Statistical Models of Sets of Curves and Surfaces based on Currents

Stanley Durrleman [a,b,*] Xavier Pennec [a] Alain Trouvé [b] Nicholas Ayache [a]

[a] *Asclepios team project, INRIA Sophia Antipolis Méditerranée, 2004 route des Lucioles, 06902 Sophia Antipolis cedex, France*
[b] *Centre de Mathématique et Leurs Applications, ENS Cachan, 61 avenue du président Wilson, 94235 Cachan cedex, France*

Abstract

Computing, visualizing and interpreting statistics on shapes like curves or surfaces is a real challenge with many applications ranging from medical image analysis to computer graphics. Modeling such geometrical primitives with currents avoids to base the comparison between primitives either on a selection of geometrical measures (like length, area or curvature) or on the assumption of point-correspondence. This framework has been used relevantly to register brain surfaces or to measure geometrical invariants. However, while the state-of-the-art methods efficiently perform pairwise registrations, new numerical schemes are required to process groupwise statistics due to an increasing complexity when the size of the database is growing.

In this paper, we propose a Matching Pursuit Algorithm for currents, which allows us to approximate, at any desired accuracy, the mean and modes of a population of geometrical primitives modeled as currents. This leads to a sparse representation of the currents, which offers a way to visualize, and hence to interpret, such statistics. More importantly, this tool allows us to build atlases from a population of currents, based on a rigorous statistical model. In this model, data are seen as deformations of an unknown template perturbed by random currents. A Maximum A Posteriori approach is used to estimate consistently the template, the deformations of this template to each data and the residual perturbations. Statistics on both the deformations and the residual currents provide a complete description of the geometrical variability of the structures.

Eventually, this framework is generic and can be applied to a large range of anatomical data. We show the relevance of our approach by describing the variability of population of sulcal lines, surfaces of internal structures of the brain and white matter fiber bundles. Complementary experiments on simulated data show the potential of the method to give anatomical characterization of pathologies in the context of supervised learning.

Key words: currents, curves, surfaces, statistics, matching pursuit algorithm, approximation, sparse decomposition, atlas estimation, template, registration, shape space, anatomical variability, group classification, computational anatomy

1. Introduction

In many communities such as medical imaging, computer vision or computer graphics, there is a substantial need for shape statistics to incorporate shape priors in image segmentation, to analyze geometrical or anatomical differences between groups, to classify new observations according to some characteristics, for shape recognition purpose, etc. Over the last years, many methods have been proposed for encoding statistical priors on curves, quite less were proposed for surfaces, even less were proposed for generic geometrical data like point sets, curves, surfaces or volumes. In this paper, for the sake of simplicity, we call simply shapes any of these geometrical data.

Approaches based on level-sets use geometrical constraints (length, area, volume, curvature, etc.)

* Corresponding author.
Email address: stanley.durrleman@sophia.inria.fr (Stanley Durrleman).

as priors, like in Leventon et al. (2000, 2003); Pardo et al. (2004); Cremers (2006), for instance. However, it is difficult to figure out how to automatically learn such priors from typical data sets. Indeed, curves in this setting are embedded into a space which does not allow us to define easily statistics. Models based on point distribution assume point-correspondences between structures (Cooper et al., 1995; Cootes et al., 2008). Such approaches can lead to statistical models as in Twining et al. (2005) for instance. In medical imaging, however, samples of curves or surfaces segmented from MR images, have never been shown to correspond from one structure to another. Assuming arbitrary point correspondences may lead to important bias in the statistical estimations. Improvements of point distribution models were proposed to relax the constraint of point correspondences, for instance Minimum Description Length approach (Marsland et al., 2008), statistical versions of the Iterative Closest Points algorithm (Hufnagel et al., 2008) or fuzzy correspondences approaches (Chui and Rangarajan, 2003). These techniques, though, are still built on the idea of point correspondences. Alternatively, Medial Axis representations have been proposed to build statistical models on surfaces (Pizer et al., 2003; Fletcher et al., 2004). They require, nevertheless, to completely specify the topology of shapes. Like in several of the previous references, Charpiat et al. (2005); Mio et al. (2007); Joshi et al. (2007) propose to do statistics on the deformations between shapes in the perspective of Grenander's approach based on group action for modeling objects (Grenander, 1994). Although such approaches have been proved useful to capture variations between anatomical structures, they lead only to a partial description of the variability. Indeed, interesting anatomical features may remain in the residual shape that has not been captured by the smooth deformations (difference between the registered shapes and the target shape). Such features are typically change of topology, matter creation of deletion. A complete description of the variability should be based therefore on both the deformations and the residual shapes. This requires to embed shapes into a metric space which enables to compute easily statistics and which may be coupled consistently with a deformation framework (an action of deformations on the metric shape should be defined).

In contrast to the previous modeling, the currents define a unifying framework to process any kind of geometrical data like sets of curves, sets of surfaces or a mix of both. This framework has the advantage not to assume point correspondences between structures, to be robust to change of connectivity within the sets of shapes and to be weakly sensitive to the sampling of shapes (Vaillant and Glaunès, 2005; Durrleman et al., 2008b,c). In this framework, the comparison between shapes relies on their global geometrical similarity and not on a set of a priori extracted features like volume, length or curvature for instance. Moreover, shapes are embedded into a vector space provided with an inner-product. This gives a mathematically well grounded framework to compute statistics like mean and principal modes of a population of shapes.

This framework has been successfully applied to perform registrations between two sets of anatomical structures segmented from MR images of two different subjects (Vaillant and Glaunès, 2005; Glaunès, 2005; Vaillant et al., 2007). If a template configuration is given, statistics may be performed on the registrations between this template and every subject within a population as in Vaillant et al. (2004); Durrleman et al. (2008c). However, the estimation of the template itself, consistently with its deformation within the population, has not yet been given a satisfactory answer. In Glaunès and Joshi (2006), a template is estimated from a population of currents, but in a way that does not allow us to perform easily joint statistics on deformations and residual currents. Such residues, however, may contain also interesting anatomical features. The purpose of this paper is precisely to present a unifying framework to (1) estimate a prototype shape from a population of shapes along with its deformation to every subject, (2) perform statistics on the deformations and (3) compute mean and modes of the residual currents. Generic statistics on deformations via tangent-space representation have already been proposed by Vaillant et al. (2004); Pennec (2006) for instance. We have now to focus on the statistics on the space of currents and the atlas estimation.

In the framework of currents, sets of curves or surfaces are represented by a set of unconnected oriented points. Each oriented point represents the segment of a polygonal line or the normal of a mesh cell. While this representation has been proved useful to perform pairwise registrations, it raises dramatic computational issues for group-wise statistics. The representation of the mean or principal modes of a population of shapes contain as many elements as the total number of tangents or normals in the database. The size of this representation tends to in-

finity as the number of subjects within the database increases, or as the sampling of the discrete shapes becomes finer. In both cases, however, the statistics (mean and principal modes) converge in the space of currents. This representation is therefore not adapted to perform statistics on currents.

In this paper, we propose to use Matching Pursuit algorithms to give a more stable representation of currents (i.e. whose size does not increase linearly with the number of subjects). This method builds on ideas from the approximation theory previously developed to decompose images in wavelet bases (Mallat and Zhang, 1993; Pati et al., 1993; Davis et al., 1997). To the very best of our knowledge, this is the first time that these signal processing techniques are used in geometric shape analysis. Given any current (like mean of a population of shapes for instance), we find a new basis on which the coefficients of the decomposition of the current decrease quickly. The basis is chosen so that the redundancy of the decomposition is the lowest possible. In this sense, we say that this decomposition is adapted to the signal. Keeping only the first terms of this decomposition provides therefore an approximation of the input current at any desired accuracy. Our experiments on real data show that we can achieve very high compression ratio, with a very low approximation error. The deformation of a mean obtained from 3 shapes of 15 000 points such as in Fig.9, which was previously taking 10 hours, is now taking about 5 minutes (using the same code as in Vaillant and Glaunès (2005); Durrleman et al. (2007)). For a set of 50 instances, representing the mean requires 1.2 Kb in our framework, versus 8 Mb originally. Deforming the former still requires 5 minutes while it is not feasible to deform the later without high performance computing.

This new representation of currents can be used now to define statistics on currents. In a first simple model, we consider that the population of shapes is a realization of a random Gaussian process. Empirical mean and empirical covariance matrix gives estimation of the mean and covariance of this Gaussian law. Our numerical tool gives a sparse representation of the empirical mean and principal modes, which are easier to interpret than their initial representation (i.e. the set of every segment and normal in the database). However, deformations have been proved useful to measure anatomical variability and one wants to include them consistently within a global statistical framework to analyze variability of a population of shapes. We define therefore a second statistical model, in which every shape is considered as a random smooth deformation of an unknown template plus a random Gaussian perturbation in the space of currents. In this setting, the Maximum A Posteriori estimation of both the template and the deformations needs the use of our new numerical tool. Once the template and deformations are estimated, joint statistics may be performed on the deformations and the residual perturbation (difference in the space of current between the input shapes and the deformed template). This model decomposes therefore the anatomical variability into a geometrical part captured by the deformations, and a "texture" part which contains everything that cannot be captured by a smooth deformation.

These statistical models are generative. This means that they not only provide descriptive measures of variability (like variance, correlation between observations and modes, hypotheses testing, etc.), but also allows us to generate new data according to the estimated variability, or to compare any new available data with the variability captured by our model. By this mean, we can visually interpret the variability captured by our model and therefore drive the search of new anatomical knowledge. This also offers a way to detect automatically pathologies from deviations to the normal variability, or to the automatic classification of populations.

This framework is also very versatile and generic. It can be applied to a large range of anatomical data such as sulcal curves, surfaces of internal structures of the brain or white matter fiber bundles. We use our statistical analysis to describe the anatomical variability of the Sylvian fissure in a normal population, to describe the anatomical differences between hippocampus of population of autistics and a population of controls, to provide a description of the variability of the cortico-bulbar tract both in term of geometrical variations and residual non-diffeomorphic perturbations. We also illustrate the potential of our method for automatic group classification in the context of supervised learning.

The paper is organized as follows. In Section 2, we introduce the currents and recall the main properties used in this paper. In Section 3, we adapt the Matching Pursuit algorithm to our framework based on currents. Synthetic examples show how it can be used to represent mean and modes of a population of shapes. In Section 4, we introduce the complete statistical model which includes deformations. We show how the Matching Pursuit algorithm helps to estimate consistently the template and the deforma-

tion of this template to every subject in the population. Synthetic examples show how this model works on 2D-curves, and how it can be used for the separation of populations. In Section 5, we gather the experiments on real anatomical data. We show that our two statistical models lead to relevant description of the anatomical variability of curves or surfaces.

2. Non-parametric representation of shapes as currents

2.1. *Currents: an object which integrates vector fields*

A current is a mathematical object which has been proved relevant to model geometrical data like curves and surfaces (Vaillant and Glaunès, 2005; Glaunès and Joshi, 2006; Durrleman et al., 2007). The idea is to characterize shapes via vector fields, which are used to probe them (shapes is here a generic word to denote curves, surfaces or sets of such objects). Given ω a square integrable 3D vector field, a surface (or a set of surfaces) S integrates ω thanks to the flux equation:

$$S(\omega) = \int_S \omega(x)^t (u \times v)(x) d\sigma(x) \qquad (1)$$

where $(u \times v)(x)$ is the normal of the surface at point x, (u, v) an orthonormal basis of its tangent plane at x, and $d\sigma$ the Lebesgue measure on the surface. Equation 1 computes the flux of the vector field ω through the surface S.

Similarly, a curve (or a set of curves) L integrates any vector fields ω thanks to the path-integral:

$$L(\omega) = \int_L \omega(x)^t \tau(x) dx \qquad (2)$$

where τ is the tangent of the curve at point x. Equation 2 computes the flux of the field of tangents through all equipotential surfaces of ω.

To characterize a shape (set of curves or surfaces), we measure how these quantities (Eq.1 or Eq. 2) vary while ω (or equivalently the equipotentials of ω) varies. For this purpose, we need to define a test space of square integrable vector fields W, in which ω varies. We will then define the space of currents W^* as the space of continuous linear mapping from W to $\mathbb{R}$ (i.e. objects which integrate vector field). Eq.1 and Eq. 2 will make therefore any sets of curves and surfaces particular cases of currents.

We choose W as the set of the convolutions between any square integrable vector fields and a smoothing kernel. This excludes from W the vector fields with too high spatial frequencies. Formally, W is the reproducing kernel Hilbert space (r.k.h.s [1]). The kernel determines the transfer function of the filter. In our applications, we will use a Gaussian kernel: $K^W(x, y) = \exp(-\|x - y\|^2 \lambda_W^2)\mathrm{Id}$ for any points (x, y). In this case, the standard deviation λ_W is the typical scale at which the vector fields ω in W may vary spatially. By construction, the space of vector fields W has two important properties: (1) W is the dense span of the vector fields of the form $\omega(x) = K^W(x, y)\beta$ for any fixed points y and vectors β; (2) W is provided with an inner product which is defined on these basis vectors by $\left\langle K^W(., x)\alpha, K^W(., y)\beta \right\rangle_W = \alpha^t K^W(x, y)\beta$. This inner product leads to the more general reproducing property:

$$\left\langle K^W(., x)\alpha, \omega \right\rangle_W = \alpha^t \omega(x) \qquad (3)$$

The space of currents, denoted W^*, is the space of the continuous linear mappings from W to $\mathbb{R}$. As a consequence, the Riesz-Frechet theorem ensures that there is a linear mapping between the space of vector fields W and its dual space W^*, the space of currents. We denote this mapping $\mathcal{L}_W : W \longrightarrow W^*$. It is defined by:

$$\mathcal{L}_W(\omega)(\omega') = \langle \omega, \omega' \rangle_W \qquad (4)$$

for all vector fields $\omega, \omega' \in W$ (for $\omega \in W$, $\mathcal{L}_W(\omega)$ is a current, i.e. a mapping from W to $\mathbb{R}$). We call $\mathcal{L}_W(\omega)$ the dual representation of the vector field ω.

The dual representation of the basis vectors $K^W(x, .)\alpha$ are called the Delta Dirac currents: $\delta_x^\alpha = \mathcal{L}_W(K^W(x, .)\alpha)$. This shows that K^W is the Green function of the differential operator [2] $\mathcal{L}_W$. Combining Eq.3 and Eq.4, we get:

$$\delta_x^\alpha(\omega) = \left\langle K^W(x, .)\alpha, \omega \right\rangle_W = \alpha^t \omega(x) \qquad (5)$$

We have therefore : $\delta_x^\alpha(\omega) = \alpha^t \omega(x)$, which is the term within integrals in Eq. 1 and Eq. 2. A delta Dirac current may be interpreted therefore as an infinitesimal segment (or normal) α entirely concentrated at point x. Since W is a dense span of the vector fields $K^W(x, .)\alpha$, the space of currents is a dense

[1] We refer the reader to Saitoh (1988) for more theoretical details on the construction of r.k.h.s.

[2] An equivalent construction would consist of fixing a differential operator $\mathcal{L}_W$ and to denote K^W its Green function. However, we prefer here to have a closed form for the kernel instead of the differential operator.

span of the Delta Dirac currents δ_x^α. This means that any currents may be decomposed into a *infinite* set of Dirac currents. If the current T is a set of curves or surfaces, it decomposed into the infinite set of its tangents or normals at each point of the shape. Its dual representation $\mathcal{L}_W^{-1}(T)$ (a vector field in W) is the convolution of every tangent (or normal) with the regularizing kernel K^W.

Although a curve (resp. a surface) has an *infinite* number of tangents (resp. normals), a polygonal lines (resp. a mesh) may be approximated by a *finite* sum $\sum_k \delta_{x_k}^{\alpha_k}$ where x_k is the center of each segment (resp. center of each mesh cell) and α_k the tangent of the line (resp. the normal of the surface) at x_k. Its dual representation is therefore given at any point x by the finite sum: $\sum_k K^W(x, x_k)\alpha_k$. With this approximation, the integrals in Eq. 1 and Eq. 2 are replaced by their Riemann sums. This approximation converges in the space of currents as the sampling of the curves and surfaces become finer. This shows that this modeling of curves and surfaces is weakly sensitive to the sampling of the geometrical objects. Moreover, the description in term of the collection of tangents and normals account only for local properties of the shapes. It makes the framework based on currents fully robust to topology changes or the change of connectivity between structures (like curves interruption or reconnection for instance).

2.2. *Correspondence-less distance between curves or surfaces*

One of the main interest of modeling curves or surfaces as currents is that it embeds these geometrical objects into a vector space provided with an inner-product. This allows us to define a distance between curves and surfaces which do not assume any point correspondences between structures.

As mappings from W to $\mathbb{R}$, the currents build a *vector space* (W^*): $(T_1 + T_2)(\omega) = T_1(\omega) + T_2(\omega)$ and $(\lambda.T)(\omega) = \lambda.T(\omega)$. If T_1 and T_2 are 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. The union of the normals of a surface is represented by the addition of their Delta Dirac currents. Scaling a surface means scaling the power of the flux through the surface.

The space of currents W^* is provided with an inner-product, denoted $\langle ., . \rangle_{W^*}$, such that the mapping $\mathcal{L}_W$ between the space of vector fields W and W^* is *isometric*. On the basis vectors, this inner-product is therefore given by:

$$\begin{aligned} \left\langle \delta_x^\alpha, \delta_y^\beta \right\rangle_{W^*} &= \langle K(.,x)\alpha, K(.,y)\beta \rangle_W \\ &= \alpha^t K^W(x,y)\beta \end{aligned} \tag{6}$$

By linearity, the inner product between two finite sets of Dirac currents $T = \sum_i \delta_{x_i}^{\alpha_i}$ and $T' = \sum_j \delta_{y_j}^{\beta_j}$ (which may model two surfaces for instance) is given by: $\langle T, T' \rangle_{W^*} = \sum_i \sum_j \alpha_i^t K^W(x_i, y_j)\beta_j$. This gives *explicit* and *easily tractable* formula to compute the inner product between two shapes. The distance between the two shapes is simply defined by the norm of their difference: $d(T, T') = \|T - T'\|_{W^*}$.

This distance on the space of currents has a geometrical interpretation. Indeed, one can show that the norm on currents is a *spectral* norm: $\|T - T'\|_{W^*}^2 = \sup_{\|\omega\|_W \leq 1} |T(\omega) - T'(\omega)|$. If T is a surface, one looks for the regular vector field ω ($\|\omega\|_W \leq 1$) which best separates the two surfaces, in the sense that the difference of the flux of ω through each surface is the largest possible. In other words, a small distance means that we cannot find any regular vector field which makes a difference between the flux through each surface. The vector field which achieves the supremum is given by the dual representation of $T - T'$: $\omega = \mathcal{L}_W^{-1}(T) - \mathcal{L}_W^{-1}(T')$. In some sense, this vector field captures the largest differences between the two surfaces. It allows us therefore to visualize where these differences occur. The smaller the standard deviation λ_W, the smaller the scale at which ω may vary spatially, the finer the geometrical details captured by this distance. This distance measures shape dissimilarity globally and do not assume point correspondence between surfaces.

2.3. *Gaussian variables in the space of currents*

To define statistical models on the space of currents, we still need to define a *Gaussian noise* on currents. There is a generic way to define random variables in Hilbert spaces, such that their restrictions on any finite dimensional spaces are Gaussian. Let $V = \sum_{k=0}^{\infty} v_k e_k$ be the decomposition of a vector V on the Hilbert basis (e_k) (as such, $\|V\|^2 = \sum_{k=0}^{\infty} |v_k|^2 < \infty$). This vector may be associated to the Gaussian variable $\sum_{k=0}^{\infty} \gamma_k v_k$, where each γ_k is a zero-mean real Gaussian variable with variance 1. The variance of this variable is given by $\|V\|^2$. In infinite dimension, a random Gaussian variable ε is defined therefore as a mapping from every vector of the Hilbert space (here our space of currents) to a

real Gaussian variable. Given the inner-product in the space of currents, the covariance between the Gaussian variables associated to two basis elements is given by $\mathrm{Cov}(\varepsilon(\delta_x^\alpha), \varepsilon(\delta_y^\beta)) = \alpha^t K^W(x,y)\beta$.

In a discrete setting, we usually set a grid of points $\Lambda = \{x_p\}_{p=1\ldots N}$. The span of the Delta Dirac currents constrained to belong to the grid $(\mathrm{Span}\{\delta_{x_p}^\alpha; x_p \in \Lambda, \alpha \in \mathbb{R}^3\})$ defines a finite dimensional space of currents. The restriction of the Gaussian variables on this finite dimensional space are real Gaussian variables of the form $\varepsilon_\Lambda = \sum_{p\in\Lambda} \delta_{x_p}^{\alpha_p}$ where (α_i) are centered Gaussian variables with covariance matrix $\mathcal{K}^{-1}$ where $\mathcal{K} = (K^W(x_i,x_j))_{i,j\in\Lambda}$. In this case, ε_Λ has a probability density function (pdf) which is proportional to $\exp(-\|\varepsilon_\Lambda\|^2_{W^*}/2)$.

2.4. *Diffeomorphic deformations of currents*

To include deformations into the analysis of variability, we need to define a deformation framework and see how deformations transport curves and surfaces modeled as currents.

Let ϕ be a diffeomorphism (a smooth deformation of the underlying 3D space, with smooth inverse) and S a surface. As a surface, S may be deformed by ϕ into $\phi(S)$ (the geometrical transportation of the points of S). If we model S as a current, we define the push-forward current $\phi_* S$ such that the flux of any vector field ω through $\phi_* S$ is equal to the flux of ω through the transported surface $\phi(S)$. A change of variable within integrals of Eq. 1 and Eq. 2 leads to the definition: $\phi_* S(\omega) = S(\phi^*\omega)$ where the pull-back vector field $\phi^*\omega$ is equal to $d_x\phi^t \omega(\phi(x))$ for curves and $|d_x\phi|^{-1}(d_x\phi)\omega(\phi(x))$ for surfaces ($d_x\phi$ denotes the Jacobian matrix of ϕ and $|d_x\phi|$ its determinant). This action replaces for curves and surfaces the usual action on images: $(\phi_* I)(x) = I(\phi^{-1}(x))$. 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).

In practice, the push-forward action on the basis vectors is simply given by:

$$\phi_* \delta_x^{u\times v} = \delta_{\phi(x)}^{d_x\phi(u)\times d_x\phi(v)} \tag{7}$$

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

$$\phi_* \delta_x^{\alpha} = \delta_{\phi(x)}^{d_x\phi(\alpha)} \tag{8}$$

if α is a tangent of a curve.

In the following, we restrict the deformations ϕ to belong to the group of diffeomorphisms set up in Trouvé (1998); Dupuis et al. (1998): the diffeomorphisms are obtained by integration of a time-varying vector field v_t: $\partial\phi_t(x)/\partial t = v_t(\phi_t(x))$. v_t is the speed vector field in the Lagrangian coordinates. The geodesic flows $(\phi_t^{v_t})_{t\in[0,1]}$ are completely determined by the initial speed vector field v_0 which belongs to a r.k.h.s. V. The final diffeomorphism at time $t = 1$ is therefore denoted ϕ^{v_0}. As for the currents, we choose for V a Gaussian kernel with variance λ_V^2. This parameter determines the spatial scale at which points move coherently (Durrleman et al., 2008c). How to find such a diffeomorphism that best matches two sets of currents is explained in detail in Vaillant and Glaunès (2005); Glaunès (2005). However, the methodology presented in this paper may work as well with an other group of diffeomorphisms, like the one presented in Marsland and Twining (2004) for instance.

3. Mean and modes in the space of currents

3.1. *Computational issue with the direct approach*

We assume that we have a collection of N discrete shapes (set of curves or surfaces): $T_1, \ldots T_N$. Our purpose is to infer the mean and covariance of this collection of shapes in the space of currents.

Each shape is modeled as a current and is therefore seen as the union of their tangents or normals: $T_i = \sum_p \delta_{x_p^i}^{\alpha_p^i}$ where x_p^i is the center of each segment (resp. mesh cell) of the curve (resp. surface) T_i, and α_p^i its segments (resp. its normals).

Since the space of currents W^* is a vector space, one may compute directly the empirical mean as $\bar{T} = \sum_i T_i/N = \sum_i \sum_p \delta_{x_p^i}^{\alpha_p^i/N}$. This is simply the union of all the tangents (resp. normals) in the database, scaled by $1/N$.

Since the space of currents is provided with an inner-product, one may compute the N-by-N empirical covariance matrix Γ: $\Gamma_{ij} = \langle T_i - \bar{T}, T_j - \bar{T}\rangle_{W^*}$ (computed in practice with Eq. 6). Let $V^1, \ldots, V^N$ be the eigenvectors of Γ. The n-th principal mode of the population of currents is therefore given by: $\bar{T} + \sum_i V_i^n (T_i - \bar{T})$. Expanding this expression leads to a double sum of the form: $\sum_i \sum_p \delta_{x_p^i}^{w_p^{i,n}\alpha_p^i}$ for some weights $w_p^{i,n}$. These expressions of the empirical mean and modes are exact and can be used directly to give quantitative measures of the variability of the population. For instance, one may

compute the projection of each observation to the first mode and analyze the correlation between modes and observations.

However, the representation of the mean and principal modes raise several problems. Both are a weighted sum of all the input Delta Dirac currents. They have as many terms as the total number of tangents and normals within the database. This number increases as the number of subjects in the database increase, while at the same time the statistics converge to their true value. This number increases also as the sampling of the shapes becomes finer, while the discrete shapes converge in the space of currents to continuous geometrical objects. This representation is therefore not stable from a computational point of view. If one wants to deform the mean to a new subject (like for an atlas to subject registration for instance), this representation may cause dramatic computational issues. Moreover, from a visualization point of view, the representation of the mean and principal modes as the union of scaled tangents and normals is particularly difficult to interpret.

Although this representation of these statistics (mean and principal modes) is exact, it is often far from being optimal. It may be highly redundant at the scale λ_W, especially if the shapes are at a distance of order λ_W one from the others. This is particularly visible in the dual representation of the statistics $\gamma = \mathcal{L}_W^{-1}(\bar{T})$. The vector field γ results from the convolution of every tangent or normal with the Gaussian kernel K^W with standard deviation λ_W. Two tangents or normals closer than λ_W contribute to the same mode of the Gaussian distribution γ. This vector field γ integrates precisely the redundancy of the mean current (or principal modes) at the scale λ_W (See Fig.1-a).

Theoretically, the two representations ($\bar{T}$ in the space of current W^* and γ in the space of vector fields W) are equivalent: one should retrieve the original current T from its dual representation $\gamma = \mathcal{L}_W^{-1}(\bar{T})$. However, while $\mathcal{L}_W^{-1}$ is a regularizing operator (a convolution by a Gaussian kernel), $\mathcal{L}_W$ is an ill-posed deconvolution problem. In general, a direct approach to compute the current T from the vector field γ is not possible and requires specific numerical scheme. The idea is here to write the true solution of the deconvolution problem ($T = \mathcal{L}_W(\gamma)$) as an infinite sum ($T = \sum_k \delta_{c_k}^{\beta_k}$), but with faster decreasing term than the initial representation (with a high but finite number of terms). The first terms of this series provide an approximation of the true solution with an increasing precision while the number of terms increases. We expect that a good approximation can be achieved with much fewer terms than the number of tangents or normals in the initial representation.

In the next section, we will present a numerical scheme to perform robustly this ill-posed deconvolution problem. Experiments on simulated examples (see Fig. 1 and Fig. 2), as well as on real anatomical data will prove the relevance of our approach.

3.2. *Orthogonal matching pursuit for currents*

Matching Pursuit Algorithms were proposed initially in Mallat and Zhang (1993); Pati et al. (1993) to find adapted wavelets bases for image decomposition. We adapt here the idea of how to find a basis adapted to a particular signal to our framework based on currents. The method proposed here estimates iteratively points (x_k) and vectors (α_k) such that the series $\sum_{k'} \delta_{x_{k'}}^{\alpha_{k'}}$ approximates the solution of the deconvolution problem: $T = \mathcal{L}_W(\gamma)$ when one knows γ and looks for T. The points and vectors are chosen so that the decomposition of T has fast decreasing terms. In this sense, the basis vectors on which T is decomposed is adapted to the signal. Each basis vector integrates locally the redundancy of the signal at the scale λ_W.

3.2.1. *The continuous case*

Given a vector field $\gamma \in W$, one wants to find a set of N points (x_i) and N vectors (α_i) such that the current $\Pi(T) = \sum_{i=1}^N \delta_{x_i}^{\alpha_i}$ is the closest possible to $T = \mathcal{L}_W(\gamma)$.

If one knows the optimal point positions (x_i), $\Pi(T)$ is the orthogonal projection of T onto the finite vector space $\mathrm{Span}\left(\delta_{x_i}^{\epsilon_k}; k = 1,2,3, i = 1 \dots N\right)$ where $(\epsilon_k)_{k=1,2,3}$ is the canonical basis of $\mathbb{R}^3$. In the following the index k always takes values $k = 1,2,3$. The orthogonality conditions are $\left\langle T, \delta_{x_i}^{\epsilon_k} \right\rangle_{W^*} = \left\langle \Pi(T), \delta_{x_i}^{\epsilon_k} \right\rangle_{W^*}$, which gives (applying the isometric mapping $\mathcal{L}_W^{-1}$ and Eq.3) the set of $3N$ linear equations:

$$\sum_{p=1}^{N} \left(K^W(x_i, x_p)\alpha_p \right)_k = \gamma(x_i)_k \qquad (9)$$

Solving this linear system leads to the optimal vectors α_i, once one knows the optimal positions x_i. Finding the optimal points (x_i) has been proved to be NP-hard in general (Davis et al., 1997). The orthogonal matching pursuit algorithm is a sub-

optimal greedy approach to this problem: the first point x_1 is the one for which the projection of T on $\mathrm{Span}(\delta_x^{\epsilon_k})$ is maximal. δ_{x_1} is the basis vector which enables to explain the largest part of the signal. Since $\langle \mathcal{L}_W(\gamma), \delta_x^{\epsilon_k} \rangle_{W^*} = \langle \gamma, K^W(.,x)\epsilon_k \rangle_W = \gamma(x)_k$, x_1 is a point where $\|\gamma(x)\|_{\mathbb{R}^3}$ is maximal. Solving Eq.9 gives $\alpha_1^1 = K^W(x_1,x_1)^{-1}\gamma(x_1)$. Then we remove from γ its orthogonal projection on $\mathcal{L}_W^{-1}(\delta_{x_1}^{\epsilon_k})$ (i.e. the first basis vector in the space W). We iterate then the procedure on the residual: $\gamma_1(x) = \gamma(x) - K^W(x,x_1)\alpha_1^1$. Eventually, the algorithm can be written as:

Algorithm 1 Orthogonal Matching Pursuit for Currents

Input: a vector field γ, a threshold $\eta > 0$
$\gamma_0 = \gamma$, $N = 0$
while $\|\gamma_N\|_\infty \geq \eta$ **do**
-$x_{N+1} = \mathrm{argmax}_{x\in\mathbb{R}^3} \|\gamma_N(x)\|_{\mathbb{R}^3}$
-Find $(\alpha_i^{N+1})_{1\leq i\leq N+1}$ by solving
$\sum_{p=1}^{N+1} \left(K^W(x_i,x_p)\alpha_p^{N+1}\right)_k = \gamma(x_i)_k$
-$\gamma_{N+1} = \gamma - \sum_{i=1}^{N+1} K^W(.,x_i)\alpha_i^{N+1}$
-$N \leftarrow N+1$
end while
Output: list of $(x_i,\alpha_i^N)_{i=1\ldots N}$

After N steps, the algorithm gives an approximation of $T = \mathcal{L}_W(\gamma)$ with N delta Dirac currents: $\Pi_N(T) = \sum_{i=1}^N \delta_{x_i}^{(\alpha_i^N)}$. The couples (x_i,α_i) are called *momenta* in the sequel. We prove in Appendix A that $\Pi_N(T)$ converges to T as N tends to infinity (e.g. $\|\Pi_N(T) - T\|_{W^*} \longrightarrow 0$). The auxiliary variable $\gamma_N = \gamma - \mathcal{L}_W^{-1}(\Pi_N)$ stores the residual vector field that remains to be explained. We prove also in Appendix A that the L^∞-norm of this residue tends towards zeros as N tends towards infinity (e.g. $\|\gamma_N\|_\infty = \sup_{x\in\mathbb{R}^3} |\gamma_N(x)| \longrightarrow 0$). This means, in particular, that the norm of the residue is below any positive threshold in finite time, thus proving that the algorithm finishes.

3.2.2. *Computations in a discrete setting*

In the continuous setting, finding the maximum of γ_N over the whole $3D$ space (step 1) and computing γ_N at every point of the $3D$-space (step 3) might be difficult and time-consuming. To efficiently implement the algorithm, we set a linearly spaced grid $\Lambda = \{x_\lambda\}$ (supposed to be large enough to assume periodic boundary conditions) with step Δ. Since the input vector field γ is band-limited (it results from a convolution with a Gaussian kernel), we sample it on the grid Λ. For this purpose, we assume that Δ/λ_W is small (typically less than 1/5). γ can be stored therefore as an image of vectors $[\gamma]_\lambda = \gamma(x_\lambda))$.

At step 1 of the algorithm, we look for the maximum of γ_N only at the points of the grid. As a result, every estimated point x_i is on the grid Λ. Since K^W is a translation invariant scalar kernel (i.e. of the form $K^W(x,y) = k(\|x-y\|)\mathrm{Id}$), the vector field $\mathcal{L}_W^{-1}(\Pi_N)(x) = \sum_{i=1}^N K^W(x,x_i)\alpha_i^N$ sampled at the points of the grid may be computed by a circular convolution between the image $[k]_\lambda = k(\lambda\Delta)$ and the image $[\alpha]_\lambda$, where $[\alpha]_\lambda = \alpha_p^N$ if $x_\lambda = x_p$ and 0 otherwise. This circular convolution may be computed efficiently by Fast Fourier Transforms (FFTs). Note that the Fourier transform of the image $[k]_\lambda$ can be pre-computed once for all. As a consequence, at step 3 one needs to (1) project the momenta (x_i,α_i) on the grid's nodes (such that the tri-linear interpolation of the vectors on the grid's node retrieves α_i at point x_i), (2) compute a FFT of this image of vectors and multiply it with the FFT of the kernel, and (3) compute the FFT inverse of the resulting image of vectors. The output Π_N is stored as a list of momenta $(x_i,\alpha_i^N)_{i=1\ldots N}$.

The threshold η has to be specified for every application. If T is a linear combination of N currents $T_1,\ldots,T_N$ (such as mean current, principal mode, difference between two currents, etc.), we can choose η as a fixed ratio of the standard deviation of the set of currents: $\eta = \tau\sigma$ where $\sigma^2 = \frac{1}{N-1}\sum_{i=1}^N \left\|T_i - \bar{T}\right\|_\infty^2$ and $\bar{T} = \frac{1}{N}\sum_{i=1}^N T_i$ (for sake of simplicity, $\|T\|_\infty$ denotes $\left\|\mathcal{L}_W^{-1}(T)\right\|_\infty$). This means that the algorithm finishes when the maximum approximation's error is smaller than $\tau\%$ of the standard deviation. In our applications, we usually fix this sparsity parameter $\tau = 5\%$.

Finally, for a given current T, we have 3 distinct representations: the initial one with N_T momenta (the total number of segments or mesh's cells in the set of currents), the projection of T onto the grid with *a priori* N_{grid} momenta (which depends on both the spreading of the input data and the step Δ supposed to be smaller than $\lambda_W/5$) and our sparse representation with N_{mom} momenta. Depending on the number of input points, their spreading and redundancy at the scale λ_W, these 3 representations can vary dramatically in size.

3.3. *Application to a simulated example*

We show here how the matching pursuit algorithm helps to approximate statistics on currents, and to give them a visual representation that we can interpret. We choose the mean of two 2D curves as a simulated example. In Fig. 1, we show the initial mean in the space of currents and the first and third iterations of the algorithm, which builds iteratively an approximation of the mean. The approximation error tends to zero as the algorithm goes on. In Fig. 2 we show how the number of momenta needed to represent the mean of the two curves vary with respect to the standard deviation of the kernel (λ_W) for the a fixed approximation error. The greater λ_W, the closer the two curves in the space of currents (i.e. their differences become small perturbations at the scale λ_W, the smaller the number of Delta Dirac currents needed to represent the mean for the same accuracy (i.e. the faster the terms of the series decrease).

4. Joint statistics on deformations and residual currents

In the previous section, we presented a method to perform statistics on a population of shapes. However, this statistical analysis may not be relevant with unregistered shapes. A common practice in medical imaging is first to register shapes into a common reference frame and then to perform statistics on the residues which remain after registration, as in voxel-based morphometry. Nevertheless, such an approach is not completely satisfactory since deformations may also capture interesting anatomical features. One wants to perform statistics on the deformations as well, like in tensor-based morphometry. In this section, we propose a unifying framework which estimates jointly a prototype shape (also called template or atlas in the literature) and the deformations of this template to each observation. We can then compute *consistently* statistics both on the deformations and on the residual shapes in the space of currents.

4.1. *Forward versus backward models for template estimation*

In the medical imaging field, atlases are used 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 (Joshi et al., 2004; Avants and Gee, 2004; Marsland and Twining, 2004; Zollei et al., 2005), much fewer were proposed for anatomical curves or surfaces (Chui et al., 2004; Glaunès and Joshi, 2006; Durrleman et al., 2008a). 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 observation back to a common reference frame (See Fig.3). However, we prefer here to base our statistical estimation on a *forward* model, as pioneered in Allassonnière et al. (2007); Ma et al. (2008), 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:

$$T_i = \phi_{i*}\bar{T} + \varepsilon_i \tag{10}$$

whereas the *backward* model is:

$$\phi_{i*}T_i = \bar{T} + \varepsilon_i \quad \Longleftrightarrow \quad T_i = {\phi_i^{-1}}_*\bar{T} + {\phi_i^{-1}}_*\varepsilon_i \tag{11}$$

The backward model considers either that the template $\bar{T}$ is noisy and the observations T_i free of noise (Eq.11-left), or that the noise added the observations (${\phi_i^{-1}}_*\varepsilon_i$) *depends* on the observations via an unknown deformation (Eq.11-right). By contrast, the forward model (Eq.10) considers that the template is not blurry, as an "ideal" object, and an independent and *identically distributed* noise ε_i is added to every observation. This models more accurately the physical acquisitions, whereas the backward model relies on less realistic but more practical assumptions, as we shall see below.

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. In the backward model, sampled observations deform a continuous template: an extrinsic *interpolation* scheme is required. By contrast, in the forward setting, the continuous template is deformed to the subject's space and then needs only to be *sampled* to be compare 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

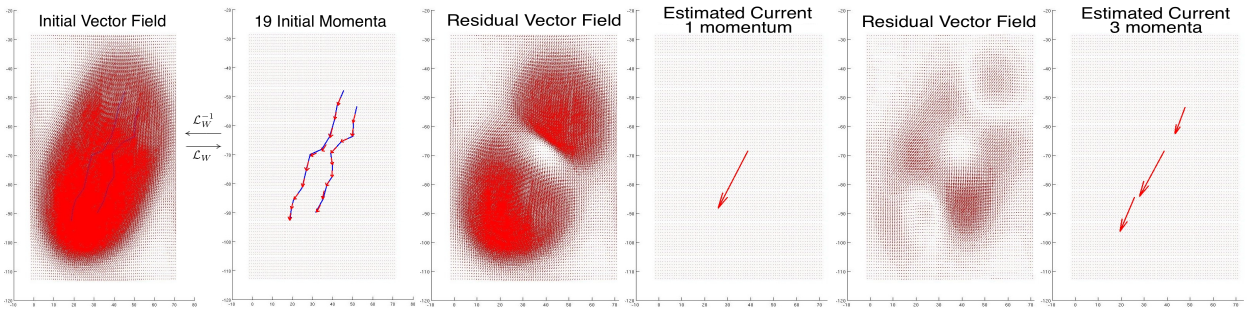

(a) Initial Configuration (b) iteration 1 ($\|R_1\|_\infty = 0.51\sigma$) (c) iteration 3 ($\|R_3\|_\infty = 0.11\sigma$)

Figure 1. A sparse deconvolution scheme for currents: (a) the initial configuration. right: two curves in blue and their mean in red: the collection of all tangents scaled by 0.5, seen as momenta in the space of currents W^*. left: the Gaussian convolution ($\mathcal{L}_W^{-1}$) of the initial momenta gives the dual representation of the mean as a dense vector field in W ($\lambda_W = 15$). (b) (resp. (c)): first (resp. third) iteration of the matching pursuit algorithm: estimated momenta on the right panel, residual vector field on the left panel (what remains to be explained by the forthcoming momenta). The momenta converge to the true solution while the residual vector field tends to zero. $\|R_n\|_\infty$ denotes the norm of the residual vector field at step n and σ the standard deviation: $\|L - L'\|_\infty / \sqrt{2}$

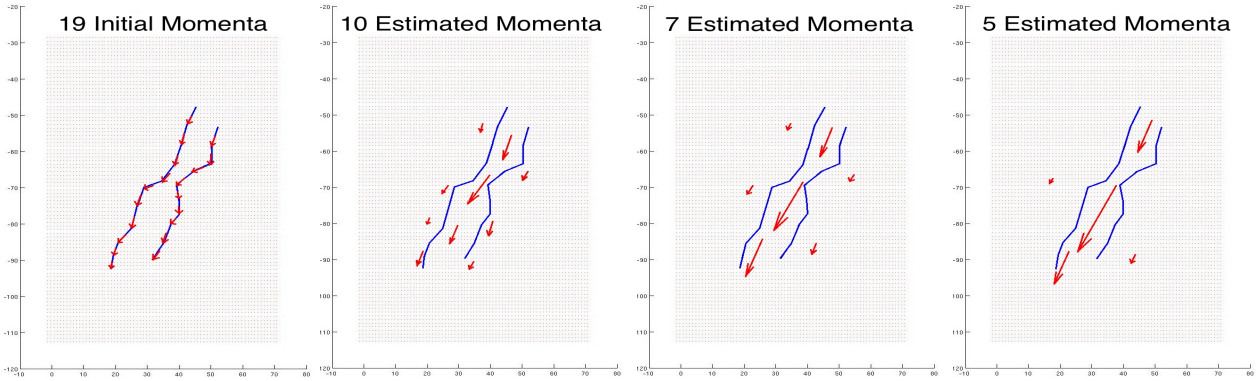

(a) Initial mean current (b) $\lambda_W = 10$ (c) $\lambda_W = 15$ (d) $\lambda_W = 20$

Figure 2. Impact of λ_W. (a)- Two curves in blue and their mean in the space of currents in red. (b) to (d) the approximation of the mean of two curves for increasing size of the kernel λ_W and fixed precision $\tau = 5\%$ (defined as the ratio between the residual norm and the standard deviation $\|L - L'\|_\infty / \sqrt{2}$). The greater λ_W, the closer the two curves in the space of currents, the more redundant the initial momenta at the scale λ_W, the sparser the estimated decomposition.

would require at least to compute the probability of having the template given a training database of T_i: $p(\bar{T}|T_i)$. Once the atlas is built, one would like to know how a new observation T_{new} is compared to the estimated variability model: one needs to compute the likelihood of this observation given the template $p(T_{\text{new}}|\bar{T})$. Because ϕ_i acts differently in Eq.10 and in Eq.11, the computational cost of these two steps varies significantly. In the backward scheme, computing $p(\bar{T}|T_i)$ is much simpler than computing $p(T_{\text{new}}|\bar{T})$ which depends on the Jacobian of the deformations ϕ_i. It is exactly the reverse for the forward scheme: computing the atlas is more difficult than to compare a new observation to the estimated variability. 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. The backward scheme seems simpler as long as one does not compute joint statistics on the deformation ϕ and the residual perturbation ε.

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 (Allassonnière et al.,

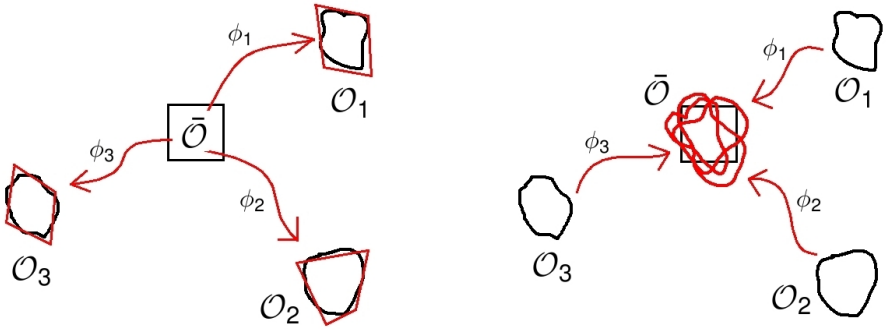

forward scheme backward scheme

Figure 3. 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.

2007). 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*. In the sequel, we show how the atlas building step, which is the most critical step in this paradigm, is possible in case of curves and surfaces. For this purpose, we take advantage of the Matching Pursuit Algorithm for currents, introduced previously in the section 3.2.

Joint statistics on deformations and residues will be used to give a complete description of the anatomical variability of white matter fiber bundles. Numerical simulations on simulated examples will also show how the atlas may be used for group classification in the context of supervised learning.

4.2. *Joint estimation of template and deformations*

4.2.1. *A Heuristic Maximum A Posteriori in infinite dimension*

From a Bayesian point of view, in Eq.10 ($T_i = \phi_{i,*}\bar{T}+\varepsilon_i$), 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 σ_W^2. To define statistics on deformations, we take advantage of the tangent-space representation of the diffeomorphisms. As explained in section 2.4, the diffeomorphisms are completely determined by their initial speed vector fields v_0 which belong to a Hilbert space V.

Let us assume 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: $p_\varepsilon(\varepsilon) = C_\varepsilon \exp(-\|\varepsilon\|^2_{W^*}/2\sigma_W^2)$ and $p_\phi(v) = C_\phi \exp(-\|v\|^2_V/2\sigma_V^2)$. In that case, a Maximum A Posteriori (MAP) estimation for independent observations maximizes $\prod_{i=1}^N p(T_i|\bar{T})$ over $\bar{T}$. Formally,

$$\begin{aligned} p(T_i|\bar{T}) &= \int p_\varepsilon(Ti|\bar{T}, v_0^i) p_\phi(v_0^i) dv_0 \\ &= \int p_\varepsilon(T_i - \phi_i^{v_0^i}{}_*\bar{T}) p_\phi(v_0^i) dv_0 \end{aligned} \tag{12}$$

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:

$$(\bar{T}, \phi^{v_0^i}) = \text{argmin}_{\bar{T}, v_0^i} \left\{ \frac{1}{\sigma_W^2} \left\| T_i - \phi^{v_0^i}{}_*\bar{T} \right\|^2_{W^*} + \frac{1}{\sigma_V^2} \sum_{i=1}^N \left\| v_0^i \right\|^2_V \right\} \tag{13}$$

called Fast Approximation with Mode (FAM).

However, as explained in Section 2.3, the Gaussian variables have no pdf in the infinite dimensional Hilbert spaces, such as the space of currents W^* and the space of initial vector fields V. Nevertheless, the restriction of these variables to finite dimensional spaces does have a Gaussian pdf. Therefore, a more rigorous MAP derivation could be done considering finite dimensional parameterization 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 Vaillant and Glaunès (2005) to account for such a discretization. Moreover, to avoid the approximation with mode (FAM), Markov Chain Monte Carlo (MCMC) approaches for sam-

pling the posterior could be also possible along the lines of Allassonnière and Kuhn (2009) but this is still challenging and out of the scope of this paper.

4.2.2. *An alternated minimization procedure*

We solve Eq.13 by minimizing it alternatively with respect to the template and to the deformations. When the template $\bar{T}$ is fixed, each term of Eq. 13 can be minimized separately. For a given observation T_i, minimizing $\frac{1}{\sigma_V^2}\left\|v_0^i\right\|_V^2+\frac{1}{\sigma_W^2}\left\|T_i-\phi_i^{v_0^i}{}_*\bar{T}\right\|_{W^*}^2$ with respect to v_0^i is exactly a registration problem, as stated and solved in Vaillant and Glaunès (2005); Glaunès (2005). 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 Eq. 13 with respect to the template $\bar{T}$ leads to the minimization of the convex function:

$$J(\bar{T}) = \frac{1}{2}\sum_{i=1}^{N}\left\|\phi_{i*}\bar{T}-T_i\right\|_{W^*}^2 \tag{14}$$

If all $\phi_i = \mathrm{Id}$ (i.e. no deformation), the minimum is reached at the empirical mean: $\bar{T} = \frac{1}{N}\sum_i^N T_i$. For arbitrary deformations, there is no closed form and we use a gradient descent scheme. The gradient of Eq.14 is precisely:

$$\nabla_{\bar{T}}J = \sum_{i=1}^{N}\phi_{i*}^{\dagger}(\phi_{i*}\bar{T}-T_i) \tag{15}$$

where $\phi_*^\dagger$ is the adjoint action of ϕ_* defined by: $\left\langle\phi_*^\dagger T, T'\right\rangle_{W^*} = \langle T, \phi_* T'\rangle_{W^*}$ for any currents T and T'. This would be a matrix transpose if the action were linear (like for an affine deformation points for instance). In this non-linear setting, standard computations lead to $\phi_*^\dagger T = \mathcal{L}_W(\phi^*\mathcal{L}_W^{-1}(T))$. In the backward scheme, Eq.14 would be: $J(\bar{T}) = \frac{1}{2}\sum_{i=1}^N\left\|\bar{T}-\phi_{i*}T_i\right\|_{W^*}^2$, whose minimum has the closed form $\bar{T} = \frac{1}{N}\sum_{i=1}^N\phi_{i*}T_i$, which is the mean of the observations pulled back into the current template configuration. By contrast, in our forward setting, the estimation of the template involves the Jacobian matrix of the deformations (via ϕ_* as we will see below). However, computing the likelihood of any new observations $p(T_i|\bar{T})$ will be much simpler and faster in this setting.

We must now make explicit the computation of the gradient in Eq. 14. 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 $\phi_{i*}\bar{T}-T_i$ is of the form $\sum_k \delta_{y_k^i}^{\beta_k^i}$ which gives finally $\phi_i^*\mathcal{L}_W^{-1}(\phi_{i*}\bar{T}-T_i) = \sum_k (d_x\phi_i)^t K^W(\phi_i(x), y_k^i)\beta_k^i$ in case of curves. Finally, the dual representation of the gradient in the space of vector field W is given at any point x by:

$$\mathcal{L}_W^{-1}(\nabla_{\bar{T}}J)(x) = \sum_{i=1}^{N}(d_x\phi_i)^t\left(\sum_k K^W(\phi_i(x), y_k^i)\beta_k^i\right) \tag{16}$$

in case of curves. For surfaces, $d_x\phi_i^t$ must be replaced by $|d_x\phi_i|\, d_x\phi_i^{-1}$. Thus, we see that we know how to compute the dual representation of the gradient (a vector field) but not the gradient itself (a current). 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.16. The sparse deconvolution method of Section 3.2 precisely provides a finite set of Dirac current which approximates ∇J at any accuracy. As a consequence, the template remains a finite set of Dirac currents at each iteration.

We initialize the algorithm by setting $\phi_i = \mathrm{Id}$, $\bar{T} = \frac{1}{N}\sum_{i=1}^N T_i$ and by computing the vector field $\gamma_{\bar{T}} = \mathcal{L}_W^{-1}(\bar{T})$ associated to $\bar{T}$ via a Gaussian convolution (computed with FFT's on images of vectors as in Section 3.2.2). The current $\bar{T}$ is encoded as a list of (position, vectors) that approximates small segments or small triangles. The dense vector field $\gamma_{\bar{T}}$ is discretized at a the points of a fixed grid: $\Lambda = \{x_p\}$ and is therefore encoded as an image of vectors. The template estimation algorithm is written as in Algorithm 2, where the auxiliary variable grad is also encoded as an image of vectors.

As a result, the algorithm returns (1) an unbiased template $\bar{T}$ and (2) the deformations of this template to every observation T_i. The residues are given in the space of currents by $T_i-\phi_{i*}\bar{T}$. The methodology developed in Section 3 can be used therefore to perform statistics on such residual currents. These residues model the variability which are not captured by regular deformations such as topology changes, matter creation or deletion, numerical or physical noise, etc.

4.3. *Atlas construction on simulated 2D-curves*

We illustrate here the template estimation procedure on a simulated example of 2D curves, and show

Algorithm 2 Atlas Construction

Input: N shapes T_i (stored as a list of oriented points (tangents or normals): $T_i = (x_i^p, \alpha_i^p)$), a grid Λ.
List $\bar{T} = (\bar{x}_k, \bar{\alpha}_k) \leftarrow$ concatenation of all $(x_i^p, \alpha_i^p/N)$ (encodes $\sum_{i=1}^N T_i/N$)
Image of vectors $\gamma_{\bar{T}} \leftarrow \mathcal{L}_W^{-1}(\bar{T})$ (Convolution: $\forall x_p \in \Lambda, \gamma_{\bar{T}}(x_p) = \sum_k K^W(x_p, \bar{x}_k)\bar{\alpha}_k$)
repeat
  for $i = 1 \dots N$ **do**
    $\phi_i \leftarrow$ registration of $\bar{T}$ to T_i.
  end for
  repeat
    Image of vectors grad $= 0$
    for $i = 1 \dots N$ **do**
      Deform $\bar{T}$ with ϕ_i: $\phi_{i*}\bar{T} = (\phi_i(\bar{x}_k), d_{\bar{x}_k}\phi_i(\bar{\alpha}_k))$ (curves) or $(\phi_i(\bar{x}_k), |d_{\bar{x}_k}\phi_i|\, d_{\bar{x}_k}\phi_i^{-t}(\bar{\alpha}_k))$ (surfaces).
      Concatenate the list $(x_i^p, -\alpha_i^p)$ with the previous one to give (y_k^i, β_k^i) (encodes $\phi_{i*}\bar{T} - T_i$).
      Deform Λ with ϕ_i
      for all $x_p \in \Lambda$ **do**
        Compute $d_{x_p}\phi_i$ by a finite difference scheme
        Compute $G = \sum_k K^W(\phi_i(x_p), y_k^i)\beta_k^i$ (convolution)
        $$\text{grad}(x_p) \leftarrow \text{grad}(x_p) + \begin{cases} 2d_{x_p}\phi_i^t G & \text{for curves} \\ 2\left|d_{x_p}\phi_i\right| (d_{x_p}\phi_i)^{-1} G & \text{for surfaces} \end{cases}$$
      end for
    end for
    $\gamma_{\bar{T}} \leftarrow \gamma_{\bar{T}} - \tau\,\text{grad}$
    Deconvolution of $\gamma_{\bar{T}}$ to give the new $\bar{T} = (\bar{x}_k, \bar{\alpha}_k)$ (See section 3.2).
  until convergence
until convergence
Output: One template $\bar{T}$, N deformations ϕ_i, N residues $\phi_{i*}\bar{T} - T_i$ (stored as a list of momenta)

how this atlas can be used for group comparison. We build a synthetic database with two classes made of 40 curves each (see Fig. 4). First, we construct the atlas (template and the variations of this template in the population) from all 80 curves. The parameters used for this simulation are the scale of deformations $\lambda_V = 0.4$, the spatial scale of currents $\lambda_W = 0.05$, the trade-off between regularity and fidelity-to-data (for registrations) $\gamma = 10^{-4}$ and the sparsity number $\tau = 5\%$. The iterative construction of the atlas is shown in Fig. 5. The difference between the estimated template and true template (used to create the database) is equal to 0.24 times the standard deviation (the norm being computed in the space of currents). With 80 samples, the t-statistics is equal to $\sqrt{80} * 0.24 = 2.15$, which is below the usual 3σ threshold to decide statistical significance.

Along with the template, we estimate the diffeomorphisms that map the template to each observation. These diffeomorphisms are entirely determined by their initial speed vector fields at the points of the template (Vaillant et al., 2004). We can use this tangent-space representation of diffeomorphisms as a feature to perform group statistics. First, to reduce the dimensionality of the feature space (we have much fewer observations than the dimension of the feature space which is three times the number of points of the template), we perform a PCA on all data together and keep the N_{modes} first principal directions. Then, we project the initial speed vector fields onto this subspace and estimate, for each class, the mean and empirical covariance matrix. We end up therefore with two estimated Gaussians laws on the same subspace which characterize the geometrical variability of each class.

To evaluate the performance of our statistical model, we measure how well it can separate populations. Given any new observations, we register the template to this observation. According to our forward setting, the likelihood of such a registration is given directly by the likelihood of its initial vector field (defined in the common template space) with respect to the estimated law. By contrast, in a backward scheme, the initial vector speed of

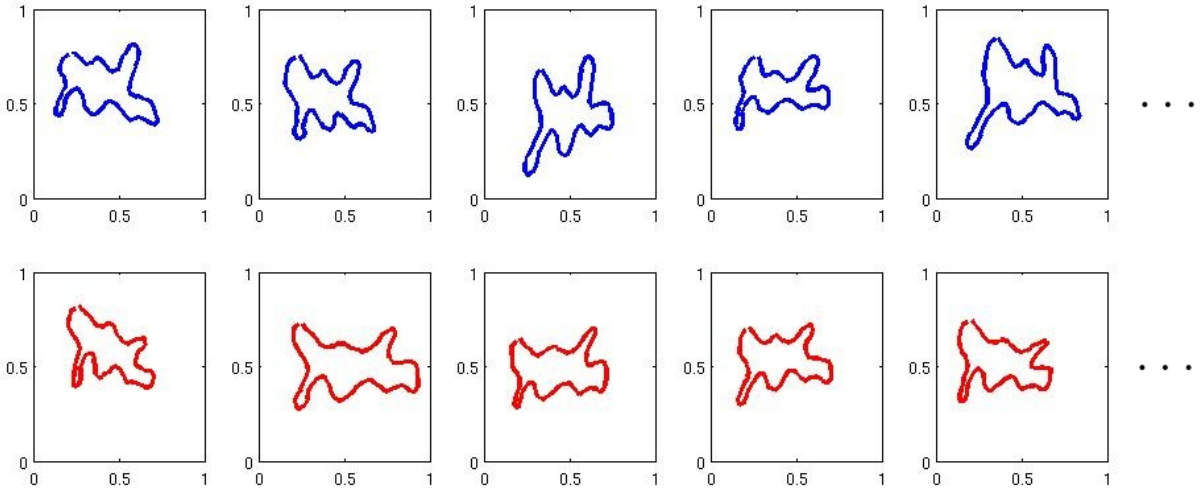

Figure 4. Five samples among the 40 samples for each class. Our goal is to show that our atlas estimation allows us to find the geometrical discriminative features between both classes.

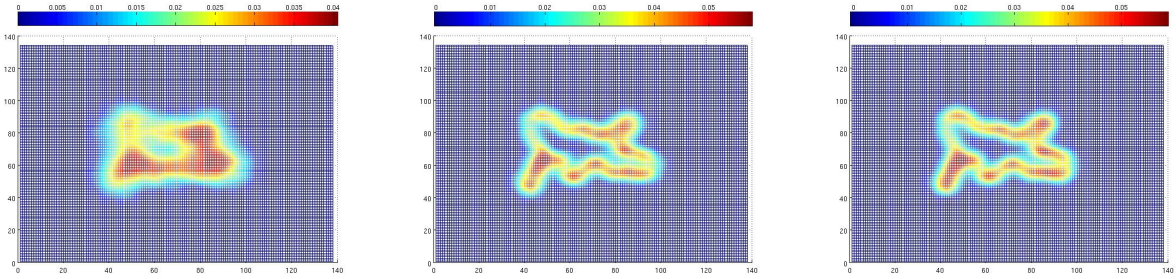

a- mean current b- iteration 1 c- iteration 2

Figure 5. Atlas construction from 80 samples. At each iteration the algorithm register the template (shown here) to each sample and update the template according to the deformations. The figures show the dual representation of the template in the space W at each iteration until convergence (we removed here the direction of the tangents and take only points into account: the dual representation is therefore a scalar field (shown here) instead of a vector field. Colors correspond to the magnitude of this scalar field). The initial template is the empirical mean current (a). The next 2 iterations are shown in (b) and (c). Along with the iterations, the bias is removed from the template. As a result, the template appear to be less and less blurred.

the deformation is defined in the subject's space and must be therefore transported on the template space before being compared with the estimated variability model: this requires to deal with the Jacobian matrix of the deformation. In our case, the Mahalonobis distance between the initial speed vector field of this registration (projected onto the common subspace) and the mean of each group determines which population this new observation is the more likely to belong to. If the difference of the two Mahalanobis distances is below some threshold, we associate the observation to a given class.

To evaluate the classification power of our atlas, we test 100 new observations for each class. For a given threshold, we compute the false positive rate and the true negative rate. The entire ROC curve, computed for all possible thresholds, is shown in Fig. 6. In these experiments, $N_{\text{modes}} = 26$ modes. The maximum good classification ratio (mean of the rate of false positive rate and true negative) is of 94%. If we train a Support Vector Machine on the initial vector speed instead of performing PCA, we can achieve a good classification ratio of 97%. This shows that our atlas estimation allows us to

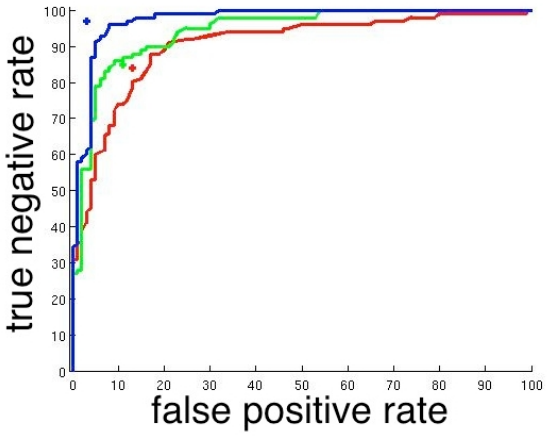

Figure 6. ROC curves measuring the capability of the statistical analysis to classify shapes. The prediction is based on the difference of Mahalanobis distances for each class. 100 data per class were used in the testing step (40 data per class in the training step). Crosses correspond to the classification results based on Support Vector Machines (SVM). In blue, our unbiased atlas was used. In red, one particular subject was used as template. In green, the mean current (initial template in our iterative estimation scheme) is used as template. This shows that the unbiased template better discriminates between both classes.

predict the class of any observations with an error of order 3%.

The same statistics on the deformations may be conducted when we replace our template by the mean current (used for the initialization of our algorithm). Fig. 6 shows that the mean current, which is biased according to our model, leads to worse classification results. If one chooses randomly one sample in the database as template, the classification results are even worse. This shows that computing the statistics in a consistent framework increases the discriminative power of the method.

Our method allows us not only to perform quantitative tests, but also to *describe* the difference between both classes. The classification can be performed using only one mode for each class. Among the 26 modes of each class, we find the two modes (one in each class) which enables to achieve the best classification ratio. In some sense, these two modes are the ones which best discriminate the two populations. In our example, the two most discriminative modes enables to achieve a good classification ratio of 91% (to be compared with the 97% ratio achieved with all possible modes (26)). Deforming the template according to these two modes gives a visual representation of the main difference between the two classes: the torque of the upper right part of

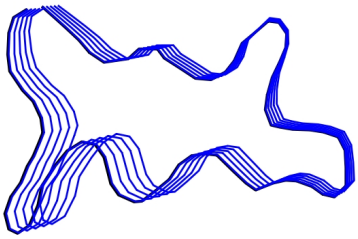

7th mode of the blue class

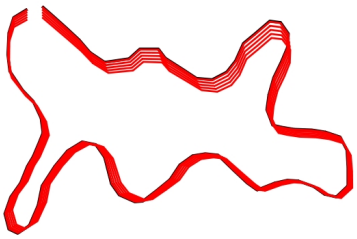

15th mode of the red class

Figure 7. The two most discriminative modes for each class. The main differences between the two populations is a torque at the upper-right part of the shape (if consider the blue class as reference). This shows how the proposed statistical analysis may detect reproducible features across the observations, whereas this signal is hidden by a high random deformation noise. For instance, this signal cannot be detected visually from the input shapes in Fig. 4

.

the shape (if we consider the blue class as reference), as shown in Fig. 7. This difference was detected as a reproducible feature across the observations, although it is almost impossible to guess it from a visual inspection of the initial data in Fig. 4.

At this stage, many other statistical computations may be conducted to quantify and describe the differences between both classes. Our purpose here was only to illustrate the potential of our method for the analysis of anatomical shapes. We proposed a consistent framework to estimate template, deformations of this template within the population and residual shapes (which were not used in this section, but will be used in the next ones). The results on simulated examples tend to show that such an atlas may put into evidence significant anatomical variations between populations, although these differences are hidden by the normal variability of the structures within the populations.

5. Numerical experiments on real anatomical data

5.1. *Choice of parameters*

The whole framework depends mostly on 4 parameters: $\lambda_W, \lambda_V, \gamma, \tau$. The spatial scale of the currents λ_W is the typical distance at which differences between shapes are considered as noise. It compares to the size the features of interest of the input shapes. The spatial scale of the deformations λ_V is the typical distance at which points move consistently during registration. It is related to the degree of rigidity of the deformation and must be therefore compared to the diameter of the data. If λ_V is large compared to the diameter of the data, the deformations would move the shape almost rigidly. On the contrary, small values of λ_V will favor uncorrelated motion of small parts of the shape while increasing the cost of a rigid motion. The trade-off γ is the ratio σ_W^2/σ_V^2 in Eq. 13: it balances the fidelity-to-data against the regularity of the deformation. In practice, these three parameters may be adjusted by computing registrations between pairs of input data. More discussions about the impact of each of these parameters on the registrations may be found in Durrleman et al. (2008c). The sparsity threshold τ is a fraction of the standard deviation of the input shape, which determines the approximation error as explained in Section 3.2.2. As a fraction of the standard deviation, this value does not depend much on the data and can be fixed for every application. Usually, it is set to 5%. A value of 1% may be used, but it increases the number of points of the template and therefore slows down the computation of the atlas.

In this current stage of the work, it is up to the user to estimate manually the values of these parameters, depending on the data and the application. These values may be assessed according to the diameter of the data and the size of the feature of interest. In the following, we give the typical values of these parameters for each experiment.

5.2. *Statistics on sulcal lines*

The sulci are the fissures on the brain surface and they are often used to measure anatomical differences between subjects (Thompson et al., 1996). We perform here statistics on a set of 70 sulci delineated in $N_{obs} = 34$ subjects. These data were provided by Paul Thompson (Laboratory of NeuroImaging, University of California, Los Angeles) as part of the collaborative project `Brain-Atlas`.

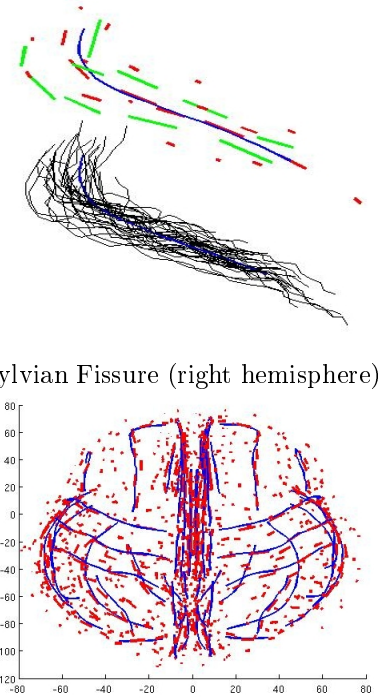

Figure 8. Statistics for 70 sulci in 34 subjects ($\lambda_W = 12\text{mm}, \tau = 5\%$). Top: every subject's Sylvian fissure (black), the mean (red) and the first eigenmode at $+\sigma$ (green) showing the spreading of the set of lines. Bottom: Mean currents (red) compared to the mean lines (blue) computed from B-spline parameterization of curves (Fillard et al., 2007b). Results are in good agreement.

For each sulcal line, we approximate the mean current $\bar{L} = \frac{1}{N_{obs}} \sum_{i=1}^{N_{obs}} L_i$ for the scale of kernel $\lambda_W = 12\text{mm}$ and sparsity parameter $\tau = 5\%$. Results are shown in Fig.8 for the Sylvian Fissure of the right hemisphere and for all 70 sulci. The initial number of momenta for the mean fissure was $N_T = 899$ (i.e. the number of segments of all lines) whereas the final approximation needs only $N_{mom} = 54$ momenta. In this case, the compression ratio is of 94%. Considering all sulci, the compression rate is on average: $94.8\% \pm 0.02$. Each line's grid has a step $\Delta = 2\text{mm}$ and typically $N_{grid} = 10^5$ points. Our mean is visually in good agreement with other mean curves computing from B-spline representation (Fillard et al., 2007b).

Then, we compute the eigenmodes of the lines sets. We find the eigenvectors (V_k) of the $N_{obs} \times N_{obs}$

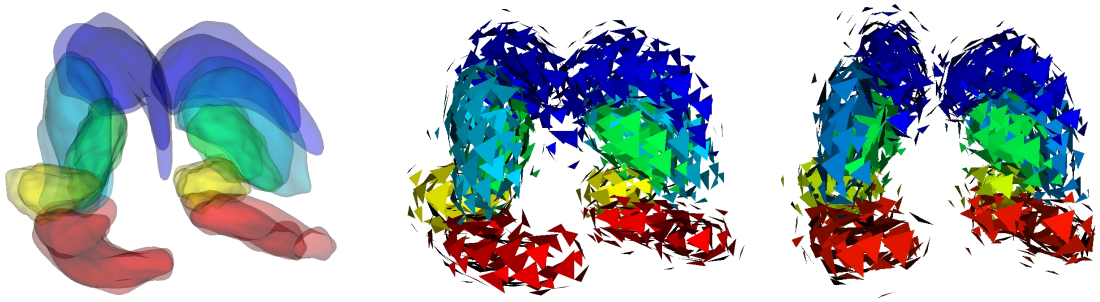

(a) Deep brain structures of 2 autistics (b) Mean of 50 autistics (c) Mean of 7 controls

Figure 9. Sparse approximation of the mean current for 10 meshes segmented in 50 autistics patients (b) and 7 controls (c). $\lambda_W = 5\text{mm}, \tau = 5\%$ and the diameter of the data is 60mm.

matrix $\Sigma = \left(\left\langle L_i - \bar{L}, L_j - \bar{L} \right\rangle_{W^*}\right)_{i,j=1\ldots N_{obs}}$. The k^{th} eigenmode is given by the linear combination of input currents: $m_k = \bar{L} \pm \sum_i (V_k)_i (L_i - \bar{L})$. We approximate the first eigenmode at $+\sigma$ of the Sylvian Fissure of the right hemisphere (Fig.8-top): this mode captures mainly the spreading of the lines set around its mean.

5.3. *Statistics on surfaces of brain structures*

In this section, we turn to statistics on 10 meshes of deep brain structures (Caudate, Putamen, Globus Pallidus, Amygdala and Hippocampus for each hemisphere) segmented from a population of 50 autistics patients and 10 controls (Hazlett et al., 2005). These data were kindly provided by Guido Gerig (Scientific Computing and Imaging Institute, University of Utah) and Martin Styner (Computer Science Department, University of North Carolina, Chapel Hill).

We compute the mean meshes for each population and approximate these means with the Matching Pursuit Algorithm. Results are presented in Fig. 9. Note that for surfaces, we represent the estimated momenta (normal of an infinitesimal mesh cell) as equilateral triangles whose normals is the momenta. The difference between both means is still a current that we approximate: the arrows of Fig.10 are the 10 first estimated momenta of this difference, suggesting that the autistic mean is more curved at the Hippocampus' extremity and thicker in the middle. Such results still need to be confirmed by rigorous statistical tests.

Within the autistic group, the compression ratio

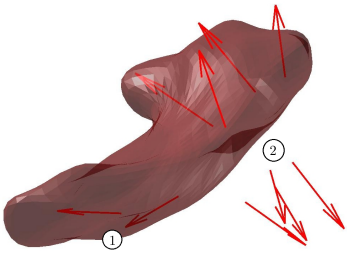

Figure 10. The arrows represent the difference between the mean of autistics and the mean of controls (shown in Fig. 9. They are superimposed with the Hippocampus of a control. This shows that the mean from autistics is more curved at hippocampus' extremity (area 1) and thicker in area 2.

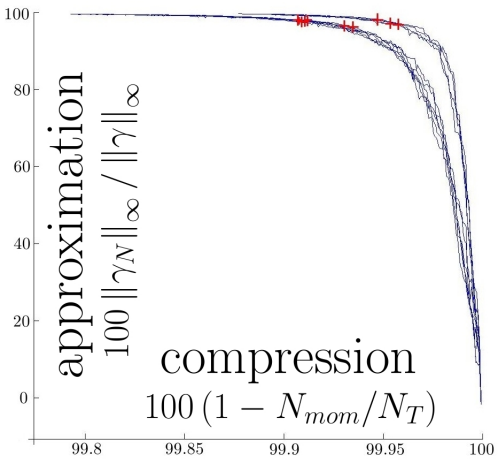

Figure 11. We approximate each of the 10 meshes for every autistics with $\lambda_W = 5\text{mm}$ (see Fig.9). The graph shows the evolution of the approximation error when the number of momenta is increasing. Very high compression rate can be achieved while the approximation error remains small. Red points correspond to the approximation error equal to $\tau = 5\%$ of the variance of the structures within the population.

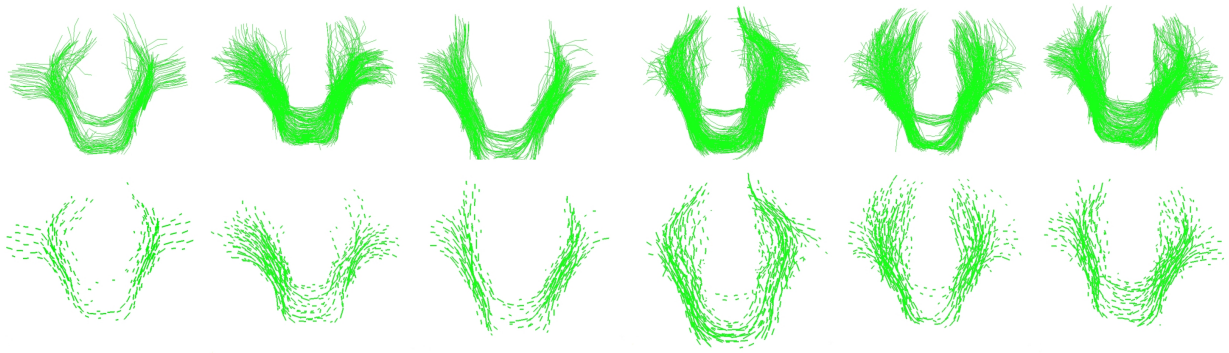

Figure 12. **Cortico-bulbar tract segmented in 6 subjects**. Original data (top row) and their approximations (bottom row) obtained from the Matching Pursuit Algorithm with the spatial scale of currents $\lambda_W = 3$mm (diameter of the data are 100mm) and sparsity number $\tau = 5\%$. For visualization purposes, the segments of the approximations have been scaled by 0.1. The length of these segments encode the local redundancy of the initial bundles at the scale λ_W.

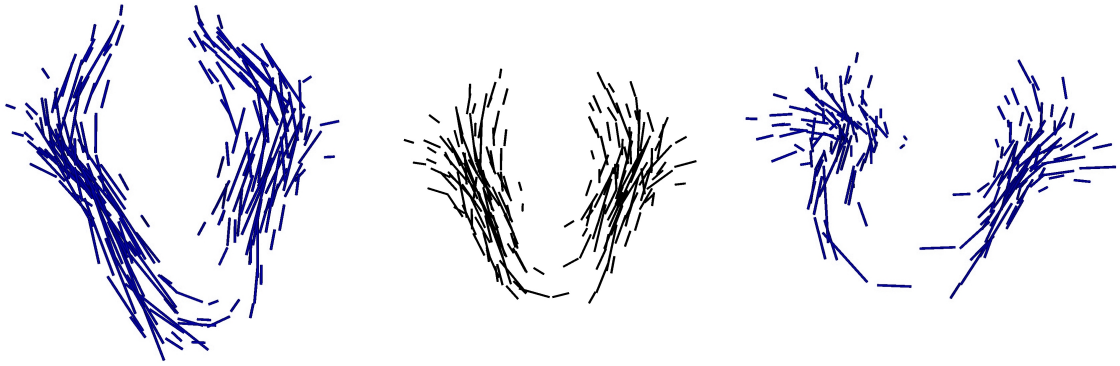

(a) def. mode at $-\sigma$ $\phi_*^{-m_v}\bar{B}$ (b) template $\bar{B}$ (c) def. mode at $+\sigma$ $\phi_*^{+m_v}\bar{B}$

Figure 13. **First deformation mode at $\pm\sigma$ of the corticobulbar bundle (view from top)**. This mode captures the diffeomorphic variability of the population around the prototype bundle (b). It shows a torque at the basis of the bundle and a stretching/shrinking effect of the left and right parts of the bundle.

between N_T and N_{mom} for the 10 structures is on average of $99.96\% \pm 10^{-4}$. The grid of each structure has a step $\Delta = 1$mm and for one structure we have the following typical values: $N_T = 50 \times 3000 = 1.5 \times 10^5$, $N_{grid} = 3 \times 10^5$ and $N_{mom} = 100$. Fig.11 shows that the quality of approximation remains good until very high compression ratio.

5.4. *Anatomical variability of white matter fiber bundles*

In this section, we apply our method to white matter fiber bundles. Diffusion Tensor Images (DTI) of 6 healthy subjects were provided by Denis Ducreux (Hôpital du Kremlin Bicêtre, Paris) and processed by Pierre Fillard (CEA, Neurospin, Saclay). Fiber tractography was performed using MedINRIA [3] (Toussaint et al., 2007), which includes a robust tensor estimation and a streamline tractography algorithm using log-Euclidean tensor interpolation (Fillard et al., 2007a). From all extracted fibers, the cortico-bulbar bundle was selected for our study.

Each extracted cortico-bulbar bundle is a set of curves. It can be represented naturally as the sum of the currents of each individual curve. This representation, in particular, is robust to change of connectivity, fiber interruption or reconnection within the bundle. It takes into account the local orientation of the fibers at each point of the bundle.

[3] www-sop.inria.fr/asclepios/software/MedINRIA

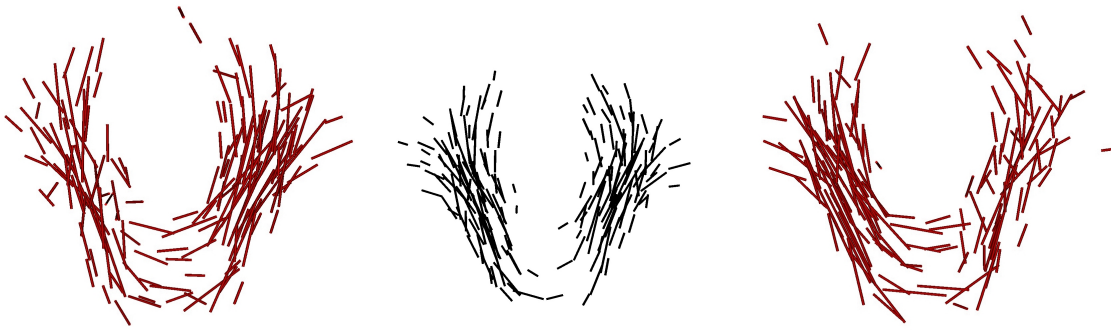

(a) texture mode at $-\sigma$ $\bar{B} - m_\varepsilon$ | (b) template $\bar{B}$ | (c) texture mode at $+\sigma$ $\bar{B} + m_\varepsilon$

Figure 14. **First "texture" mode at $\pm\sigma$ of the corticobulbar bundle.** This mode captures the residual variability of the population, once the diffeomorphic variability has been discounted. This mode is added to (c) or removed from (a) the prototype bundle (b). The mode at $+\sigma$ (resp. $-\sigma$) shows that the left (resp. right) part of the bundle becomes thicker, while its right (resp. left) part becomes thiner. Further investigation should determine whether this asymmetry is a true anatomical feature or an artifact of the tracking method.

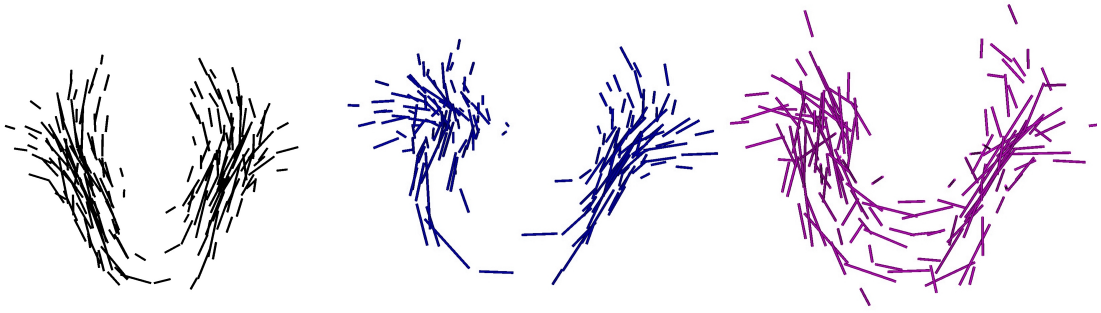

(a) template $\bar{B}$ | (b) def. mode at $+\sigma$ $\phi_*^{+m_v}\bar{B}$ | (c) def. mode + texture mode $\phi_*^{+m_v}\bar{B} + m_\varepsilon$

Figure 15. **Simulation of synthetic corticobulbar bundle.** Our statistical model learn how to generate new bundles from the prototype, which compare to the data. First we deform the template (a) according to the first mode of deformation (b). We then add the first texture mode to this deformed template, to give a new instance of the bundle (c). Such simulations reproduce the common features that our model detect across the subjects.

First, we approximate each bundle via the Matching Pursuit Algorithm with the spatial scale of currents $\lambda_W = 3$mm and the sparsity parameter $\tau = 5\%$. This reduces the total number of segments in the database from 26011 to 2927, namely a reduction of 88.7% while the approximation error remains below 5% of the variance. The approximated momenta, represented by segments in Fig. 12, account for the redundancy of the fibers at the scale λ_W. This approximation allows us to speed up the registrations.

The application of our atlas estimation method leads to the computation of a template bundle, the deformations of this template to each subject and the residual perturbations that cannot be explained by these smooth deformations. Analyzing the deformations and the residual "texture" enables to describe the variability in terms of smooth deformations (torque, stretching, shrinking, etc.) and in terms of non diffeomorphic variations (matter creation or deletion, change of topology, etc.) For this experiment, we set the spatial scale of deforma-

tions to $\lambda_V = 20$mm, the spatial scale of currents $\lambda_W = 5$mm, the trade-off between regularity and fidelity-to-data: $\gamma = 10^{-4}$ and the sparsity parameter to $\tau = 5\%$.

The "geometrical" variability is captured by the deformations. As a result of the MAP estimation, the deformations appear to be centered: the norm of the mean parameters is 0.42 times the standard deviation, which is not significantly different from 0 (t-statistics is equal to $\sqrt{6} * 0.42 = 1.03$). The first mode of the deformations at $\pm\sigma$ is shown in Fig. 13. Variations are mostly a torque of the frontal part of the bundle, as well as a stretching/shrinking of its lateral parts. Further investigation should determine whether this torque is related to the one of the occipital lobe of the cortex.

The variability in terms of "texture" is captured in the residual perturbations. The residues are centered: the mean current is 0.36 times its standard deviation, which is not significantly different from 0. The first residual mode m_ε is shown in Fig. 14. It shows an asymmetry in terms of the number of fibers in each lateral part of the bundle. This result shows, undoubtedly, that the variability left aside from the diffeomorphisms is not pure noise, but may contain interesting anatomical feature. In our case, further investigation is needed to determine whether this fiber creation/deletion effect is due to a true anatomical variability or to an artifact of the tracking algorithm.

Our statistical model is generative: we can combine the deformation and the texture variability to simulate new data. This procedure is shown in Fig. 15: The template is deformed according to the first mode of deformation, and the first texture mode is added to the result. This new synthetic bundle can be compared to the input data in order to get an insight into the common features which has been detected in the population. More details on the application of our methodology to white matter fiber bundles can be found in Durrleman et al. (2009).

6. Discussion and conclusion

In this paper, we proposed a Matching Pursuit algorithm for currents to give a sparse representation of any currents. This sparse representation, in some sense, is the optimal representation of the currents for a given spatial scale λ_W. It allows us to give a stable representation of mean and modes of a population of currents, even if the number of observations is growing in the database, or if the sampling of the shapes becomes finer. This representation enables also a better visualization and therefore interpretation of the statistics on currents. Moreover, this light-weighted representation of currents is used to integrate the representation of shapes as current into a general atlas estimation scheme. We define a statistical model in which data are seen as random deformations of an unknown template plus a random perturbation in the space of current. A Maximum A Posteriori approach leads to a consistent estimation of the template (in the space of currents), the deformations of the template to each data and the residual perturbations. Statistics on the diffeomorphic deformations describe the variability in terms of smooth variations like torque, elongation or shrinking effect. Statistics on the residual perturbation in the space of currents describe the variability that cannot be captured by smooth deformations: topology changes, matter creation or deletion, physical or numerical noise, etc. In some sense, this framework completes the instance of the Grenander's pattern theory of shapes (which amounts to do statistics on the deformations only) for the analysis of shape variability. The method is very versatile and may be applied to any kind of geometrical data like point sets, curves, surfaces or volumes.

This framework depends on mainly 4 parameters: the spatial scale of the currents λ_W which gives the typical scale at which anatomical features are taken into account (far below this scale geometric variations are considered as noise), the scale of the deformation λ_V which measure the typical scale at which a structure may deform (this controls the rigidity of the transformation used in our framework) and the trade-off γ which controls the decomposition of the variability into smooth variations and residual perturbations which remain after registration. As a consequence, the template is the representation of a shape with no geometrical details at a scale much smaller than λ_W. If γ tends to infinity, deformations are more and more constrained to remain close to the identity (no deformation), more and more variability are captured by the residuals, less and less by the deformations. These parameters determine therefore the decomposition of the variability in a geometrical part and a "texture" part. In this work, the parameters were set manually for every application, taking into account the size of the data and the size of the features of interest. However, future work would focus on a statistical estimation of the best parameters in the sense of maximum likelihood. Such

an estimation has been performed for images and small deformations in Allassonnière et al. (2007). This would provide an optimal decomposition of the variability. Finally, the sparsity parameter τ controls the numerical precision of the template estimation. The smaller, the more precise the computations, the more points in the representation of the template, the slowest the computation of the atlas.

The Orthogonal Matching Pursuit adapted to our framework of currents allows us to achieve very good compression ratio, while ensuring small approximation errors. However, even better compression ratio could be achieved, for instance by adding more vectors in the basis on which the signal is decomposed. In this paper, we consider the basis made of the Delta Dirac currents δ_x^α which models infinitesimal segments or normals. One way to extend this basis is to include the directional derivative of the Delta Dirac currents $\partial_u \delta_x^\alpha$ (defined by $\partial_u \delta_x^\alpha(\omega) = \alpha^t \frac{d\omega(x)}{dx} u$ for any vector u), which model torques. We could also take advantage of the recent developments of the matching pursuit techniques such as the Compressive Sampling Matching Pursuit of Needell and Tropp (2008), the subspace Matching Pursuit of Dai and Milenkovic (2008) or the stagewise Orthogonal Matching Pursuit of Donoho et al. (2006) for instance.

The approximation of currents based on matching pursuit techniques allows us to give a representation of mean and modes which is much more interpretable than their initial representation (the union of all weighted tangents or normals in the database). However, from a pure visualization point of view, this representation is not completely satisfactory, since the mean of a surface modeled as currents is represented visually by sets of small triangles, mean of a fiber tracts with sets of unconnected segments. One would like to have a better rendering of the currents. Several ideas may be followed: approximating the current with a smooth surface via minimal surfaces or using computer graphics tools like splats for instance. This would be useful if one wants to use these statistics to constraint the segmentation of anatomical structures in images via atlas to subject registration.

Our first results on real anatomical data shows the relevance of this description of the variability. Even if there were obtained with relatively small datasets, our variability analysis could be interpreted from an anatomical point of view. Of course, this interpretation should be strengthen by using larger datasets. In this case, we can expect new anatomical knowledge to come up. Last but not least, the description of the variability of population of patients may drive the search of anatomical characterization of pathologies, as illustrated by our experiments on sulcal curves, deep brain structures and white matter fiber bundles. In turn, this analysis could be used to perform discrimination between healthy and pathologic subjects, or automatic classification of patients according to their pathologies. Unsupervised clustering on the deformations and residual perturbations may also lead to the decomposition of pathologies into consistent subtypes.

7. Acknowledgments

We would like to thank P. Thompson, G. Gerig, M. Styner, D. Ducreux and P. Fillard for kindly providing us with the data. This work was partially supported the European IP project Health-e-child (IST-2004-027749) and Microsoft Research.

Appendix A. Proof of convergence of the orthogonal matching pursuit algorithm

Given a vector field $\gamma = \mathcal{L}_W^{-1}(T)$, the algorithm detailed in Section 3.2 estimates at each step n a current Π_n that is supposed to approximate the true solution T. To prove this fact, we show that the infinite norm of the residual vector field $\left\|\mathcal{L}_W^{-1}(\Pi_n) - \gamma\right\|_\infty$ as well as the W^*-norm of the residual current $\|T - \Pi_n\|_{W^*}$ $(= \left\|\gamma - \mathcal{L}_W^{-1}(\Pi_n)\right\|_W)$ tends to zero as n tends towards infinity.

We denote by $E_n \doteq \mathrm{Span}(\delta_{x_i}^\alpha; 1 \leq i \leq n, \alpha \in \mathbb{R}^d)$ for the iteratively estimated points (x_i). The estimated current Π_n is defined as the W-orthogonal projection onto E_n. We denote R_n the residue so that $\mathcal{L}_W(\gamma) = \Pi_n + R_n$. The corresponding residual error in W is denoted by $\gamma_n = \mathcal{L}_W^{-1}(R_n)$. K denotes here a generic kernel.

A.1. *Convergence for the L^∞ norm*

We assume that there is a constant $c < 0$ such that $\langle K(x,x)u, u\rangle \geq c|u|^2$ for any $x, u \in \mathbb{R}^d$. This is the case if K is translation invariant and definite positive since then $K(x,x) = K(0) > 0$. At each iteration, due to the numerical implementation, we choose $x_{n+1} = x_*$ so that

$$|\gamma_n(x_*)| \geq \max_{x \in \mathbb{R}^3} |\gamma_n(x)| / 2 \qquad \text{(A.1)}$$

instead of the point that reaches exactly the maximum.

For any $\alpha \in \mathbb{R}^d$, since $\Pi_n + \delta^{\alpha}_{x_*} \in E_{n+1}$

$$\|R_{n+1}\|^2_{W^*} \leq \left\|\Pi_n + \delta^{\alpha}_{x_*} - T\right\|^2_{W^*} \qquad \text{(A.2)}$$

Minimizing the right-hand side of Eq.A.2 with respect to α leads to $\alpha_* = K(x_*, x_*)^{-1}\gamma_n(x_*)$ for which $\left\|\Pi_n + \delta^{\alpha_*}_{x_*} - T\right\|^2_{W^*} = \left\|R_n + \delta^{\alpha_*}_{x_*}\right\|^2_{W^*} = \|R_n\|^2_{W^*} - \langle K(x_*, x_*)^{-1}\gamma_n(x_*), \gamma_n(x_*)\rangle \leq \|R_n\|^2_{W^*} - c|\gamma_n(x_*)|^2$. Thus we get from (A.1) and (A.2) that $\|R_{n+1}\|^2_{W^*} \leq \|R_n\|^2_{W^*} - c\|\gamma_n\|^2_\infty/2$. Therefore the series $\|R_n\|_{W^*}$ is monotonically decreasing and hence converges. Moreover, $\frac{c}{2}\sum_{k=1}^{\infty}\|\gamma_k\|^2_\infty \leq \|T\|^2_{W^*}$ and $\|\gamma_n\|_\infty \to 0$.

A.2. *Convergence for the W norm*

Introducing F_n such that $E_{n+1} = E_n \overset{\perp}{\oplus} F_n$ and p_n the W-orthogonal projection on F_n we have for $n \geq m$: $\|R_n - R_m\|^2_{W^*} = \|\Pi_n - \Pi_m\|^2_{W^*} = \sum_{k=m}^{n-1}\|p_k\|^2_{W^*} \leq \|T\|^2_{W^*}$. Writing $R_m = \sum_{i=1}^{m}\delta^{\alpha_i}_{x_i}$, we get $\|R_n\|^2_{W^*} \leq \langle R_n, R_m\rangle_{W^*} + \langle R_n, R_n - R_m\rangle_{W^*} \leq \sum_{i=1}^{m}|\alpha_i|\|\gamma_n\|_\infty + \|T\|^2_{W^*}\sum_{k=m}^{\infty}\|p_k\|^2_{W^*}$. Thus, $\lim\|R_n\|^2_{W^*} \leq \|T\|^2_{W^*}\sum_{k=m}^{\infty}\|p_k\|^2_{W^*}$. Since $\sum_{k=0}^{\infty}\|p_k\|^2_{W^*} < \infty$ we get the result for $m \to \infty$.

References

Allassonnière, S., Amit, Y., Trouvé, A., 2007. Towards a coherent statistical framework for dense deformable template estimation. Journal of the Royal Statistical Society Series B 69 (1), 3–29.

Allassonnière, S., Kuhn, E., 2009. Stochastic algorithm for bayesian mixture effect template estimation. ESAIM Probability and StatisticsIn Press.

Avants, B., Gee, J., 2004. Geodesic estimation for large deformation anatomical shape averaging and interpolation. NeuroImage 23, 139–150.

Charpiat, G., Faugeras, O. D., Keriven, R., 2005. Approximations of shape metrics and application to shape warping and empirical shape statistics. Foundations of Computational Mathematics 5 (1), 1–58.

Chui, H., Rangarajan, A., 2003. A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding 89 (2-3), 114–141.

Chui, H., Rangarajan, A., Zhang, J., Leonard, C., 2004. Unsupervised learning of an atlas from unlabeled point-sets. IEEE Trans. on Pattern Analysis and Machine Intelligence 26 (2), 160–172.

Cooper, D., Cootes, T., Taylor, C., Graham, J., 1995. Active shape modelsâtheir training and application. Computer Vision and Image Understanding 61, 38–59.

Cootes, T., Twining, C., Babalola, K., Taylor, C., 2008. Diffeomorphic statistical shape models. Image and Vision Computing 26 (3), 326–333.

Cremers, D., August 2006. Dynamical statistical shape priors for level set-based tracking. Transactions on Pattern Analysis and Machine Intelligence (PAMI) 28 (8), 1262–1273.

Dai, W., Milenkovic, O., 2008. Subspace pursuit for compressive sensing: Closing the gap between performance and complexity. CoRR abs/0803.0811.

Davis, G., Mallat, S., Avellaneda, M., 1997. Greedy adaptive approximations. Journal of Constructive Approximation 13 (1), 57–98.

Donoho, D. L., Drori, I., Tsaig, Y., Starck, J. L., 2006. Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit. Tech. rep.

Dupuis, P., Grenander, U., Miller, M., 1998. Variational problems on flows of diffeomorphisms for image matching. Quaterly of Applied Mathematics 56 (3), 587–600.

Durrleman, S., Fillard, P., Pennec, X., Trouvé, A., Ayache, N., 2009. A statistical model of white matter fiber bundles based on currents. In: Prince, J. L., Pham, D. L., Myers, K. J. (Eds.), Proc. of Information Processing in Medical Imaging (IPMI). LNCS. Springer.

Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., 2007. Measuring brain variability via sulcal lines registration: a diffeomorphic approach. In: Ayache, N., Ourselin, S., Maeder, A. (Eds.), Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI). Vol. 4791 of Lecture Notes in Computer Science. Springer, pp. 675–682.

Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., September 2008a. A forward model to build unbiased atlases from curves and surfaces. In: Pennec, X., Joshi, S. (Eds.), Proc. of the International Workshop on the Mathematical Foundations of Computational Anatomy (MFCA-2008).

Durrleman, S., Pennec, X., Trouvé, A., Ayache, N., 2008b. Sparse approximation of currents for statistics on curves and surfaces. In: Metaxas, D., Axel, L., Székely, G., Fichtinger, G. (Eds.), Proc. Medical Image Computing and Computer As-

sisted Intervention (MICCAI), Part II. Vol. 5242 of LNCS. Springer, pp. 390–398.

Durrleman, S., Pennec, X., Trouvé, A., Thompson, P., Ayache, N., 2008c. Inferring brain variability from diffeomorphic deformations of currents: an integrative approach. Medical Image Analysis 12/5 (12), 626–637.

Fillard, P., Arsigny, V., Pennec, X., Ayache, N., Nov. 2007a. Clinical DT-MRI estimation, smoothing and fiber tracking with log-Euclidean metrics. IEEE Transactions on Medical Imaging 26 (11), 1472–1482.

Fillard, P., Arsigny, V., Pennec, X., Hayashi, K., Thompson, P., Ayache, N., January 2007b. Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines. NeuroImage 34 (2), 639–650.

Fletcher, P., Lu, C., Pizer, S., Joshi, S., 2004. Principal geodesic analysis for the study of nonlinear statistics of shape. Transactions on Medical Imaging 23 (8), 995–1005.

Glaunès, J., 2005. Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l'anatomie numérique. Ph.D. thesis, Université Paris 13, http://cis.jhu.edu/ joan/TheseGlaunes.pdf.

Glaunès, J., Joshi, S., 2006. 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).

Grenander, U., 1994. General Pattern Theory: a Mathematical Theory of Regular Structures. Oxford University Press.

Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., Gilmore, J., Piven, J., 2005. Magnetic resonance imaging and head circumference study of brain size in autism. The Archives of General Psychiatry 62, 1366–1376.

Hufnagel, H., Pennec, X., Ehrhardt, J., Ayache, N., Handels, H., 2008. Generation of a statistical shape model with probabilistic point correspondences and em-icp. International Journal for Computer Assisted Radiology and Surgery (IJCARS) 2 (5), 265–273.

Joshi, S., Davis, B., Jomier, M., Gerig, G., 2004. Unbiased diffeomorphic atlas construction for computational anatomy. NeuroImage 23, 151–160.

Joshi, S. H., Klassen, E., Srivastava, A., Jermyn, I., 2007. A novel representation for riemannian analysis of elastic curves in r^n. In: Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society.

Leventon, M., Faugeras, O., Grimson, E., Wells, W., 06 2000. Level set based segmentation with intensity and curvature priors. In: Mathematical Methods in Biomedical Image Analysis, MMBIA 2000.

Leventon, M., Grimson, E., Faugeras, O., Wells, S., Kikinis, R., 2003. Knowledge-based segmentation of medical images. Geometric Level Set Methods in Imaging, Vision, and Graphics.

Ma, J., Miller, M., Trouvé, A., Younes, L., 2008. Bayesian template estimation in computational anatomy, to appear in NeuroImage.

Mallat, S., Zhang, Z., 1993. Matching pursuits with time-frequency dictionaries. IEEE Transactions on Signal Processing 41 (12), 3397–3415.

Marsland, S., Twining, C., 2004. Constructing diffeomorphic representations for the groupwise analysis of non-rigid registrations of medical images. IEEE Transactions on Medical Imaging 23 (8), 1006–1020.

Marsland, S., Twining, C., Taylor, C., 2008. A minimum description length objective function for groupwise non-rigid image registration. Image and Vision Computing 26 (3), 333–346.

Mio, W., Srivastava, A., Joshi, S. H., 2007. On shape of plane elastic curves. International Journal of Computer Vision 73 (3), 307–324.

Needell, D., Tropp, J. A., Apr 2008. Cosamp: Iterative signal recovery from incomplete and inaccurate samples.
URL `http://arxiv.org/abs/0803.2392`

Pardo, X. M., Leboran, V., Dosil, R., 2004. Integrating prior shape models into level-set approaches. Pattern Recognition Letters 25 (6), 631 – 639.

Pati, Y., Rezaifar, R., Krishnaprasad, P., November 1993. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: Conf. Record of the 27th Asilomar Conference on Signals, Systems and Computers. Vol. 1. pp. 40–44.

Pennec, X., Jul. 2006. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. Journal of Mathematical Imaging and Vision 25 (1), 127–154.

Pizer, S., Fletcher, P., Joshi, S., Thall, A., Chen, J., Fridman, Y., Fritsch, D., Gash, A., Glotzer, J., Jiroutek, M., Lu, C., Muller, K., Tracton, G., Yushkevich, P., Chaney, E., 2003. Deformable m-reps for 3d medical image segmentation. International Journal of Computer Vision 55 (2-3), 85–106.

Saitoh, S., 1988. Theory of Reproducing Kernels and Its Applications. Vol. 189 of Pitman Research Notes in Mathematics Series. Wiley.

Thompson, P., Schwartz, C., Lin, R., Khan, A., Toga, A., 1996. 3D statistical analysis of sulcal variability in the human brain. Journal of Neuroscience 16 (13), 4261–4274.

Toussaint, N., Souplet, J., Fillard, P., 2007. Medinria: Medical image navigation and research tool by inria. In: Proc. of MICCAI'07 Workshop on Interaction in medical image analysis and visualization. Brisbane, Australia.

Trouvé, A., 1998. Diffeomorphisms groups and pattern matching in image analysis. International Journal of Computer Vision 28, 213–221.

Twining, C. J., Cootes, T., Marsland, S., adn Roy Schestowitz, V. P., Taylor, C. J., 2005. A unified information-theoretic approach to groupwise non-rigid registration and model building. In: Information Processing in Medical Imaging (IPMI). Lecture Notes of Computer Vision, Springer-Verlag.

Vaillant, M., Glaunès, J., 2005. Surface matching via currents. In: Proceedings of Information Processing in Medical Imaging. Vol. 3565 of Lecture Notes in Computer Science. Springer, pp. 381–392.

Vaillant, M., Miller, M., Younes, L., Trouvé, A., 2004. Statistics on diffeomorphisms via tangent space representations. NeuroImage 23, 161–169.

Vaillant, M., Qiu, A., Glaunès, J., Miller, M., 2007. Diffeomorphic metric surface mapping in subregion of the superior temporal gyrus. NeuroImage 34 (3), 1149–1159.

Zollei, L., Learned-Miller, E., Grimson, E., Wells, W., 2005. Efficient population registration of data. In: Computer Vision for Biomedical Image Applications. Vol. 3765 of Lecture Notes in Computer Science. Springer, pp. 291–301.