paths to extract 3-d tubular surfaces and centerlines. *IEEE Transactions on Medical Imaging* **26**(9) (2007) 1213–1223
20. Cohen, L.D., Kimmel, R.: Global minimum for active contour models: A minimal path approach. In: *CVPR*, *IEEE Computer Society* (1996) 666–673
21. Schmahmann, J., Pandya, D.: *Fiber Pathways of the Brain*. Oxford University Press (2006)
22. Melonakos, J., Pichon, E., Angenent, S., Tannenbaum, A.: Finsler active contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (To Appear)
23. Sundaramoorthi, G., Yezzi, A.J., Mennucci, A., Sapiro, G.: New possibilities with Sobolev active contours. In: *SSVM*. (2007) 153–164
24. Christensen, G.E., Sonka, M., eds.: *Information Processing in Medical Imaging*, 19th International Conference, *IPMI 2005*, Glenwood Springs, CO, USA, July 10-15, 2005, Proceedings. In Christensen, G.E., Sonka, M., eds.: *IPMI*. Volume 3565 of *Lecture Notes in Computer Science.*, Springer (2005)

Modeling the Remaining Flexibility of Partially Fixed Statistical Shape Models

Thomas Albrecht, Reinhard Knothe, and Thomas Vetter

Computer Science Department, University of Basel, Switzerland
{Thomas.Albrecht, Reinhard.Knothe, Thomas.Vetter}@unibas.ch

Abstract. Statistical shape models are widely used to model the variability of biological shapes. They can be used to reconstruct missing information given partial or noisy data. In case of partial data, many different reconstructions are possible, and one is not only interested in a plausible reconstruction but also the remaining flexibility within the model and the reliability of the reconstruction. We present a method to model the remaining flexibility when some part of a statistical shape model is fixed. Using such a flexibility model, we can give answers to questions like: Does one half of a human femur bone determine the other half? or How much is the shape of a face determined by its contour?

1 Introduction

Statistical shape models are widely used in medical image analysis, computational anatomy, and computer vision to model the variability of biological shapes, see [1–6] for instance. The variability of a certain class of shapes is deduced from a representative set of example data from this class.

If the example data sets represent the class of shapes well, the model can be fitted to virtually any individual shape within the class. It is even possible to fit the model to partial data from an individual, [7, 8]. The missing remaining data is automatically reconstructed by the model, yielding the most plausible reconstruction.

However, there may be many other possible reconstructions which fit the partial data equally well. The partial data only determines a part of the model, while the rest of the model may remain flexible. This paper focuses on modeling this remaining flexibility of the partially determined model.

Keeping a part \mathbf{x}_b of the complete model \mathbf{x} fixed, how much flexibility remains for the remaining part \mathbf{x}_a ? In principle, as PCA models can be statistically interpreted by a multivariate Gaussian distribution with probability density function p , we can model the remaining flexibility by the conditional distribution $p(\mathbf{x}_a|\mathbf{x}_b)$. We will see however, that there is no nontrivial conditional distribution if more components are kept fixed than there are degrees of freedom in the model, as is usually the case in models built from a small set of examples.

In this case, we propose a method which models the remaining flexibility of the variable points when the fixed points are allowed to move slightly instead of

being completely fixed. The method leads to a generalized eigenvalue problem which can be solved efficiently.

The remaining paper is organized as follows. In Section 2 we will review the well-known concept of PCA-based statistical shape models. In Section 3 we will derive an expression of the conditional distribution $p(\mathbf{x}_a|\mathbf{x}_b)$ and investigate when this can actually be used. In Section 4, we will model the remaining flexibility when the fixed points are allowed to move slightly. In Section 5 we will use the proposed method to model the remaining flexibility of a shape model of the human femur bone when its distal part is fixed, and the remaining flexibility of a face model, when its contour is fixed.

Related Work. Virtually all cited papers deal with statistical shape models with applications in medicine or computer vision, [1–8]. Most notably, [5] uses Canonical Correlation Analysis to model the correlation between different parts of a statistical model, computing model coefficients which maximize the correlation between these parts. However, they do not address our problem of modeling the remaining flexibility when one part of the model is fixed.

2 Principle Component Analysis

The 3D statistical shape models used in this paper are based on a Principle Component Analysis (PCA) of a set of training data comprised of n 3D surfaces. Each surface is represented by a triangular mesh with the same number $m \in \mathbb{N}$ of vertices, which are stacked into a data vector $\mathbf{x} = (x_1, y_1, z_1, \dots, x_m, y_m, z_m)^T \in \mathbb{R}^M$ with $M = 3m$. The training surfaces need to be in correspondence.

As in all PCA models, the samples are assumed to be i.i.d. samples drawn from a multivariate normal distribution that is approximated by the estimated multivariate normal distribution $\mathcal{N}(\bar{\mathbf{x}}, \Sigma)$. The mean is estimated by the arithmetic mean $\bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$ of all samples. The covariance matrix $\Sigma \in \mathbb{R}^{M \times M}$ can be estimated from the mean-free data matrix $\mathbf{X} := [\mathbf{x}_1 - \bar{\mathbf{x}}, \dots, \mathbf{x}_n - \bar{\mathbf{x}}] \in \mathbb{R}^{M \times n}$ as $\Sigma = \frac{1}{n} \mathbf{X} \mathbf{X}^T$.

When \mathbf{X} is decomposed with a (reduced) Singular Value Decomposition $\mathbf{X} = \mathbf{U} \mathbf{W} \mathbf{V}^T$ into the product of a column-orthonormal matrix $\mathbf{U} \in \mathbb{R}^{M \times n}$, a diagonal matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$, and an orthonormal matrix $\mathbf{V} \in \mathbb{R}^{n \times n}$, the covariance matrix can be expressed as $\Sigma = \frac{1}{n} \mathbf{U} \mathbf{W}^2 \mathbf{U}^T$.

The columns \mathbf{u}_i of the matrix \mathbf{U} are the eigenvectors of Σ . They are known as the principal components of the model and describe the main modes of variation of the training data. Their corresponding eigenvalues $\sigma_i^2 := \frac{1}{n} w_i^2$ describe the variance of the model projected onto these eigenvalues. They are arranged according to size so that \mathbf{u}_1 is the direction with maximal projected variance σ_1^2 .

One individual \mathbf{x} in the object class modeled by the PCA model can be identified by its coefficients $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)^T \in \mathbb{R}^n$:

$$\mathbf{x} = \bar{\mathbf{x}} + \sum_{i=1}^n \alpha_i \sigma_i \mathbf{u}_i = \bar{\mathbf{x}} + \frac{1}{\sqrt{n}} \mathbf{U} \mathbf{W} \boldsymbol{\alpha}. \quad (1)$$

Under the assumption that the data \mathbf{x} is distributed according to a multivariate normal distribution $\mathcal{N}(\bar{\mathbf{x}}, \Sigma)$, the coefficient vector α is distributed according to $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$.

In this paper two PCA models are used: A model of the femur bone, built from 21 CT scans of femur bones, which were hand-segmented and brought into correspondence with [9], and a model of the human face built from 100 face scans that were brought in correspondence with a modification of [10].

3 Conditional Distribution

In a PCA-based statistical model, we now wish to fix a certain number $l \in \mathbb{N}$ of points in the model in order to investigate how flexible the model remains with respect to the remaining $m - l$ variable points. This means that we fix $L = 3l$ components of the model vector \mathbf{x} . Without loss of generality, we can assume that they are the last l components and \mathbf{x} can be partitioned as $\mathbf{x} = (\mathbf{x}_a, \mathbf{x}_b)^T$.

As \mathbf{x} is distributed according to a multivariate normal distribution, the conditional distribution is also a multivariate normal distribution $\mathcal{N}(\mu_{a|b}, \Sigma_{a|b})$. Its mean and covariance can be calculated from $\bar{\mathbf{x}}$ and Σ .

The matrix of principal components $\mathbf{U} \in \mathbb{R}^{M \times n}$ can be partitioned according to \mathbf{x} . For simplicity's sake, we define the matrix of the principal components scaled by the diagonal matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$ as $\mathbf{Q} := \frac{1}{\sqrt{n}} \mathbf{U} \mathbf{W} \in \mathbb{R}^{M \times n}$. We have

$$\mathbf{U} = \begin{pmatrix} \mathbf{U}_a \\ \mathbf{U}_b \end{pmatrix}, \quad \mathbf{Q} = \frac{1}{\sqrt{n}} \mathbf{U} \mathbf{W} = \frac{1}{\sqrt{n}} \begin{pmatrix} \mathbf{U}_a \mathbf{W} \\ \mathbf{U}_b \mathbf{W} \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_a \\ \mathbf{Q}_b \end{pmatrix} \quad (2)$$

The covariance matrix $\Sigma \in \mathbb{R}^{M \times M}$ can be calculated from \mathbf{Q} by $\Sigma = \frac{1}{n} \mathbf{U} \mathbf{W}^2 \mathbf{U}^T = \mathbf{Q} \mathbf{Q}^T$. We can partition Σ as follows:

$$\Sigma = \begin{pmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{pmatrix} = \begin{pmatrix} \mathbf{Q}_a \mathbf{Q}_a^T & \mathbf{Q}_a \mathbf{Q}_b^T \\ \mathbf{Q}_b \mathbf{Q}_a^T & \mathbf{Q}_b \mathbf{Q}_b^T \end{pmatrix}. \quad (3)$$

Note that it is only possible to calculate a nontrivial conditional distribution if the matrix Σ_{bb} is invertible. A necessary condition for Σ_{bb} to be invertible is that the number of fixed degrees of freedom L is less the number of training examples n of the statistical model. Let us, for the moment, assume that Σ_{bb} is invertible and calculate the conditional distribution according. According to [11], the covariance matrix $\Sigma_{a|b}$ of the can be expressed as:

$$\Sigma_{a|b} = \Sigma_{aa} - \Sigma_{ab} \Sigma_{bb}^{-1} \Sigma_{ba} \quad (4)$$

$$= \mathbf{Q}_a \mathbf{Q}_a^T - \mathbf{Q}_a \mathbf{Q}_b^T (\mathbf{Q}_b \mathbf{Q}_b^T)^{-1} \mathbf{Q}_b \mathbf{Q}_a^T \quad (5)$$

$$= \mathbf{Q}_a \left(\mathbf{I}_n - \mathbf{Q}_b^T (\mathbf{Q}_b \mathbf{Q}_b^T)^{-1} \mathbf{Q}_b \right) \mathbf{Q}_a^T, \quad (6)$$

where \mathbf{I}_n denotes the $(n \times n)$ identity matrix. While the original matrix \mathbf{U} is column-orthonormal, the sub-matrix \mathbf{U}_b defined in Equation (2) is *not*. Therefore, we perform an additional singular value decomposition $\mathbf{Q}_b = \bar{\mathbf{U}}_b \mathbf{W}_b \mathbf{V}_b^T$ with a column-orthonormal $\bar{\mathbf{U}}_b$. Thus, the expression can be expanded to:

$$\Sigma_{a|b} = \mathbf{Q}_a \left(\mathbf{I}_n - \mathbf{V}_b \mathbf{W}_b \bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b \mathbf{W}_b^{-2} \bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b \mathbf{W}_b \mathbf{V}_b \right) \mathbf{Q}_a^T \quad (7)$$

$$= \mathbf{Q}_a \left(\mathbf{I}_n - \mathbf{V}_b^T \mathbf{V}_b \right) \mathbf{Q}_a^T. \quad (8)$$

Where, in the last step, we have used $\bar{\mathbf{U}}_b^T \bar{\mathbf{U}}_b = \mathbf{I}_n$. Similarly, the mean $\mu_{a|b}$ is given as:

$$\mu_{a|b} = \bar{\mathbf{x}}_a + \Sigma_{ab} \Sigma_{bb}^{-1} (\mathbf{x}_b - \bar{\mathbf{x}}_b) \quad (9)$$

$$= \bar{\mathbf{x}}_a + \mathbf{Q}_a \mathbf{V}_b \mathbf{W}_b^{-1} \bar{\mathbf{U}}_b^T (\mathbf{x}_b - \bar{\mathbf{x}}_b), \quad (10)$$

with $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_a, \bar{\mathbf{x}}_b)^T$.

Provided that Σ_{bb} is invertible, \mathbf{x}_a can be reconstructed from \mathbf{x}_b as $\mu_{a|b}$. This is the reconstruction presented in [7] as the maximum a posteriori reconstruction. All other reconstructions which fit \mathbf{x}_b are modeled by $\mathcal{N}(\mu_{a|b}, \Sigma_{a|b})$. The more flexibility this distribution allows, the less reliable the reconstruction by $\mu_{a|b}$ is.

For each of these reconstructions, the fixed values \mathbf{x}_b are matched equally well. Indeed, if we take a closer look at the expression for the covariance matrix $\Sigma_{a|b}$ in Equation (8) we notice that the inner part $(\mathbf{I}_n - \mathbf{V}_b^T \mathbf{V}_b)$ is a projection onto the orthogonal complement of the column space of \mathbf{V}_b . As \mathbf{V}_b is the “input matrix” of the SVD of \mathbf{Q}_b , this projection is a projection onto the kernel of \mathbf{Q}_b . This means that the distribution $\mathcal{N}(\mu_{a|b}, \Sigma_{a|b})$ models only linear combinations $\mathbf{Q}_a \alpha$ of the scaled principal components for which α is in the kernel of \mathbf{Q}_b , i.e. $\mathbf{Q}_b \alpha = 0$. So for all $\alpha \sim \mathcal{N}(\mu_{a|b}, \Sigma_{a|b})$, the deformation of the fixed points is zero.

Invertibility of Σ_{bb} . So far, we have assumed that $\Sigma_{bb} = \mathbf{Q}_b \mathbf{Q}_b^T \in \mathbb{R}^{L \times L}$ is invertible, i.e. $\text{rank}(\Sigma_{bb}) = L$. However, if \mathbf{Q}_b has less than L linearly independent columns we have $\text{rank}(\Sigma_{bb}) < L$. In particular, if the statistical model is built from less than L examples or less than L principal components are used, the rank of Σ_{bb} will be less than L and the above calculations involving an inversion of Σ_{bb} are not valid. In this case there is no nontrivial conditional distribution $p(\mathbf{x}_a, \mathbf{x}_b)$.

In this case the mean $\mu_{a|b}$ can still be approximated by using the pseudo-inverse of Σ_{bb} instead of its inverse, [7]. However, if we try to use the pseudo-inverse for calculating the covariance matrix according to Equation (4), we get $\Sigma_{a|b} = 0$.

The problem is that the conditional distribution $\mathcal{N}(\mu_{a|b}, \Sigma_{a|b})$ models only coefficient vectors from the kernel of \mathbf{Q}_b . Without a nontrivial kernel of \mathbf{Q}_b , it is not possible to calculate a nontrivial conditional distribution. Therefore, we propose that instead of looking for coefficient vectors $\alpha \in \ker \mathbf{Q}_b$, i.e. $\mathbf{Q}_b \alpha = 0$,

we will look for coefficients for which $\mathbf{Q}_b \boldsymbol{\alpha}$ is small. This means that we relax the constraint of keeping the fixed points completely fixed to allowing them to move slightly.

It also has to be noted that even if there is a nontrivial kernel of \mathbf{Q}_b , it will most likely still be of interest to consider deformations which allow slight changes of the fixed points \mathbf{x}_b when considering the remaining flexibility of the model.

4 Generalized Eigenvalues

The aim is to model the flexibility of the variable points \mathbf{x}_a when the fixed points \mathbf{x}_b are allowed to move slightly. The deformations are given as $\mathbf{Q}_a \boldsymbol{\alpha}$ and $\mathbf{Q}_b \boldsymbol{\alpha}$. A measure for the change caused by these deformations is the squared Euclidean norm of these vectors:

$$\|\mathbf{Q}_a \boldsymbol{\alpha}\|^2 = \boldsymbol{\alpha}^T \mathbf{Q}_a^T \mathbf{Q}_a \boldsymbol{\alpha}, \quad \|\mathbf{Q}_b \boldsymbol{\alpha}\|^2 = \boldsymbol{\alpha}^T \mathbf{Q}_b^T \mathbf{Q}_b \boldsymbol{\alpha}. \quad (11)$$

As we are first and foremost interested in the coefficients $\boldsymbol{\alpha}$ which change the variable points \mathbf{x}_a as much as possible, we can formulate our aim as a constrained maximization problem:

$$\max_{\boldsymbol{\alpha} \in \mathbb{R}^k} \boldsymbol{\alpha}^T \mathbf{Q}_a^T \mathbf{Q}_a \boldsymbol{\alpha} \quad (12)$$

$$\text{subject to } \boldsymbol{\alpha}^T \mathbf{Q}_b^T \mathbf{Q}_b \boldsymbol{\alpha} = c, \quad (13)$$

where $c \in \mathbb{R}^+$ quantifies the amount of change allowed in the fixed coefficients.

Introducing a Lagrangian multiplier λ and differentiating with respect to $\boldsymbol{\alpha}$ leads to the generalized eigenvalue problem:

$$\mathbf{Q}_a^T \mathbf{Q}_a \boldsymbol{\alpha} = \lambda \mathbf{Q}_b^T \mathbf{Q}_b \boldsymbol{\alpha}. \quad (14)$$

Both matrices $\mathbf{Q}_a^T \mathbf{Q}_a$ and $\mathbf{Q}_b^T \mathbf{Q}_b$ are positive definite and symmetric. The generalized eigenvalue problem can be solved efficiently with standard software (LAPACK, MATLAB), yielding a set of generalized eigenvectors $\{\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_k\}$ arranged according to the size of their corresponding generalized eigenvalues $\{\lambda_1, \dots, \lambda_k\}$. The eigenvectors are scaled so that $\|\mathbf{Q}_b \boldsymbol{\alpha}_i\|^2 = \boldsymbol{\alpha}_i^T \mathbf{Q}_b^T \mathbf{Q}_b \boldsymbol{\alpha}_i = 1$ for $i = 1, \dots, k$. If we pre-multiply Equation (14) by $\boldsymbol{\alpha}_i^T$, we see that for an eigenvector $\boldsymbol{\alpha}_i$ and its eigenvalue λ_i , we have

$$\|\mathbf{Q}_a \boldsymbol{\alpha}_i\|^2 = \lambda_i \|\mathbf{Q}_b \boldsymbol{\alpha}_i\|^2. \quad (15)$$

This means that, measured in the squared Euclidean norm, the deformation determined by the coefficient vector $\boldsymbol{\alpha}_i \in \mathbb{R}^k$ changes the variable model points λ_i times as much the fixed ones. Therefore, the eigenvector $\boldsymbol{\alpha}_1$ corresponding to the largest eigenvalue λ_1 is the coefficient vector which causes the largest change on the variable points (with a squared Euclidean norm of λ_1), changing the fixed points only slightly (with a deformation with a squared Euclidean norm of 1). The last eigenvectors change the fixed points more than the variable ones. In




Fig. 1. In a statistical shape model of the human femur bone, the two first flexibility components α_1 and α_2 model those deformations that change the proximal part (gray) as much as possible, while changing the distal part (colored) as little as possible.

fact, exchanging fixed and variable points leads to the same eigenvalues, only inverted.

We call the coefficient vectors α_i the flexibility components. α_1 is the first flexibility component. Just like the principal components of the PCA model, they can be scaled and added together to achieve different deformations from the mean. When a flexibility component α_i is scaled by \sqrt{r} it causes a deformation of squared Euclidean norm r of the fixed points and of $\lambda_i r$ of the variable points. The larger the eigenvalues λ_i are, the more flexibility remains for the fixed points and the less reliable a reconstruction of \mathbf{x}_a from \mathbf{x}_b is considered. In order to meet the constraint from Equation (13), the vectors have to be scaled by \sqrt{c} .

The squared Euclidean norm is not a very intuitive measure of the deformation as it implicitly depends on the number l of fixed points. A more intuitive measure is the mean squared norm, averaged over all l fixed points:

$$\frac{1}{l} \sum_{k=1}^l (\mathbf{Q}_b \alpha_i)_k^2 = \frac{1}{l} \|\mathbf{Q}_b \alpha_i\|^2. \quad (16)$$

Therefore, if we scale the coefficient vectors α_i by \sqrt{l} , and the coordinates of the model are given in millimeters, the squared deformation of one model point is 1 millimeter on average.




Fig. 2. For the examples shown in Figures 4 and 1, the deformation of the fixed vertices (blue) is compared with the flexibility of the full shape (red). The first flexibility components strongly affect the full shape, while they hardly alter the fixed vertices.

Regularization. The flexibility components α_i are calculated as generalized eigenvectors. As outlined in Section 2, the coefficients α of the PCA model are distributed according to a multivariate normal distribution $\mathcal{N}(0, \mathbf{I}_n)$. Without regularization, it is possible that some of the entries of α_i can be extremely large. In terms of the distribution $\mathcal{N}(0, \mathbf{I}_n)$, such extremely large values are very unlikely and in practice they cause unnatural deformations of the modeled shape.

Up to a normalizing factor, the probability of a deformation caused by α can be calculated by $e^{-\frac{1}{2}\|\alpha\|^2}$. By minimizing $\|\alpha\|^2 = \alpha^T \alpha$, the probability of the deformation is maximized. Therefore, we propose replacing the constraint in Equation (13) by the following constraint:

$$\alpha^T Q_b^T Q_b \alpha + \eta \alpha^T \alpha = c, \quad (17)$$

with a regularizing parameter $\eta \in \mathbb{R}^+$. In this way, we limit not only the deformation of the fixed points given by $\|Q_b \alpha\|^2$ but also the improbability of the deformation, given by $\|\alpha\|^2$. The corresponding generalized eigenvalue problem is a regularized version of Equation (14) and is given by:

$$Q_a^T Q_a \alpha = \lambda (Q_b^T Q_b + \eta \mathbf{I}_n) \alpha. \quad (18)$$

This system no longer admits extremely large values in the generalized eigenvectors α_i . The parameter η controls the balance between the original problem and the regularizing effect of allowing only probable shapes.

Note that when we deform not the mean but a certain individual \mathbf{x} , which is defined by coefficients β , the deformation coefficients α are not distributed according to $\mathcal{N}(0, \mathbf{I}_n)$ but according to $\mathcal{N}(-\beta, \mathbf{I}_n)$. Therefore it would make more sense to minimize $\|\alpha + \beta\|^2 = (\alpha + \beta)^T (\alpha + \beta)$. However, it is not obvious how to include such a constraint into a generalized eigenvalue problem like




Fig. 3. The vertices on the contour in profile view are fixed (a). The 1st (b) and 2nd (c) flexibility components (see Figure 4) don't change the contour in profile view.

Equation (18). Therefore, we content ourselves with the proposed regularization, which penalizes large deformations, even though its statistical motivation is perfectible.

Nontrivial kernel of \mathbf{Q}_b . We have introduced the generalized eigenvalue problem Equation (14) in order to model the remaining flexibility of the variable points when it is not possible to calculate the conditional distribution $p(\mathbf{x}_a|\mathbf{x}_b)$, which is the case if and only if \mathbf{Q}_b has only the trivial kernel $\ker \mathbf{Q}_b = \{\mathbf{0}\}$.

But what happens to the generalized eigenvalue problem if \mathbf{Q}_b does have a nontrivial kernel? In this case, we can split up the space of all coefficients α into the kernel of \mathbf{Q}_b and its orthogonal complement. In the kernel, we have all deformations which do not change the fixed points at all, which can be modeled by the conditional distribution. The covariance matrix in the kernel is simply $\mathbf{Q}_a \mathbf{Q}_a^T$, as in the kernel, the projection term from Equation (8), $(\mathbf{I}_n - \mathbf{V}_b^T \mathbf{V}_b) = \mathbf{I}_n$. In the complement, we can compute the generalized eigenvalue problem in order to additionally allow deformations which change the fixed points slightly.

From a practical point of view, when $\ker \mathbf{Q}_b \neq \{\mathbf{0}\}$ and therefore $\mathbf{Q}_b^T \mathbf{Q}_b$ is singular, the aforementioned LAPACK or MATLAB routines return a basis for the kernel as generalized eigenvectors with eigenvalues infinity, as for these vectors we have $\mathbf{Q}_b^T \mathbf{Q}_b \alpha_i = 0$. The remaining eigenvalues span the complement of the kernel and are computed as usual.

If we use the regularized form of the generalized eigenvalue problem Equation (18), there will be no infinite eigenvalues, as the matrix $(\mathbf{Q}_b^T \mathbf{Q}_b + \eta \mathbf{I}_n)$ is always nonsingular. However, for vectors $\alpha \in \ker \mathbf{Q}_b$, we have $(\mathbf{Q}_b^T \mathbf{Q}_b + \eta \mathbf{I}_n) \alpha = \eta \alpha$. Therefore, on $\ker \mathbf{Q}_b$, we effectively solve the eigenvalue problem:

$$\alpha^T \mathbf{Q}_a^T \mathbf{Q}_a \alpha = \lambda \eta \alpha, \quad (19)$$

whose eigenvectors corresponding to the largest eigenvalues maximize $\|\mathbf{Q}_a\|^2$. In this regularized case, there is no strict decomposition into $\ker \mathbf{Q}_b$ and its complement.




Fig. 4. A statistical shape model of human faces is used to model the remaining flexibility with fixed contour. The vertices on the contour in profile view (see also Figure 3) are fixed. The flexibility components model the remaining variability within the model. The 1st/2nd flexibility components is shown in the 1st/2nd row.

5 Experiments and Conclusion

We demonstrate the use of the flexibility models for two different scenarios with two different models, one for the shape of human faces and one for the human femur bone.

In the case of the femur bone, we are interested in determining how well the distal (bottom) part of the femur determines the proximal (top) part, in order to estimate how reliable the model can be for reconstruction of missing or injured parts of the bone. In Figure 1, the distal part of the bone is colored. It corresponds to the fixed points \mathbf{x}_b . From the generalized eigenvalue problem Equation (18), the flexibility components α_i are calculated. The regularization parameter was chosen as $\eta = 10$. Figure 1 illustrates the effect of deforming the mean by $\pm 2.2\sqrt{l}$ times the first two flexibility components α_1 and α_2 . Clearly, both the distal and the proximal part are changed, but the proximal part is much more heavily deformed. In Figure 2, the amount of deformation is plotted, measured in the mean and the maximum of the Euclidean norm of the deformation at each point. For the mean norm, the ratio is approximately equal to

the corresponding eigenvalue. For the mean squared norm, which is not plotted here, it is of course exactly equal to this eigenvalue.

In case of the face model we are interested in the question: Given the vertices on the contour, how much is the shape of the face determined? Here, the fixed vertices \mathbf{x}_b are the vertices of the occluding contour in profile view, as shown in Figure (3 a)). This is the contour for one individual represented by its model coefficients β . Again, the flexibility components α_i are computed using Equation (18), with $\eta = 10$. To visualize the result, the surface with the coefficients $\beta \pm 1.1\sqrt{\lambda}\alpha_i$ is shown in Figure 4. We see that the first two flexibility components heavily deform the model, while the vertices at the contour are almost fixed, see Figure 3 b,c). The resulting deformations measures are plotted in Figure 2. The ratio is much higher than in the femur case as less points are kept fixed and the model is built from many more examples, making it more expressive.

Conclusion. We have introduced a way to model the remaining flexibility of a statistical shape model when a part of the model is kept as fixed as possible, even in the absence of a nontrivial conditional distribution. In future work, we will apply this technique in the fitting of statistical models to partial data.

References

1. Fleute, M., Lavalée, S.: Building a complete surface model from sparse data using statistical shape models: Application to computer assisted knee surgery. *Medical Image Computing and Computer-Assisted Intervention-MICCAI* **98** 880–7
2. Rajamani, K., Styner, M., Haydar, T., Nolte, L., Gonzalez Ballester, M.A.: Statistical deformable bone models for robust 3d surface extrapolation from sparse data. submitted to Elsevier Science (2006)
3. Subsol, G., Thirion, J., Ayache, N.: A scheme for automatically building three-dimensional morphometric anatomical atlases: application to a skull atlas. *Medical Image Analysis* **2**(1) (1998) 37–60
4. Dam, E., Fletcher, P., Pizer, S.: Automatic shape model building based on principal geodesic analysis bootstrapping. *Medical Image Analysis* (2007)
5. Rao, A., Babalola, K., Rueckert, D.: Canonical Correlation Analysis of Sub-cortical Brain Structures Using Non-rigid Registration. *Lecture Notes in Computer Science* **4057** (2006) 66
6. Shen, D., Herskovits, E., Davatzikos, C.: An adaptive-focus statistical shape model for segmentation and shape modeling of 3-D brain structures. *IEEE Transactions on Medical Imaging* **20**(4) (2001) 257–270
7. Blanz, V., Vetter, T.: Reconstructing the complete 3d shape of faces from partial information. *Informationstechnik und Technische Informatik* **44**(6) (2002) 1–8
8. Keller, M., Knothe, R., Vetter, T.: 3D Reconstruction of Human Faces from Occluding Contours. *Lecture Notes in Computer Science* **4418** (2007) 261
9. Dedner, A., Lüthi, M., Albrecht, T., Vetter, T.: Curvature guided level set registration using adaptive finite elements. In: *Pattern Recognition*. (2007) 527–536
10. Amberg, B., Romdhani, S., Vetter, T.: Optimal Step Nonrigid ICP Algorithms for Surface Registration. *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on* (2007) 1–8
11. Bishop, C.: *Pattern recognition and machine learning*. Springer (2006)

A Hypothesis Testing Framework for High-Dimensional Shape Models

Joshua Cates, P. Thomas Fletcher, Ross Whitaker

Scientific Computing and Imaging Institute
University of Utah
Salt Lake City, Utah

Abstract. Statistical shape models are powerful tools for describing anatomical structures and are increasingly being used in a wide variety of clinical and biological contexts. One of the promising applications of this technology is the testing of hypotheses that entail shape differences, and visualization of those differences between cohorts. Statistical testing of shapes, however, is difficult due the large numbers of degrees of freedom and challenge of obtaining sufficient numbers of subjects to ensure statistical power. To date, research in statistical shape modeling has focused mainly on the construction of representative models, and the field has not reached a consensus on the best approach to statistical hypothesis testing. This paper illustrates some problems inherent in the statistical analysis of high-dimensional shape models, and suggests a systematic approach to hypothesis testing that avoids those problems. The proposed framework is based on research in the factor analysis statistics literature, with permutation testing in the PCA space of the model, and dimensionality reduction via a simulation-based analysis. We also describe two methods for visualizing group mean differences, first by direct visualization of the linear discriminant implicit in the hypothesis test metric, and second, by visualizing strain tensors from a deformation computed between the group means. We illustrate the proposed analysis and visualization framework on several clinical and biological datasets.

1 Introduction

Many important fields of basic research in medicine and biology now routinely employ statistical models of shape to quantify the anatomical variation in populations. Often, researchers are also interested in hypothesis testing to evaluate or demonstrate shape differences between populations. One such application area, for example, is the study of gene function as it pertains to human development and disease. Modern gene targeting technology allows researchers to create specific alterations in a mouse genome that result in different patterns of anatomical growth and form, or *phenotypes*, which can be modeled as shape and compared with normal populations to gain insight into the functionality of the targeted genes [1, 2]. Many areas of clinical psychiatric and neurological research also employ statistical shape analysis. The study of autism and its impact on brain

regions is one notable example [3, 2]. Statistical models that capture the variation in the shape of brain structure, including the *covariation* among multiple substructures, are increasingly necessary for researchers to gain understand into the development and progression of neurological disorders [4, 5].

Anatomical shape from images can be represented and computed using a variety of tools. Traditional representations of shape for phenotyping, for example, have relied on explicitly chosen *landmark* positions to define relatively *low-dimensional* parameterized models (e.g. [6]). Researchers in the 3D medical imaging community, however, have more recently pioneered the use of *high-dimensional* models of shape, which consist of very large collections of surface parameters that are derived automatically from the 3D images. High-dimensional shape models are appealing because they provide a much more detailed description of anatomy than landmark-based models, and do not require an a-priori choice of surface homologies. One common strategy for high-dimensional shape modeling is to consider shapes as embedded in images, and then to deformably register images and perform statistics on those deformations (e.g. [7]). Another common approach is to construct parameterized or point-based shape descriptors and compute statistics on those descriptions (e.g. [4, 3]). In the latter case, technologies exist to construct compact models that are optimized with respect to the information content of the population [8, 2], an important consideration for subsequent statistical analysis.

While high-dimensional models offer significant advantages for shape representation, their statistical analysis is not straightforward. The large number of degrees of freedom in the shape space, often coupled with a relatively low sample size (HDLSS), means that traditional low-dimensional statistical metrics cannot be directly applied [9]. While the shape modeling literature has proposed methods for analysis, it has not reached a consensus regarding a systematic approach that addresses the HDLSS problem. The statistics literature addresses the HDLSS problem with a variety of useful techniques, but these tools must be applied systematically in order to avoid either under-powered studies or over-optimistic conclusions. Through examples on clinical and biological datasets, this paper illustrates some of the potential difficulties that are encountered in high-dimensional shape analysis. We focus on the problems of hypothesis testing for group differences, and the visualization of those group differences. For hypothesis testing, we suggest permutation testing in a lower-dimensional PCA subspace of the model. For the dimensionality reduction, we propose using a simulation-based method to choose dimensions whose variance is distinguishable from noise. For visualization of group differences, we describe two approaches. The first is a direct visualization of the linear discriminant vector implicit in the hypothesis test, and the second is a visualization of strain tensors derived from a thin-plate spline deformation between the group mean shapes.

2 Related Work

Hypothesis testing for group differences using high-dimensional shape models has been most extensively investigated for comparative studies of brain anatomy. To date, however, researchers have typically each chosen different strategies for analysis, and there has not been a systematic treatment of the effects of HDLSS and dimensionality reduction choices on the statistical tests. Davies, et al. investigate hippocampus shape in schizophrenia populations with spherical harmonic and MDL-based shape models, and perform classification trials based on linear discriminant analysis in the high-dimensional shape space [8]. Terriberry [10] proposes a framework for multivariate, permutation-based hypothesis testing using nonlinear m-rep shape models, with an application to the lateral ventricles of twins. Styner, et al. propose point-wise hypothesis testing on shape correspondence points derived from spherical harmonic parameterizations [11, 12]. Golland [13] proposes a Support Vector Machine algorithm for training group classifiers of distance-transform shape representations. More recently, Gorczowski, et al. use medial axis shape representations and distance-weighted discriminants to compare complexes of brain structures in autistic and normal populations [3].

Several researchers have explored hypothesis testing using high-dimensional models for mouse phenotyping. Chen, et al. perform hypothesis testing using image-based metrics on deformation-based shape models to illustrate differences between strains of mutant mice [7]. In a phenotyping study of *Hoxd11* knock-out mice, the authors employ univariate hypothesis testing with regression analysis on PCA bases of point-based models of the mouse forepaw [14].

3 Challenges for HDLSS Shape Statistics

This section describes some problems and solutions in HDLSS shape statistics as they relate to shape analysis. The concepts presented here are applicable to linear statistics using any high-dimensional model, though we will use surface-point correspondence models as examples, computed by the particle-based optimization described in [2]. Point-based models represent shape by sampling each shape surface in a consistently ordered fashion so as to define homologous object surface points called *correspondences*. The set of 3D positions of all m correspondences on a shape is $3m$ shape vector, and the positions of the individual shapes in this $3m$ -dimensional shape space give rise to the statistical analysis. Hypothesis testing is done on a single shape model constructed from all data without knowledge of the group classification, which we refer to as a *combined* model.

In the context of point-based models, or surface samplings obtained from parameterized models, one approach to shape statistics is point-wise analysis of correspondences, which are elements of R^3 . These data are drawn from the *marginal distributions* of the full shape space, and the mean shape is obtained by computing the Euclidean averages of correspondence positions, with point-wise differences defining local shape variation [11]. Hypothesis tests in this case reveal *regions* of significant differences between groups, which can be directly visualized




Fig. 1. Point-wise hypothesis test results for the putamen

as p -value maps on the mean shapes. Styner [11], for example, proposes statistical analysis of the correspondences that uses a nonparametric permutation test with the Hotelling T^2 metric with an FDR correction for the multiple-comparison problem inherent in the analysis.

Figure 1 is an illustration of point-wise hypothesis testing as proposed in [11] on a combined model of the right putamen from normal control subjects and autism patients. The data is taken from an ongoing longitudinal pediatric autism study[15]. For the test, we had 10 autism structures available with 15 matched normals, and used 1024 correspondence points and 20,000 test permutations. The uncorrected p -values that indicate significance at the 5% level are colored in red on the mean normal putamen surface in Fig 1a, and suggest several distinct areas of shape differences. Fig. 1b shows that in this case, however, which is not uncommon in neurological shape analysis, no significant p -values remain after FDR correction (5% bound). This example illustrates a major difficulty encountered in point-wise analysis: the large number of comparisons results in a very conservative correction of the hypothesis test results, significantly reducing the statistical power of the test.

To avoid the multiple-comparisons problem, we can analyze high-dimensional shape model data in the full shape space, i.e. the joint-space of the correspondences. The analysis in this case, however, is also difficult because traditional statistical metrics no longer apply [9]. At issue is the fact that the convergence of any estimator in very high dimensional space is prohibitively slow with the respect to the number of samples. A common solution is to employ dimensionality reduction by choosing a subspace of the $3m$ -dimensional shape in which to project the data for traditional multivariate analysis. Principal component analysis (PCA) is often an attractive choice because the basis vectors are orthogonal and determined solely from the data. With PCA, we can find no more than $n - 1$ modes that have non-zero variance, meaning that the problem is reduced to $d < n$ without loss of information. Other basis functions such as wavelets [12] have also been used for dimensionality reduction, with the difference being that they impose an a-priori choice of how the space of the model should be decomposed.




Fig. 2. Hotelling T^2 test results with increasing numbers of PCA modes for 3 brain structures from a pediatric autism study.

In a suitably low-dimensional shape space, such as basis vectors from PCA, we can apply traditional statistical methods such as nonparametric Hotelling T^2 testing. There are two challenges for dimensionality reduction in shape analysis, however. First, is the factor analysis problem of *how many* basis vectors to choose, which can be hard to resolve when the choice of different numbers of factors leads to different statistical results. The second challenge is how to visualize group differences, which is important for researchers in order to relate the outcomes of statistical tests to scientific hypotheses regarding the growth and form of anatomy. The remainder of this section addresses the problem of choosing the number of bases, and then proposes two visualization strategies for understanding the group differences.

To illustrate the number-of-bases problem, Figure 2 shows the p -value results of Hotelling T^2 permutation tests using increasing numbers of PCA modes on three brain structures from the pediatric autism study referenced above. Several trends can be observed that pose a challenge for the analysis. First, is the trend at higher numbers of modes towards increasing p -values, which is due to the cumulative effects of noise in these lower-variance modes of the PCA. The second trend is that the p -value curves do not smoothly decrease to a minimum value, but rather, tend to bounce around with lower numbers of higher-variance modes. The challenge is to choose as many modes as possible that contain meaningful variation, i.e. variation that is distinguishable from noise, with the caveat that too few modes may result in the loss of information that is useful for group discrimination.

Many methodologies have been proposed to address the number-of-bases problem, and good reviews such as [16], are available. Methodologies range from simple conventions, such as choosing only PCA modes that account for at least 5% of the total variance, to more systematic approaches that attempt to model the problem based on its dimensionality and sample sizes. Notable among these latter approaches, *parallel analysis* is commonly recommended as the best method for determining modes with variance that is distinguishable from noise[17], and is described in more detail in Section 4.

In contrast to the point-wise statistical method illustrated in Figure 1, a significant drawback of hypothesis testing in a PCA subspace is that the group differences in this space are not necessarily easy to visualize and interpret from an anatomical perspective. The hypothesis test poses the question of whether there are significant group differences. The next logical question of interest to researchers is *what* are the group differences. One possible approach to gain insight into this question is to transform the group differences measured in the PCA space back into the full shape space, where they can be visualized on the mean shape surfaces. Implicit in the Hotelling T^2 metric, for example, is a linear discriminant which indicates the direction in the PCA space along which the maximum group difference is observed. This discriminant vector can be rotated back into the full shape space for visualization. Another standard approach for visualizing group differences is a comparison of the differences in the mean shapes. Thin-plate spline analysis is commonly used in morphometric studies, for example, to visualize the deformations between shapes parameterized with landmarks (sparse correspondences) [18], and a similar approach can be applied in the context of high-dimensional point-based shape models. We discuss these visualization strategies further in the next section, along with the development of the dimensionality reduction and hypothesis testing.

4 Methodology

For a correspondence point shape model in 3D, we have a $3m \times n$ shape matrix \mathbf{P} , where columns of \mathbf{P} are the shape vectors of correspondence points for the set of all samples. For dimensionality reduction, we first project \mathbf{P} onto the basis vectors determined by PCA analysis, i.e. $\tilde{\mathbf{P}} = \mathbf{E}\mathbf{P}$, where columns of \mathbf{E} are the eigenvectors of the covariance matrix of correspondences, in decreasing order of the magnitude of their eigenvalues.

Following projection into the PCA space, we perform parallel analysis to choose the number of PCA bases for hypothesis testing. In the context of principal components analysis (PCA) of n , vector-valued data samples of dimensionality $3m$, the goal of parallel analysis is to identify the subset of the components that contain variation distinguishable from the expected variation resulting from noise, where noise is modeled by an isotropic, multivariate unit Gaussian, i.e. a random $3m$ -vector $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. To make such a distinction, we need an estimator \mathbf{E} for the expected values of the variances in the *ordered* PCA modes of random samplings from X , given the fixed sample size n . Due to the ordering problem, there is no obvious closed-form expression for E , so it is estimated using Monte Carlo simulation. Many random sample sets of size n are independently drawn from X , followed by PCA on each sample set and ordering of the associated eigenvalues. The ordered eigenvalues are then averaged to produce an estimate of the Gaussian noise variance profile across modes. Note that the eigenvalues in this case quantify variance, and the percentage of total variance for a PCA mode is equivalent to the ratio of its eigenvalue to the sum of all eigenvalues.

In order to determine the number of modes to use from parallel analysis, the percent-total-variance profiles from the Monte Carlo simulation and the PCA of the true data are compared, and only modes where the percent-total-variance in the true data is greater than the simulation data are retained. Figure 3, for example, is a scree plot of the percent-variances associated with shape data of a putamen brain structure ($n = 25, 3m = 3000$) [19] (solid line) and the variances from the Monte Carlo noise variance simulation (dashed line). The two lines intersect just before mode 6, and so we would consider only modes 1-5 in the analysis.




Fig. 3. Parallel analysis for the putamen data

Once we have chosen the set of k PCA basis vectors by parallel analysis, we project the correspondences into that subspace. Hypothesis testing for group differences can now be done using a nonparametric, permutation test with the Hotelling T^2 metric, with the null hypothesis that the two groups are drawn from the same distribution.

The Hotelling T^2 , two-sample metric is given by

$$T^2 = \frac{(n_a n_b)(n_a + n_b - 2)}{n_a + n_b} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}})^T \mathbf{w}, \quad (1)$$

$$\mathbf{w} = (\Sigma_a + \Sigma_b)^{-1} (\mu_{\mathbf{a}} - \mu_{\mathbf{b}})$$

where $\mu_{\mathbf{a}}$ and $\mu_{\mathbf{b}}$ are the means, Σ_a and Σ_b are the covariances, and n_a and n_b are the sample sizes of the two groups, respectively. Note that the vector \mathbf{w} is also Fisher's linear discriminant, which is well known to be the line along which the between-group variance is maximized with respect to within-group variance[20]. The Hotelling T^2 metric is therefore a scaled projection of the group difference onto the discriminant line. We therefore propose to visualize the morphological differences that are driving the statistical test results by transforming \mathbf{w} back from PCA space into the full-dimensional shape space, i.e. $\hat{\mathbf{w}} = \mathbf{E}^{-1} \mathbf{w}$, where $\hat{\mathbf{w}}$




Fig. 4. LDA visualization for the *HoxD11* phenotype study

is \mathbf{w} padded to n -dimensions with $n - k$ zeros. The components of the $3m$ vector $\hat{\mathbf{w}}$ can then be mapped onto a visualization of the n mean correspondence point positions, such as a surface mesh. This resulting linear discriminant analysis (LDA) visualization indicates group differences in shape that the test metric identified as the most effective for discrimination.

To visualize deformations between the group mean shapes, we can compute metrics on the displacement field describing the mapping from points \mathbf{x} on one group mean to corresponding points \mathbf{x}' on the another. Using the set of correspondence points, a smooth transformation $T(\mathbf{x}) = \mathbf{x}'$, can be computed using a thin-plate spline interpolation. Details for computing $T(\mathbf{x})$ are omitted here for brevity, and good descriptions can be found elsewhere (e.g. [18, 21]). We propose to visualize strain, a measure on the Jacobian J of the deformation field $\mathbf{x} - T(\mathbf{x})$ that describes the local stretching and compression caused by the deformation. The Lagrangian strain tensor is a symmetric, second order tensor given by

$$E = \frac{1}{2}(J + J^T + J^T J). \quad (2)$$

The eigenvectors of E indicate the principal directions of strain, and the eigenvalues of E indicate the unit elongations in those directions. An effective visualization for the strain tensor is an ellipsoid with principal axes given by the eigenvalues and oriented along the eigenvector directions.

5 Results and Discussion

This section presents two shape analysis experiments that illustrate the effectiveness of the proposed hypothesis testing and visualization methodology for phenotyping studies of gene function, and for the analysis of clinical neurological datasets. As an example of the application of our method to phenotyping, we analyzed differences between wild-type mice and a population of mice deficient in the gene *Hoxd11*. The data for this study are segmentations of the second phalange of the first digit of the right forepaw, derived from micro-CT images of normal and mutant strains acquired by the authors [14]. We computed a combined shape model of the normal ($n=20$) and mutant population ($n=20$) using the particle system method from [2], and applied the hypothesis testing framework from Section 4.




Fig. 5. LDA visualization of the right putamen from an autism study

For the mouse data, we have a clear biological hypothesis as to the group differences between mutant and normal mouse strains. Through a series of gene-targeting experiments, Boulet, Davis and Capecchi have shown that *Hoxd11* is important for the normal development and patterning of the appendicular skeleton, and suggest group differences between wild-type and mutant mice in the length of several of the digits of the forepaw [1, 22]. The proposed hypothesis test method applied to the phalange model indicates highly significant group differences ($p \ll .01$), with the parallel analysis choosing the first two PCA modes for the hypothesis test. Figure 4 depicts the length in the surface normal direction of each of the point-wise vector components of the discriminant $\hat{\mathbf{w}}$ on a surface reconstruction of the mean wild-type population correspondence points. In the figure, the direction of the arrows are from the wild-type to the mutant population. The linear discriminant visualization reveals two clear morphological differences: a reduction in bone length and an increase in bone thickness in the *Hoxd11*-deficient population. This analysis has, therefore, quantified and statistically validated one of the major conclusions from empirical studies of the *Hoxd11* gene, as well as revealing a new significant phenotypic effect.

As a second example, we present the analysis of the three brain structure model described from the pediatric autism study described in Sect.3. The p -value results for the amygdala, putamen, and hippocampus models, respectively, are 0.003, 0.046, and 0.100, with the number of PCA modes chosen as 5, 6, and 5. Of particular interest is the fact that the result for the putamen indicates group differences at the 5% significant level, which is in contrast to the point-wise hypothesis testing shown in Fig. 1 that indicates no significance. This difference illustrates the increased statistical power of the proposed testing method, which avoids the multiple-comparisons problem. The discriminant vector is visualized for the putamen in Fig.5 for the mean normal population correspondence points, with arrows indicating the direction from patient to the normal control populations. The visualization indicates a shortening of the anterior and posterior regions of the putamen, with a thickening in the medial region.

Figure 6 is a visualization of the strain tensors computed from the deformation from the mean patient shape to the mean normal control shape for




Fig. 6. Strain tensors for the right putamen from an autism study. Tensor scale is exaggerated for visualization.

the putamen data. The strain tensors were computed from a thin-plate spline deformation, as described in the previous section. In the figure, the three principal axes of each ellipsoid are scaled by the three principal eigenvalues of the strain tensor at each correspondence position, and oriented according to their corresponding eigenvectors. Ellipsoids and the surrounding surface are colored according to the value of the first principal eigenvector (the longest axis), with yellow indicating negative (compression) and blue indicating positive (stretching). While a clinical interpretation of this result is beyond the scope of this paper, this visualization clearly offers a more detailed insight into *how* groups differ than, for example, the LDA visualization. Note, however, that in this case we have no indication of the statistical significance of these differences.

In summary, the proposed hypothesis testing and visualization methodologies offer an intuitive approach to analysis of anatomical shape from images for biological and clinical research, and avoid the problems inherent in HDLSS statistics. We have illustrated the effectiveness of the framework by statistically validating a biological hypothesis regarding *Hoxd11* gene function that was previously only based on empirical evidence, and we have shown how the method can be useful in exploring new patterns in shape from clinical data which have not previously been observed, or that are not observable with lower-powered statistical methods.

Acknowledgments

This work was funded by the Center for Integrative Biomedical Computing, National Institutes of Health (NIH) NCRR Project 2-P41-RR12553-07. This work is also 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.

References

1. Davis, A., Capecchi, M.: Axial homeosis and appendicular skeleton defects in mice with targeted disruption of *hoxd-11*. *Development* **120** (1995) 2187–2198
2. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: *Information Processing in Medical Imaging (IPMI 2007)*, LNCS 4584. (2007) 333–345
3. Gorczowski, K., Styner, M., Jeong, J., Marron, J., Piven, J., Hazlett, H., Pizer, S., Gerig, G.: Statistical shape analysis of multi-object complexes. In: *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, IEEE (2007)* 1–8
4. Styner, M., Lieberman, J.A., Pantazis, D., Gerig, G.: Boundary and medial shape analysis of the hippocampus in schizophrenia. *Medical Image Analysis* (2004)
5. Pizer, S.M., Jeong, J.Y., Lu, C., Muller, K.E., Joshi, S.C.: Estimating the statistics of multi-object anatomic geometry using inter-object relationships. In: *Deep Structure, Singularities, and Computer Vision*. Volume 3753 of LNCS. (2005) 60–71
6. Klingenberg, C.P.: Morphometrics and the role of the phenotype in studies of the evolution of developmental mechanisms. *Gene* **287** (2002) 3–10
7. Chen, X.J., Kovacevic, N., Lobaugh, N.J., Sled, J.G., Henkelman, R.M., Henderson, J.T.: Neuroanatomical differences between mouse strains as shown by high-resolution 3d mri. *NeuroImage* **29** (2005) 99–105
8. Davies, R.H., Twining, C.J., Allen, P.D., Cootes, T.F., Taylor, C.J.: Shape discrimination in the hippocampus using an mdl model. In: *IPMI*. (2003) 38–50
9. Ahn, J., Marron, J.S., Muller, K.M., Chi, Y.: The high-dimension, low-sample-size geometric representation holds under mild conditions. *Biometrika* **94**(3) (2007) 760–766
10. Terriberry, T., Joshi, S., Gerig, G.: Hypothesis testing with nonlinear shape models. In: *IPMI'05*. (2005) 15–26
11. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J., Shenton, M., Gerig, G.: Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *The Insight Journal* (2006)
12. Nain, D., Niethammer, M., Levitt, J., Shenton, M., Gerig, G., Bobick, A., Tannenbaum, A.: Statistical shape analysis of brain structures using spherical wavelets. In: *IEEE Symposium on Biomedical Imaging ISBI*. (2007) in print
13. Golland, P., Grimson, W., Shenton, M., Kikinis, R.: Detection and analysis of statistical differences in anatomical shape. *Medical Image Analysis* **9** (2005) 69–86
14. Cates, J., Fletcher, P.T., Warnock, Z., Whitaker, R.: A shape analysis framework for small animal phenotyping with application to mice with a targeted disruption of *hoxd11*. In: *Proc. 5th IEEE International Symposium on Biomedical Imaging (ISBI '08)*. (2008) 512–516
15. Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., Gilmore, J., Piven, J.: Magnetic resonance imaging and head circumference study of brain size in autism: Birth through age 2 years. *Arch Gen Psych* **62** (2005) 1366–1376
16. Fabrigar, L.R., Wegener, D.T., MacCallum, R.C., Strahan, E.J.: Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods* **4** (1999) 272–299
17. Glorfeld, L.W.: An improvement on horn's parallel analysis methodology for selecting the correct number of factors to retain. *Educational and Psychological Measurement* **55** (1995) 377–393
18. Bookstein, F.: Principal warps: Thin plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **11**(6) (1989)

19. Cates, J., Fletcher, P., Styner, M., Hazlett, H., Whitaker, R.: Particle-based shape analysis of multi-object complexes. In: Proceedings of the 11th International Conference on Medical Image Computing and Computer Assisted Intervention, MIC-CAI (2008) to appear
20. Timm, N.H.: Applied Multivariate Analysis. Springer-Verlag (2002)
21. Whitbeck, M., Guo, H.: Multiple landmark warping using thin-plate splines. In: IPCV. (2006) 256–263
22. Boulet, A.M., Capecchi, M.R.: Duplication of the *hoxd11* gene causes alterations in the axial and appendicular skeleton of the mouse. *Developmental Biology* **249** (2002) 96–107

Author Index

- Albrecht, Th., 160
Allasonnière, S., 80
Ayache, N., 68
- Barysheva, M., 48
Bigot, J., 92
Bossa, M., 13
Brun, C., 48
- Cahill, N.D., 117
Cates, J., 170
Chan, T., 36, 56
Chou, Y.Y., 48
Chung, M.K., 128
- Davidson, R.J., 128
Durrleman, S., 68
- Fletcher, P.T., 170
- Gorelick, L., 117
Gu, X., 36
Gutman, B., 56
- Hernandez, M., 24
- Kim, P.T., 140
Knothe, R., 160
Kubicki, M., 150
Kuhn, E., 80
- Le Guyader, C., 1
Lee, A., 48
Leow, A., 1
Leporé, N., 48
Liu, X., 105
Loubes, J.-M., 92
- McMahon, K.L., 48
Melonakos, J., 150
- Meredith, M., 48
Mio, W., 105
Mohan, V., 150
- Nacewicz, B.M., 128
Niethammer, M., 150
Noble, J.A., 117
- Olmos, S., 13, 24
- Pennek, X., 24, 68
Pollak, S., 128
- Qiu, A., 128
- Richards, D.St.P., 140
- Shi, Y., 105
Sundaramoorthi, G., 150
- Tannenbaum, A., 150
Thompson, P.M., 1, 36, 48, 56
Toga, A.W., 1, 48, 56
Trouvé, A., 68, 80
- Vese, L., 1
Vesom, G., 117
Vetter, Th., 160
Vimond, M., 92
- Wang, Y., 36, 56
Whitaker, R., 170
Wright, M.J., 48
- Yanovsky, I., 1
Yau, S.T., 36
Yin, X., 36
- Zhang, J., 36
Zubiaray (de), G.I., 48