

Sparse Approximation of Currents for Statistics on Curves and Surfaces

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

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

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

Abstract. Computing, processing, visualizing statistics on shapes like curves or surfaces is a real challenge with many applications ranging from medical image analysis to computational geometry. Modelling such geometrical primitives with currents avoids feature-based approach as well as point-correspondence method. This framework has been proved to be powerful to register brain surfaces or to measure geometrical invariants. However, if the state-of-the-art methods perform efficiently pairwise registrations, new numerical schemes are required to process groupwise statistics due to an increasing complexity when the size of the database is growing. Statistics such as mean and principal modes of a set of shapes often have a heavy and highly redundant representation. We propose therefore to find an adapted basis on which mean and principal modes have a sparse decomposition. Besides the computational improvement, this sparse representation offers a way to visualize and interpret statistics on currents. Experiments show the relevance of the approach on 34 sets of 70 sulcal lines and on 50 sets of 10 meshes of deep brain structures.

1 Introduction

There is a substantial need for shape statistics in many communities ranging from medical imaging to computer vision or computational geometry, for example to incorporate shape priors in image segmentation, to analyze geometrical or anatomical differences between groups, to classify new observations according to some characteristics, to shape recognition purpose, etc. The method we propose here precisely enables to learn a statistical shape model (mean and principal modes) from a database of curves or surfaces. Over the last years, many methods have been proposed for encoding statistical priors on curves, quite less were proposed for surfaces. Among them, level-sets for instance [1], use geometrical constraints (length, area, volume, curvature, etc.) as priors. However, it is difficult to figure out how to automatically learn such priors from typical datasets. Modelling based on Point Distribution Model [2] assumes exact or at least fuzzy correspondences between point sets. Medial axis representations [3] requires to completely specify the topology of shapes.

To define a mean shape, a “covariance” matrix and the principal modes via principal component analysis (PCA), one usually requires some metric properties

on the shape space itself. Currents were precisely introduced as an interesting alternative to do that for shape registrations [4–6]. This modelling enables to define an inner product and hence an easy to compute distance between shapes *without assuming any point correspondences* between discrete structures and without *selecting arbitrary features*. Since the space of currents is linear, one could *theoretically* compute directly standard statistics on shapes such as mean or PCA. However, there is a drawback to this mathematically well grounded method: the polynomial computational complexity in the number of points in shapes (which was not critical for pairwise registration) becomes a clear bottleneck for group-wise statistics. We propose in this paper an efficient computational framework that overcomes these limitations by providing a sparse representation of currents at any desired accuracy. Our algorithm builds on ideas from the approximation theory previously developed to decompose images in wavelet bases [7, 8]. To the very best of our knowledge, this is the first time that these signal processing techniques are used in geometric shape analysis. Experimental results clearly demonstrate the interest of our method: the deformation of a mean obtained from 3 shape instances such as in Fig.3, each with 15 000 points, which was previously taking 10 hours, is now taking about 5 minutes (using the same code as in [4, 6]). 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 offers a way to the automated learning of shape priors from typical databases, which can be used, for instance, to adapt the generic model to each particular observation for image segmentation. This also gives a way to visualize and interpret mean and principal modes of a set of anatomical structures.

In the first section, we present how shapes like curves and surfaces may be modeled geometrically as currents. We then introduce a sparse deconvolution scheme to approximate currents. In the last section we show on real anatomical data how this tool may compress shapes and may help to visualize and interpret statistics on shapes.

2 Non-parametric Representation of Shapes as Currents

In the framework of currents, a curve or a surface is characterized by the way it integrates vector fields. Their construction is based therefore on a space of square integrable vector fields W . The space of currents W^* is the set of continuous mappings from W to \mathbb{R} (see [5, 9, 10] for details). A piecewise smooth surface S in \mathbb{R}^3 like a set of smooth surfaces or meshes defines a current based on:

$$S : \omega \longrightarrow \int_S \langle \omega(x), u_x \times v_x \rangle_{\mathbb{R}^3} d\sigma(x) \quad (1)$$

where (u_x, v_x) is an orthogonal basis of the tangent plane of the surface at the point x and $n_x = u_x \times v_x$ the normal of the surface defined almost everywhere. This measures the flux of the vector field through the surface. Similarly, a piecewise smooth curve L in \mathbb{R}^3 such as a set of smooth curves or polygonal lines

is a current based on $L(\omega) = \int_L \langle \omega(x), \tau(x) \rangle_{\mathbb{R}^3} dx$ for all $\omega \in W$, where $\tau(x)$ is the tangent vector at point x defined almost everywhere. This is the integral of the vector field along the path defined by L . To ensure that these mappings are continuous (and hence $S, L \in W^*$) we need to assume that any $\omega \in W$ is bounded and that for a fixed constant $C: \|\omega\|_\infty \leq C \|\omega\|_W$. This technical condition guarantees in particular that small errors measured in W are numerically small. As a consequence, W is reproducing kernel Hilbert space (r.k.h.s.) with kernel K [10]: for every $\omega \in W$ and $\alpha \in \mathbb{R}^3$:

$$\langle \omega(x), \alpha \rangle_{\mathbb{R}^3} = \langle \omega, K(\cdot, x)\alpha \rangle_W \quad (2)$$

We choose here W as the r.k.h.s. whose kernel is Gaussian: $K(x, y) = \exp(-\|x - y\|^2 / \lambda_W^2) \text{Id}$.

In this setting, any current in W^* decomposes into an infinite sum of Dirac currents: δ_x^α that may be seen as a tangent (or normal) α entirely concentrated at point x :

$$\delta_x^\alpha(\omega) = \langle \omega(x), \alpha \rangle_{\mathbb{R}^3} \quad (3)$$

for any $\omega \in W$. A curve (resp. a surface) is decomposed into its infinite set of tangents (resp. normals). Combining Eq.3 and 2 leads to: $\delta_x^\alpha(\omega) = \langle \omega, K(x, \cdot)\alpha \rangle_W$ for any ω . This shows that we can associate to every basis element δ_x^α a **dual representation** in terms of vector field: $K(x, \cdot)\alpha$ which is simply the Gaussian convolution of the vector α . Thanks to the theory [5, 9], we can extend this result: there is a one-to-one mapping $\mathcal{L}_W : W \rightarrow W^*$ that maps any vector field ω to a current T such that $T(\omega) = \langle \mathcal{L}_W^{-1}(T), \omega \rangle_W$. With this notation, $\mathcal{L}_W(K(\cdot, x)\alpha) = \delta_x^\alpha$ showing that K is the Green function of the operator \mathcal{L}_W . δ_x^α is called the momentum associated to the vector field $K(\cdot, x)\alpha$. \mathcal{L}_W enables also to provide W^* with an inner product:

$$\langle T, T' \rangle_{W^*} = \langle \mathcal{L}_W^{-1}(T), \mathcal{L}_W^{-1}(T') \rangle_W \quad (4)$$

On basis elements, this gives: $\langle \delta_x^\alpha, \delta_y^\beta \rangle_{W^*} = \alpha^t K(x, y)\beta$.

Suppose that we have a set of N discretized shapes (polygonal lines or meshes). These shapes may be approximated in W^* by the finite set of their tangents (resp. normals) at the center of the segments (resp. the center of mass of the mesh cell). Their mean and principal modes are linear combinations of the input shapes: $T = \sum_{i=1}^N \lambda_i T_i = \sum_{k=1}^{N_T} \delta_{x_k}^{\alpha_k}$ where N_T (the total number of segments or mesh cells in the database) may be very large. This representation, if exact, is far from being optimal: it may be highly redundant at the scale λ_W . The Gaussian convolution of the N momenta: $\gamma = \mathcal{L}_W^{-1}(T)$ precisely integrates this redundancy (See Fig.1-a). The two representations (T and $\gamma = \mathcal{L}_W^{-1}(T)$) are equivalent: theoretically, the deconvolution of γ would retrieve exactly T . However, the deconvolution is an ill-posed problem and requires specific numerical schemes as the matching pursuit algorithm. This method finds iteratively adapted Dirac currents $\delta_{x_{k'}}^{\alpha_{k'}}$ such that the series $\sum_{k'} \delta_{x_{k'}}^{\alpha_{k'}}$ approximates the true solution T but with faster decreasing terms than the initial decomposition. The first terms of this series provide therefore an approximation of T with an increasing precision (see Fig.1-b,c).

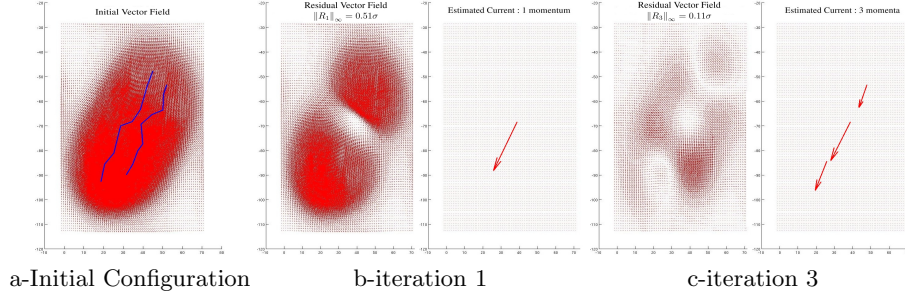


Fig. 1. A sparse deconvolution scheme: Gaussian convolution of the mean of the two blue lines with $\lambda_W = 15$ (a). First (b) and third (c) iterations: estimated momenta on the right, residual vector field on the left. σ is the standard deviation $\|L - L'\|_\infty / \sqrt{2}$. The momenta converge to the true solution while the residue tends to zero.

3 A Sparse Deconvolution Scheme

Matching Pursuit Algorithms were proposed in [11, 8] 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. Given a vector field $\gamma \in W$, we want to find an approximation of T that solves $\mathcal{L}_W^{-1}(T) = \gamma$. The goal is to find N points (x_i) and vectors (α_i) such that the current $\Pi(T) = \sum_{i=1}^N \delta_{x_i}^{\alpha_i}$ is the closest possible to T . If one knows the optimal points, $\Pi(T)$ is the orthogonal projection of T onto $\text{Span}(\delta_{x_i}^{\epsilon_k}; k = 1, 2, 3, i = 1 \dots N)$ 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 $\langle T, \delta_{x_i}^{\epsilon_k} \rangle_{W^*} = \langle \Pi(T), \delta_{x_i}^{\epsilon_k} \rangle_{W^*}$, which gives (applying \mathcal{L}_W^{-1} and Eq.4) the linear set of $3N$ equations:

$$\sum_{p=1}^N (K(x_i, x_p) \alpha_p)_k = \gamma(x_i)_k \quad (5)$$

The search for the optimal points (x_i) has been proved to be NP-hard in general [7]. The orthogonal matching pursuit algorithm is a suboptimal greedy approach to this problem: the first point x_1 is the one for which the projection of T on $\delta_x^{\epsilon_k}$ is maximal. Since $\langle \mathcal{L}_W(\gamma), \delta_x^{\epsilon_k} \rangle_{W^*} = \langle \gamma, K(\cdot, x) \epsilon_k \rangle_W = \gamma(x)_k$ (Eq.2), x_1 is a point where γ is maximal. Solving Eq.5 gives: $\alpha_1^1 = K(x_1, x_1)^{-1} \gamma(x_1)$. We then remove from γ its orthogonal projection on $\mathcal{L}_W^{-1}(\delta_{x_1}^{\epsilon_k})$ and we iterate this residue: $\gamma_1 = \gamma - K(\cdot, x_1) \alpha_1^1$. Finally, the algorithm can be then written as:

- Input: a vector field γ , a threshold $\eta > 0$
- Initialization: $\gamma_0 = \gamma$, $N = 0$
- While $\|\gamma_N\|_\infty \geq \eta$ do:
 1. $x_{N+1} = \operatorname{argmax}_{x \in \mathbb{R}^3} \|\gamma_N(x)\|_{\mathbb{R}^3}$
 2. Find $(\alpha_i^{N+1})_{1 \leq i \leq N+1}$ by solving $\sum_{p=1}^{N+1} (K(x_i, x_p) \alpha_p^{N+1})_k = \gamma(x_i)_k$
 3. $\gamma_{N+1} = \gamma - \sum_{i=1}^{N+1} K(\cdot, x_i) \alpha_i^{N+1}$ (Gaussian convolution), $N = N + 1$

After N_{mom} steps, the algorithm gives an approximation of $T = \mathcal{L}_W(\gamma)$ with N_{mom} momenta: $\Pi_{N_{\text{mom}}}(T) = \sum_{i=1}^{N_{\text{mom}}} \delta_{x_i}^{\alpha_i^{N_{\text{mom}}}}$. It has been proved [12] that this algorithm converges to the true solution, and more precisely that both $\|T - \Pi_N\|_{W^*}$ and $\|\gamma - \mathcal{L}_W^{-1}(\Pi_N(T))\|_{\infty}$ converges to 0 as N_{mom} tends toward infinity. This means, in particular, that the norm of the residue is smaller than any positive threshold in finite time, thus proving that the algorithm finishes.

From a computational point of view, we suppose that the input vector field γ is sampled is therefore stored as an image of 3D vectors. At step 1, we restrict ourselves to find the maximum of $\|\gamma(x)\|_{\mathbb{R}^3}$ on the grid's points, such that the estimated position x_i always belong to the grid. At step 3, one computes the vector field: $\mathcal{L}_W^{-1}(\Pi_N)(x) = \sum_{i=1}^N K(x, x_i) \alpha_i^N$ at grid's points. Since K is translation-invariant, this Gaussian convolution can be efficiently computed by FFT's. γ_N is also stored as an image of vectors. The output $\Pi_N(T)$ is stored as a list of (position,vector).

One wants to use this algorithm to find a sparse representation of an input current $T = \sum_{s=1}^{N_T} \delta_{y_s}^{\beta_s}$ where N_T is a priori very large. For this purpose, we computes $\gamma(x) = \mathcal{L}_W^{-1}(T)(x) = \sum_{s=1}^n K(x, x_s) \beta_s$. Since γ is band-limited, we sample it on a linearly spaced grid Λ with a step Δ such that Δ/λ_W is small (typically less than 1/5). We fix the grid such that the data are further than λ_W of the grid's borders (so that we can assume periodic boundary conditions and compute Gaussian convolution by FFT's). This sampling leads to an image of vectors used as input of the algorithm. The threshold η has to be specified for every applications. If T is a linear combination of N_{obs} currents T_1, \dots, T_N (such as the mean current, a principal mode, the 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 $\sigma^2 = \frac{1}{N-1} \sum_{i=1}^N \|T_i - \bar{T}\|^2$ and $\bar{T} = \frac{1}{N} \sum_{i=1}^N T_i$. This means that the algorithm finishes when the approximation's error is smaller than $\tau\%$ of the variance.

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) 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.

4 Numerical Experiments

Curves in 3D We use a set of 70 sulci delineated in $N_{\text{obs}} = 34$ subjects. The sulci are the fissures on the brain surface and they are often used to measure anatomical differences between subjects [13]. For each sulcal line, we approximate the mean current $\bar{L} = \frac{1}{N_{\text{obs}}} \sum_{i=1}^{N_{\text{obs}}} L_i$ for $\lambda_W = 12\text{mm}$ and $\tau = 5\%$. Results are shown in Fig.2-a for the Sylvian Fissure of the right hemisphere and 2-b for all 70 sulci. The initial number of momenta for the mean fissure was $N_T = 899$ (i.e.

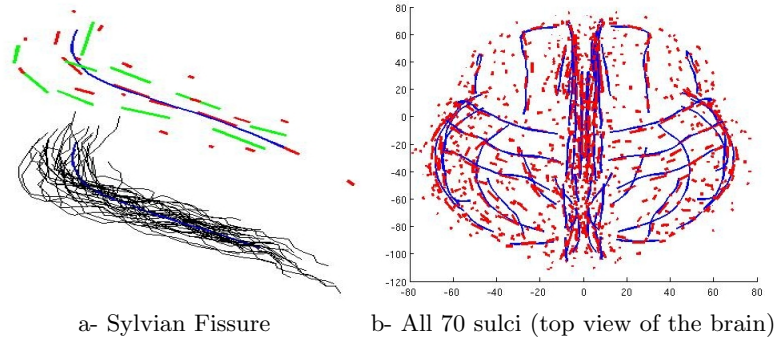


Fig. 2. Statistics for 70 sulci in 34 subjects ($\lambda_W = 12\text{mm}, \tau = 5\%$). Left: all set of lines (black), their mean (red) and first eigenmode at σ (green) showing the spreading of the lines set. Right: Mean currents (red) compared to the mean lines (blue) computed from B-spline parametrization of curves [14]. Results are in good agreement.

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 lines' grid have 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 [14].

We then compute the eigenmodes of the lines sets by PCA. We find the eigenvectors (V_k) of the $N_{obs} \times N_{obs}$ matrix $\Sigma = (\langle L_i - \bar{L}, L_j - \bar{L} \rangle_{W^*})_{i,j=1 \dots N_{obs}}$. The k^{th} eigenmode is given by the linear combination of input currents: $m_k = \sum_i (V_k)_i (L_i - \bar{L})$. We approximate the first eigenmode of the Sylvian Fissure of the right hemisphere (Fig.2-a): this mode captures the spreading of the lines set.

Surfaces We compute the mean current of 10 meshes of deep brain structures (Caudate, Putamen, Globus Pallidus, Amygdala and Hippocampus for each hemisphere) for $N_{obs} = 50$ subjects [15]. For surfaces, we represent the estimated momenta (normal of an infinitesimal mesh cell) as equilateral triangles whose normals is the momenta. Results are shown in Fig.3. The compression ratio between N_T and N_{mom} for the 10 structures is on average of $99.96\% \pm 10^{-4}$. Each structures' grid has a step $\Delta = 1\text{mm}$ and for one structure we have the following typical values: $N_T = 50 \times 3000 = 1.5e^5$, $N_{grid} = 3e^5$ and $N_{mom} = 100$. Fig.3-c shows that the quality of approximation remains good until very high compression ratio. Similarly, we compute mean from meshes of 7 controls. The difference between both means is a current that we approximate: the arrows of Fig.4-b 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.

5 Discussion and Conclusion

The method presented here enables to approximate, at any desired accuracy, the mean and principal modes of a set of curves or surfaces. On the one hand, this




Fig. 3. Approximation of the mean current for 10 structures segmented in 50 subjects (b) ($\lambda_W = 5\text{mm}$, $\tau = 5\%$ with data of diameter 60mm). Good approximation's quality can be achieved until very high compression ratio (c). Red points correspond to the approximation shown in b. for which the error equals $\tau = 5\%$ of the variance.

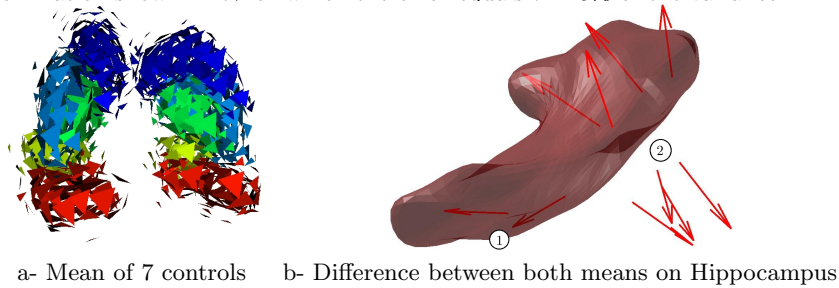


Fig. 4. a-Mean of 7 controls. b-Difference between mean of autistics and mean of controls (arrows) superimposed with the Hippocampus of a control. Mean from autistics is more curved at hippocampus' extremity (area 1) and thicker in area 2.

approximation gives a way to visualize and hence to interpret such statistics on shapes as emphasized by our results on real anatomical datasets. On the other hand, the very high compression ratio we achieve on real data offers a way to include such statistics on registration scheme for example. A registration algorithm can indeed deforms an input shape with a spatial and temporal complexity of $f(N_T)$ [4] (f is between linear and quadratic function) or solve evolution differential equations on a grid [16] with a complexity $f(N_{grid})$. In the first case our representation with N_{mom} momenta improves greatly both spatial and temporal complexity of the method while guaranteeing a fixed approximation error. This makes now possible to fit statistical surfaces to image data for segmentation purpose (deformable models). This would also be useful to adapt the atlas building scheme set up in [17] in case of images and small deformations to estimate shape prototypes. Our method could have also a significant impact in several other fields including surface rendering in computer graphics, surface reconstruction and re-meshing in computational geometry, statistics on the cortex surfaces for detecting activations from fMRI in neuroscience.

For pure visualization purposes, however, the method presented here need further improvements for a better rendering of the approximated currents. For

this purpose, we investigate how to bridge the gap with standard computer graphics methods like splats for instance.

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

References

1. Baillard, C., Hellier, P., Barillot, C.: Segmentation of brain 3D MR images using level sets and dense registration. *Medical Image Analysis* **5**(3) (2001) 185–194
2. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding* **89**(2-3) (2003) 114–141
3. Gorczowski, K., Styner, M., Jeong, J.Y., Marron, J.S., Piven, J., Hazlett, H.C., Pizer, S.M., Gerig, G.: Statistical shape analysis of multi-object complexes. In: *Computer Vision and Pattern Recognition CVPR, IEEE* (2007) 1–8
4. Vaillant, M., Glaunès, J.: Surface matching via currents. In: *Proc. of Information Processing in Medical Imaging. Volume 3565 of LNCS., Springer* (2005) 381–392
5. Glaunès, J.: Transport par difféomorphismes de points, de mesures et de courants pour la comparaison de formes et l’anatomie numérique. PhD thesis (2005)
6. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Measuring brain variability via sulcal lines registration: a diffeomorphic approach. In: *Proc. MICCAI. Volume 4791 of LNCS., Springer* (2007) 675–682
7. Davis, G., Mallat, S., Avellaneda, M.: Greedy adaptive approximations. *Journal of Constructive Approximation* **13**(1) (1997) 57–98
8. Pati, Y., R.Rezaifar, Krishnaprasad, P.: Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In: *27th Asilomar Conference on Signals, Systems and Computers. Volume 1.* (1993) 40–44
9. Schwartz, L.: Sous espaces hilbertiens d’espaces vectoriels topologiques et noyaux associés (noyaux reproduisants). *J. Analyse Math.* **13** (1964) 115–256
10. Aronszajn, N.: Theory of reproducing kernels. *Transactions of the American Mathematical Society* (68) (1950) 337–404
11. Mallat, S., Zhang, Z.: Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing* **41**(12) (1993) 3397–3415
12. Durrleman, S., Pennec, X., Trouvé, A., Ayache, N.: Sparse approximation of currents for statistics on curves and surfaces. Technical report, INRIA (2008)
13. Thompson, P., Schwartz, C., Lin, R., Khan, A., Toga, A.: 3D statistical analysis of sulcal variability in the human brain. *J. of Neuroscience* **16**(13) (1996) 4261–4274
14. Fillard, P., Arsigny, V., Pennec, X., Hayashi, K., Thompson, P., Ayache, N.: Measuring brain variability by extrapolating sparse tensor fields measured on sulcal lines. *NeuroImage* **34**(2) (January 2007) 639–650
15. Hazlett, H., Poe, M., Gerig, G., Smith, R., Provenzale, J., Ross, A., Gilmore, J., Piven, J.: Magnetic resonance imaging and head circumference study of brain size in autism. *The Archives of General Psychiatry* **62** (2005) 1366–1376
16. Aubert, G., Kornprobst, P.: Mathematical problems in image processing - Partial differential equations and the calculus of variations. Volume 147 of *Applied Mathematical Sciences.* Springer (2001)
17. Allasonnière, S., Amit, Y., Trouvé, A.: Towards a coherent statistical framework for dense deformable template estimation. *J. Roy. Stat. Soc.-B* **69**(1) (2007) 3–29