

Measuring Brain Variability by Extrapolating Sparse Tensor Fields Measured on Sulcal Lines^{*}

Pierre Fillard¹, Vincent Arsigny¹, Xavier Pennec¹, Kiralee M. Hayashi², Paul M. Thompson²,
and Nicholas Ayache¹

¹ INRIA Sophia Antipolis - EPIDAURE/ASCLEPIOS Project
2004 Route des Lucioles BP 93
06902 Sophia Antipolis Cedex, France

{Pierre.Fillard, Vincent.Arsigny, Xavier.Pennec, Nicholas.Ayache}@sophia.inria.fr

² Laboratory of Neuro Imaging, Dept. of Neurology, UCLA School of Medicine,
225E Neuroscience Research Building,
Los Angeles, CA, USA
{khayashi, thompson}@loni.ucla.edu

Abstract. Modeling and understanding the variability of brain structures is a fundamental problem in the neurosciences. Improved mathematical representations of structural brain variation are needed to help detect and understand genetic or disease related sources of abnormality, as well as to improve statistical power when integrating functional brain mapping data across subjects. In this paper, we develop a new mathematical model of normal brain variation based on a large set of cortical sulcal landmarks (72 per brain) delineated in each of 98 healthy human subjects scanned with 3D MRI (age: 51.8 +/- 6.2 years). We propose an original method to compute an average representation of the sulcal curves, which constitutes the mean anatomy. After affine alignment of the individual data across subjects, the second order moment distribution of the sulcal position is modeled as a sparse field of covariance tensors (symmetric, positive definite matrices). To extrapolate this information to the full brain, one has to overcome the limitations of the standard Euclidean matrix calculus. We propose an affine-invariant Riemannian framework to perform computations with tensors. In particular, we generalize radial basis function (RBF) interpolation and harmonic diffusion partial differential equations (PDEs) to tensor fields. As a result, we obtain a dense 3D variability map which agrees well with prior results on smaller subject samples. Moreover, "leave one (sulcus) out" tests show that our model is globally able to recover the missing information on brain variation when there is a consistent neighboring pattern of variability. Finally, we propose an innovative method to analyze the asymmetry of brain variability. As expected, the greatest asymmetries are found in regions that includes the primary language areas. Interestingly, any such asymmetries in anatomical variance, if it remains after anatomical normalization, could explain why there may be greater power to detect group activation in one hemisphere versus the other in fMRI studies.

^{*} This is a preprint version of the paper to appear in NeuroImage, 2006.

1 Introduction

Brain structures differ greatly in shape and size even among normal subjects, and these variations make it difficult to identify abnormal differences due to disease. Understanding the degree and quality of brain variation is vital for distinguishing signs of disease from normal variations. Neuroscientists are interested in identifying the causes of brain variability at a genetic or environmental level. An efficient, parsimonious model of the complex patterns of brain variation would help in identifying factors that contribute to it. Measuring brain asymmetry (i.e. differences between hemispheres) is of special interest as it sheds light on how the functions of the two hemispheres become specialized [Toga and Thompson, 2003]. Improved modeling of the range of variations in brain structure could make it easier to isolate specific effects of genetic polymorphisms on these normal variations and asymmetries [Cannon et al., 2005, Geschwind et al., 2002, Thompson et al., 2001a]. Finally, geometric variability of anatomy also makes the automated segmentation and labeling of brain structures difficult. Statistical information on brain variability would make this task easier [Pitiot et al., 2004, Fischl et al., 2002], and could be used in Bayesian approaches for inter-subject nonlinear registration ([Mangin et al., 2004a, Gee and R.K. Bajcsy, 1998, Ashburner and Friston, 2005]) as well (which adjust for anatomical variations across subjects prior to group analysis of brain function or metabolism). In this paper, we focus only on understanding and building a model of brain variability. Applications of this model for inter-subject registration will be the topic of a future study.

A major class of anatomical variations can be thought of as arising from the smooth deformation of a reference anatomy, where the deformation is represented as a 3D displacement field, after affine differences are factored out. Ideally, one could model the joint variability of all pairs of points to see how the displacement of one any point in a specific subject with respect to the reference anatomy covaries with the displacement of neighboring or distant points in the brain (e.g. symmetric ones in the opposite hemisphere). In this article, we simply model the variability of each anatomical point independently. Assuming that the mean deformation of the reference anatomy is null, the first moment of the 3D displacement distribution is its covariance matrix, which will be called a *variability tensor*. Thus, our goal is to compute the field of variability tensors within the brain from information that may be sparsely distributed. The reason that tensor representations are used, rather than simple scalar fields, is that variation may not be the same in all directions (there is some evidence that structural variation is greatest along certain preferred directions [Thompson et al., 2001b, Kindlmann, 2004]).

However, working with tensors is not so easy as manipulating scalar fields, as the underlying space is a manifold that is not a standard Euclidean space. As tensors constitute a convex half-cone in the vector space of matrices, many operations (like computing the mean) are stable. Nonetheless, the Euclidean framework is not satisfactory as one can easily reach the boundary of the space (singular symmetric matrices) with a classical gradient descent. Moreover, the arithmetic mean of a tensor and its inverse is not the identity matrix. This lack of symmetry is unsatisfactory: in many cases, one would like the mean to be geometric [Woods, 2003].

In section 2 we present a consistent Riemannian framework to compute with tensors. Then, we show how to extend these tools to implement harmonic diffusion partial differential equations (PDEs) and extrapolate tensors that are sparsely distributed in space. Solving these PDEs is computer intensive, so we provide a practical but efficient initialization by extending the radial basis functions (RBF) concept to tensors. In section 3, we consider low dimensional but anatomically readily defined and delineated features (sulcal landmark curves) as a way to obtain meaningful brain variability tensors. We show how to compute the mean sulcal curve and its correspondence with the sulcal instances of each subject. To extract only the relevant information and minimize the number of parameters, we fit a parametric tensor model to these data. Then,

we come back to our original goal by extrapolating this sparse tensor model to the whole brain, and the validity of our extrapolated model is analyzed. Finally, in section 4, we generalize our statistical model to examine the correlation and asymmetry of the observed variations at symmetric points in the two brain hemispheres.

2 A Mathematical Framework to Extrapolate Tensors

Most of the literature addresses tensor computing problems in the context of diffusion tensor image (DTI) regularization. In these articles, the spectral decomposition of the tensors is exploited. For instance, [Coulon et al., 2004] anisotropically restore the principal direction of the tensors, while [Tschumperlé and Deriche, 2002] independently restore the eigenvalues and eigenvectors. This last approach requires an additional step to perform re-orientation of the eigenvectors due to the non-uniqueness of the decomposition.

More recently, differential geometry approaches have been developed to generalize principal components analysis (PCA) to tensor-valued data [Fletcher and Joshi, 2004], for statistical segmentation of tensor images [Lenglet et al., 2006], for computing a geometric mean and an intrinsic anisotropy index [Batchelor et al., 2005], for preliminary results on brain variability modeling [Fillard et al., 2005], or as the basis for a full framework for Riemannian tensor calculus [Pennec et al., 2006]. In [Pennec et al., 2006], we endowed the space of tensors with an affine-invariant Riemannian metric to obtain results that are independent of the choice of the spatial coordinate system. In fact, this metric had already been proposed in statistics [Skovgaard, 1984], and turns out to be the basis of all the previous differential geometric approaches. Other Riemannian strategies are possible, like Log-Euclidean metrics [Arsigny et al., 2006b].

2.1 A Riemannian Framework for Tensor Calculus

The invariant metric provides a new framework to overcome the limitations of Euclidean calculus: it endows the tensor space with a highly regular structure, in which matrices with null or negative eigenvalues are at an infinite distance from any positive definite matrix. Moreover, the geodesic path between any two tensors is uniquely defined, leading to interesting properties such as the existence and uniqueness of the (geometric) mean [Pennec et al., 2006].

On Riemannian manifolds, geodesics realize a local diffeomorphism, called the exponential map, from the tangent space at a given point to the manifold itself. This allows us to (locally) identify points of the manifold with tangent vectors (see Fig. 1). With the invariant metric on tensors, the geodesic starting at Σ and with tangent vector W can be expressed simply with the classical matrix exponential and the (Riemannian) exponential map realizes a global diffeomorphism [Pennec et al., 2006]:

$$\exp_{\Sigma}(W) = \Sigma^{\frac{1}{2}} \exp\left(\Sigma^{-\frac{1}{2}} W \Sigma^{-\frac{1}{2}}\right) \Sigma^{\frac{1}{2}} \quad \text{and} \quad \log_{\Sigma}(\Lambda) = \Sigma^{\frac{1}{2}} \log\left(\Sigma^{-\frac{1}{2}} \Lambda \Sigma^{-\frac{1}{2}}\right) \Sigma^{\frac{1}{2}}.$$

These two diffeomorphisms are the key to the numerical implementation and generalization to manifolds of numerous algorithms that work on a vector space. For instance, the “difference vector” between two tensors Σ_1 and Σ_2 may be expressed in the tangent space at the identity as:

$$Z = \Sigma_1^{-1/2} \log_{\Sigma_1}(\Sigma_2) \Sigma_1^{-1/2} = \log(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2}). \quad (1)$$

Note that Z is a symmetric but not necessarily positive matrix. The distance between the two tensors is simply given by:

$$\text{dist}^2(\Sigma_1, \Sigma_2) = \|Z\|_2^2 = \text{trace} \left(\left[\log \left(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2} \right) \right]^2 \right). \quad (2)$$

Likewise, the Euclidean gradient descent scheme $\Sigma_{t+1} = \Sigma_t - \varepsilon \nabla C(\Sigma_t)$, which could easily lead to a non-positive matrix, is advantageously replaced by the *geodesic marching scheme* $\Sigma_{t+1} = \exp_{\Sigma_t}(-\varepsilon \nabla C(\Sigma_t))$.




Fig. 1. Geodesic and tangent space in a Riemannian manifold. The geodesic joining x and y is a straight line in the tangent space $T_x M$. The distance is conserved, i.e. the length of the geodesic is equal to the norm of the vector xy . The operation that turns the geodesic into a straight line is called the **logarithmic map**. The inverse mapping is called the **exponential map**.

2.2 Dense Extrapolation of Sparse Tensors

Let us consider a set of N measures Σ_i of a tensor field $\Sigma(x)$ at spatial positions $x_i \in \mathbb{R}^d$. To access the value of the tensor field at any point, one could think of interpolating or approximating these measures. We proposed in [Pennec et al., 2006] a least-squares attachment term to the sparsely distributed tensors, combined with a regularization term to perform an estimation of the extrapolated tensor: $C(\Sigma) = \text{Sim}(\Sigma) + \text{Reg}(\Sigma)$. In a continuous setting, the data attachment term is:

$$\text{Sim}(\Sigma) = \frac{1}{2} \sum_{i=1}^N \text{dist}^2(\Sigma(x_i), \Sigma_i) = \frac{1}{2} \int_{\Omega} \sum_{i=1}^N \text{dist}^2(\Sigma(x), \Sigma_i) \delta(x - x_i) dx.$$

The Dirac distributions $\delta(x - x_i)$ are problematic when numerically differentiating the criterion. To regularize the problem, we consider them as the limit of a Gaussian function G_σ when σ goes to zero. Practically, σ has to be of the order of the spatial resolution of the grid on which $\Sigma(x)$ is estimated, so that each measure influences its immediate neighborhood. After differentiating the criterion, one obtains: $\nabla \text{Sim}_\sigma(x) = -\sum_i G_\sigma(x - x_i) \log_{\Sigma(x)}(\Sigma_i)$.

Basically, the attachment term prevents the tensor field from deviating too much from the measures at the points x_i . In between these points, we need to add a regularization term that ensures a smooth, relatively homogeneous interpolation. The simplest criterion is the harmonic regularization: $\text{Reg}(\Sigma) = \frac{1}{2} \int_{\Omega} \|\nabla \Sigma(x)\|_{\Sigma}^2$. We showed in [Pennec et al., 2006] that the gradient of this criterion is $\nabla \text{Reg}(\Sigma)(x) = -\Delta \Sigma(x)$, and we provided a practical implementation of this Laplace-Beltrami operator on a tensor field. Using the geodesic marching scheme, we compute at each point x of our estimation grid the following intrinsic gradient descent:

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)}(-\varepsilon \nabla \text{Sim}(x) - \varepsilon \nabla \text{Reg}(x)). \quad (3)$$

Finally, we can evaluate the extrapolated field Σ at any point x by tri-linear interpolation of the values at the grid nodes.

However, due to the large number of tensors and the large domain of diffusion used here (see next section), this algorithm converges slowly, even with a multi-resolution implementation. To improve the initialization and enable faster convergence, in this article we develop an RBF interpolation.

2.3 Extending RBFs to Extrapolate Tensors

RBFs provide a family of methods to extrapolate sparsely defined observations [Sun, 1995]. The extrapolated field is expressed as a linear combination of translated versions of a single radial function (the *basis*). Thus, if (y_i) is a set of scalar measures of the field $y(x)$ at points x_i , we find a set of scalar coefficients (λ_i) such that $y(x) = \sum_i \lambda_i h(x - x_i)$. To interpolate the data, the coefficients need to yield $y(x_i) = y_i$, i.e. be solutions of the linear system $\forall j : y_j = \sum_i \lambda_i h(x_j - x_i)$. There is a unique solution for any set of measurements at any set of spatial positions if the symmetric matrix $[H]_{i,j} = h(x_i - x_j)$ is always *positive definite*.

Scalar RBF extrapolation can be extended to vectors by simply running the extrapolation on each component independently. To apply this method to tensors, we map all tensors into the tangent space $T_\Sigma M$ of a reference tensor Σ . We then run the RBF extrapolation on the vectors $\log_\Sigma(\Sigma_i)$ and map the resulting values back into tensor space by the inverse mapping \exp_Σ . Among the many possible choices for a common reference tensor, we chose the mean $\bar{\Sigma}$ of all tensor measurements. Also, rather than letting the extrapolated values explode at infinity as with Thin Plate Splines, we use an interpolating function that decreases toward zero at infinity, namely from the family $h(x) = 1 / (1 + (\|x\|^2 / \alpha^2)^\gamma)$. The asymptotic value for the interpolation will be the reference tensor $\bar{\Sigma}$. However, this choice does not influence the result, because we are interested in values close (in distance) to the measures Σ_i . All our experiments with other references gave identical results. Actually, the real motivation of this choice is to remain mathematically consistent: the mean tensor is a good candidate for the reference, as it has globally the same size than the measures.

2.4 A Log-Euclidean Framework for Tensor Calculus

Major advances have been made recently in tensor calculus. Log-Euclidean (LE) metrics [Arsigny et al., 2006b] are a novel family of Riemannian metrics, which combine the benefits of the affine-invariant property with the lower computational cost of the Euclidean calculus. The price to pay for that is rather cheap: LE metrics are only similitude-invariant (i.e. invariant under rigid body transforms, including translations and rotations). They work on the tensor logarithm, which turns the tensor space into a vector space. The expression of the L_2 LE metric is simply:

$$\text{dist}^2(\Sigma_1, \Sigma_2) = \text{trace} \left((\log(\Sigma_1) - \log(\Sigma_2))^2 \right).$$

Surprisingly, results with this metric are very similar to the affine-invariant family, as shown in [Arsigny et al., 2006b] (for a more in-depth mathematical description of LE metrics, please refer to [Arsigny et al., 2006a]). The results are so close that it is almost impossible to quantify the difference. Thus, we adapted our diffusion PDE and the RBF concept to LE metrics and found no significant difference. However, we chose to keep the affine-invariant family of Riemannian metrics in this work, to generate results that are as invariant as possible. A computational trick consists in running the extrapolation process using Log-Euclidean metrics until convergence, and then to perform a few iterations of the affine-invariant version to obtain an affine-invariant result.

3 Modeling Brain Variability from Sulcal Lines

One initial idea for how to measure inter-subject brain variability might be to gather statistics on 3D displacement fields computed between a reference anatomy and many individuals. Such data could be obtained using an inter-subject registration algorithm. However, we would need to estimate the influence of the chosen registration method. Moreover, the influence of the target chosen for registration may be important (see Sec. 3.2). Nevertheless, the image intensity only constraints the registration in the direction of the image gradient. In other words, the matching is generally uncertain in directions parallel to edges. Only by adding the regularization, one can obtain a smooth deformation field and help recovering the missing displacement components. To provide information that is completely independent of volumetric image registration algorithms, we chose in this work to rely on lower dimensional structures, such as cortical landmarks identified by expert neuroscientists following a formalized protocol, with known inter- and intra-rater reliability. We could have taken surfaces, e.g. the hippocampus, corpus callosum, or even the whole cortex as in [Thompson et al., 2000]. However, even for surface-based deformations, we would still face the problem of finding appropriate correspondences between the surfaces (this is currently an active area of research [Pitiot et al., 2003], [Wang et al., 2005]). Moreover, these surfaces may be difficult to extract accurately. Thus, we chose to focus on anatomically well defined 3D curves that could be manually delineated by neuroanatomists and considered as ground truth data. This choice naturally led us to the primary anatomical landmarks on the cortex: the sulci. A large number of of sulcal landmarks consistently appear in all normal individuals and allow a consistent subdivision of the cortex into major lobes and gyri [Mangin et al., 2004b]. In the absence of individual functional imaging data, sulci also provide an approximate guide to the functional subdivisions of the cortex, for all of the lobes. They are also used to guide intersubject functional activations registration [Corouge et al., 2003].

We use a dataset of sulcal lines manually delineated in 98 subjects by expert neuroanatomists according to a precise protocol³ (an example of tracing is shown in Fig. 2). The dataset consisted of 47 men and 53 women (age: 51.8 \pm 6.2 years), all normal controls. The lines are traced in 3D on the cortical surface, using an interface that allows curves to be traced interactively on surfaces. The sulcal lines appear to be traced out as curved lines in the intrasulcal CSF, near the exterior of the cortex. This is because the cortical surfaces we use are regularized and they clip across the sulcal CSF. The cortical surfaces are extracted with an active surface algorithm [MacDonald et al., 1998] that creates a surface whose geometry is well adapted to a single intensity isovalue in the image while at the same time minimizing a curvature penalty that prevents complete penetration of the surface into the sulci. As such the sulcal lines are more like the gyral divisions on the exterior cortical hull as seen in the classical cytoarchitectonic maps of Brodmann. By using these as landmarks in our prior studies for the normalization of fMRI and other cortical signals [Thompson et al., 2004, Sowell et al., 2004], we have improved the registration of gyral crests although it has to be admitted that the surface does not represent the full depths of the cortical sulci. In the following, we abusively call these sulcal lines *sulci* to simplify the description. We included the maximal subset of all sulcal curves that consistently appear in all normal subjects, 72 in total. Image analysts blind to subject sex and age traced the sulci on the lateral brain surface (including the Sylvian fissure, and central, pre-central, post-central, superior temporal sulcus (STS) main body, STS ascending branch, STS posterior branch, primary intermediate sulcus, secondary intermediate sulcus, inferior temporal, superior frontal, inferior frontal, intraparietal, transverse occipital, olfactory, occipito-temporal, and collateral sulci) in each hemisphere on the surface rendering of each subject’s brain. An additional set of sulci were outlined on each interhemispheric surface (including the callosal sulcus, inferior

³ http://www.loni.ucla.edu/~khayashi/Public/medial_surface/




Fig. 2. Example of sulcal lines drawn on the cortical surface.

callosal outline, superior rostral sulcus, inferior rostral sulcus, paracentral sulcus, anterior and posterior segments of the cingulate sulcus, outer segment double parallel cingulate sulcus when present, parieto-occipital sulcus, anterior and posterior segments of the calcarine sulcus, and the subparietal sulcus). In addition to contouring the major sulci, a set of 6 midline landmark curves bordering the longitudinal fissure were outlined in each hemisphere to establish hemispheric gyral limits. Spatially registered gray-scale image volumes in coronal, axial, and sagittal planes were available simultaneously to help disambiguate brain anatomy. We have developed detailed criteria for delineating the cortical lines, and for starting and stopping points for each sulcus using brain surface atlases as references. By repeated training on test sets of brain images, the maximum allowed inter- and intra-rater error (reliability) was ensured to be better than 2mm everywhere, in terms of r.m.s. distance, and in most regions less than 1mm, far less than the intersubject anatomical variance. MR images used for delineations were linearly aligned to the ICBM stereotactic space [Collins et al., 1995], thus providing a common coordinate system for all traced curves. Next, we determined the mean curve for each sulcal line by modeling samples as deformations of a single average curve. Based on the mean sulcal line, for each sulcus, and the mapping from this curve to its instance in each subject image, we can easily compute the local covariance matrix to create our second order statistical model of the sulcal line. Other works related to the statistics of sulcal lines include [Goualher et al., 2000] and [Barillot et al., 1998].

3.1 Learning Local Variability from a Sulcal Lines Dataset

Statistical models have frequently been constructed for objects such as open or closed curves and surfaces [Cootes et al., 1995, Trouve and Younes, 2000, Paulsen and Hilger, 2003]. In each of these examples, the *aperture problem occurs*: as we do not have point landmarks, the point-to-point correspondences between instances of a surface or a curve cannot be recovered exactly. For instance, the correspondences of two instances of a sulcus are intrinsically subject to error, with greater tangential than normal uncertainty. Here, we propose a one-to-one correspondence mapping that minimizes the influence of this error.

First, we denoise the sample lines by approximating them with B-splines: the manual sampling of 3-dimensional curves is only precise up to the voxel size (about $1mm^3$), which is lower than the inter-rater variability of 2mm. In this continuous setting, the number of degrees of

freedom can be adjusted to increase robustness to noise while avoiding resampling problems [Baumberg and Hogg, 1994]. Typically, we reduce the number of control points to one third of the original sampling points, with a mean distance of 0.25mm and a maximum error of 2.7mm.

Many criteria have been proposed in the literature to evaluate the mean curve for a set of curves and to assess the appropriateness of one-to-one correspondences between geometric objects. They usually invoke local differential characteristics such as the tangent space, curvature, the local Frenet frame for a curve on a surface

[Guéziec and Ayache, 1994, Bakircioglu et al., 1998], regional shape information [Pitiot et al., 2003].

In our case, the variability is so large (see Fig. 3), that using such refined measures is difficult. In general, sulcal curves do not have internal geometric features, along their length, that occur consistently from subject to subject. Therefore, we simply use the total variance of curve models as a criterion:

$$C(z) = \frac{1}{N-1} \sum_{k=1}^N \int_0^1 \|y^k(\phi^k(s)) - z(s)\|^2 \|z'(s)\| ds, \quad (4)$$

where y^k is the given sulcus of subject k , z the mean sulcus and ϕ^k the correspondence function between the subject's curve and the mean curve and s the curve length.

Minimizing this variance greatly reduces the variability due to inadequate correspondences. Practically, we alternately improve the correspondences between the mean curve and each sample by dynamic programming and optimize the average curve position by a first-order gradient descent with an adaptive step. This optimization strategy converges toward the mean curve after a few iterations (dashed curve in Fig. 4).

For each of the 72 sulci, we end up with the mean curve $z(s)$, and one-to-one mappings $\phi^k(s)$ that give the corresponding position $y^k(\phi^k(s))$ in each subject k . The variability tensor $\Sigma(s)$ is given by:

$$\Sigma(s) = \frac{1}{N-1} \sum_{k=1}^N \left[y^k(\phi^k(s)) - z(s) \right] \left[y^k(\phi^k(s)) - z(s) \right]^\top. \quad (5)$$

Results of covariance tensors estimated along the 72 sulci are shown in Fig. 3. Variability is greater at the extremities of the curves. These points are landmarks identified by neuro-anatomists. However, we suspect that the main part of their variability is due to a bias when estimating the position of the end points of the mean curve. To remain consistent, we chose in this paper to remove the variability information at the extremities of the sulci from our model, and focus only on the inner parts of the sulci.

3.2 Estimation of the Affine Transform from Correspondences

Initially, images were affinely registered onto a common reference image (in our case the ICBM305 space). To remove the influence of the chosen target and build unbiased atlases, one often needs to use more elaborate strategies, like in [Guimond et al., 2000], or [Kochunov et al., 2002]. With our curves, one can simply rely for each subject on the established correspondences between its sulci and the mean curves to refine the affine transform. This will also further reduce overall sulcal variability.

Let (A^k, t^k) be the affine transform of subject k . Let us assume that the mean curves and the mappings between each subject's curve and the mean are known. We are looking for the optimal affine transform $(A_{(opt)}^k, t_{(opt)}^k)$ in the least-squares sense, i.e. the one that minimizes the sum of the squared differences between the transformed sulci and the mean curves:

$$C(A^k, t^k) = \sum_{i=1}^N \int_0^1 \|A^k y_i^k(\phi_i^k(s)) + t^k - z_i(s)\|^2 ds \quad (6)$$




Fig. 3. Sulcal variability. **Top:** The sulcal lines alone. Mean curves are shown in red, and traces from 98 healthy normal individuals are displayed in green and yellow. **Bottom:** Covariance matrices (ellipsoids at one σ) are overlapped at regularly sampled spatial positions along the mean sulci. The color codes for the trace: Note that tensors close to the sulci extremities are larger.

To optimize this energy, we first fix A^k and look for the optimal translation. We find that the minimum value is reached for:

$$t_{(opt)}^k = \bar{z} - A^k \bar{y}^k, \quad (7)$$

where $\bar{z} = 1/N \sum_{i=1}^N \int_0^1 z_i(s) ds$ and $\bar{y}^k = 1/N \sum_{i=1}^N \int_0^1 y_i^k(\phi_i^k(s)) ds$. Second, one introduces $t_{(opt)}^k$ back into Eq. 6. Calling $\tilde{y}_i^k = y_i^k - \bar{y}^k$ and $\tilde{z}_i = z_i - \bar{z}$, one can rewrite Eq. 6 as:

$$\begin{aligned} C(A^k) &= \sum_{i=1}^N \int_0^1 \|A^k \tilde{y}_i^k(s) - \tilde{z}_i(s)\|^2 ds \\ &= \sum_{i=1}^N \int_0^1 \text{Trace} \left[\left(A^k \tilde{y}_i^k(s) - \tilde{z}_i(s) \right) \left(A^k \tilde{y}_i^k(s) - \tilde{z}_i(s) \right)^\top \right] ds \end{aligned}$$

Using the property that $\text{Trace}(AB) = \text{Trace}(BA)$, we find that:

$$A_{(opt)}^k = \left(\sum_{i=1}^N \int_0^1 \tilde{z}_i(s) \tilde{y}_i^k(s)^\top ds \right) \left(\sum_{i=1}^N \int_0^1 \tilde{y}_i^k(s) \tilde{y}_i^k(s)^\top ds \right)^{-1}. \quad (8)$$

The overall framework now consists in alternating the positioning of the mean curve, the computation of the correspondence functions, and the evaluation of the optimal affine transform for each subject using Eq. 7 and 8. Curves in Fig. 4 represent the energy 4 averaged over the 72 sulci and for each iteration. The optimization of mean curves and correspondences reduces the amount of variability to 70% of the initial value. Reestimating the affine transform further reduces the amount of variability to 60% of its initial value.




Fig. 4. Mean curves calculation: Energy vs. iterations. Dashed curve: Without affine refinement. **Solid curve:** With affine refinement every 10 iterations. Note that the process has converged after 40 iterations.

Figure 5 shows the resulting variability tensors computed using Eq. 5 and overlapped along the mean sulcus. One notices that the correspondence optimization greatly reduces the tangential components of the variability, thus minimizing the aperture problem. The affine correction reduces the variability more globally.

3.3 Model Simplification using Tensor Interpolation

In the interior part of the sulci, the tensors are highly regular in size and shape. Some of this information is therefore redundant and could be simplified by selecting only a few tensors at




Fig. 5. Variability tensors for the Sylvian Fissure. **Top:** Without correspondence optimization. **Middle:** After correspondence optimization. **Bottom:** After correspondence optimization and affine transform optimization. Note that the curves get slightly more concentrated around the mean, showing less variability. This is reflected by the reduction of covariance matrices in the direction orthogonal to the mean curve. The reduction of variation in the tangential direction is mainly due to the optimization of correspondences between curves.

specific points along the mean sulcus, and interpolating in between them. Moreover, to avoid corrections for multiple testing problem, we had to downsample the representation of the continuous curves to a point where each tensor is independent of its neighbors. A second computational

reason is that the tensor extrapolation process is very time consuming, and reducing the number of seed tensors decreases the computational time drastically. We use geodesic interpolation to interpolate between successive tensors, because it preserves the monotone (i.e., consistently increasing or decreasing) evolution of the determinant. This is crucial when interpolating two probability densities and is not possible in general with direct interpolation. For efficiency reasons, we also selected the tensor values among the observed data rather than optimizing them as free parameters. This operation has been automated in an algorithm called *tensor picking*.

Let $(\Sigma_i)_{1 \leq i \leq N}$ be a set of N variability tensors defined at abscissa s_i along a mean sulcus. The geodesic interpolation between 2 successive tensors is given by: $\tilde{\Sigma}(s) = \exp_{\Sigma_i}[(s-s_i)/(s_{i+1}-s_i) \log_{\Sigma_i}(\Sigma_{i+1})]$ for $s_i \leq s < s_{i+1}$. As we are working only on the interior of the sulcus, s takes its values between s_2 and $s_N - 1$ (we remove the extremities), so that the *interpolated variability* $\tilde{\Sigma}(s)$ is always defined. The tensor picking operation consists of finding the optimal set (Σ_i) such that the least-square error between the observed and interpolated variability tensors is minimized: $C(\Sigma_1, \dots, \Sigma_N) = (\int_{s_2}^{s_N-1} \text{dist}^2(\Sigma(s), \tilde{\Sigma}(s)) ds)^{1/2}$. To minimize this energy, we first set N equal to 2. Then, an exhaustive search for the optimal set of N tensors is done. If the energy obtained is below a threshold t (0.2 in our experiments), the tensors are picked. Otherwise the number N is increased and the search is reiterated.

Results of this operation are presented in Fig. 6 (middle panel): by choosing tensors at adequate positions, one can qualitatively reconstruct the full variability of each sulcus using 4 to 10 covariance matrices, depending on its length and shape. Results of the quantification of such reconstruction for 3 major sulci are presented in Table 1 (interpolation line). The variability of all the major sulci can be adequately represented by 366 variability tensors out of 2000 initially.

3.4 Extrapolating the Variability to the Full Brain

The next step consists of extrapolating these selected tensors to the full brain, using the framework developed in Sec. 2.2. There are some assumptions if this is done necessarily we have chosen certain consistent cortical features to model variability, and if a greater number of structures were chosen (including subcortical structures such as the ventricles), the variation may be greater, and the spatial correlation in the variations may be slightly reduced. However, a comprehensive set of sulci was chosen, perhaps including all those that consistently occur in normal subjects. Fig. 6 presents the result extrapolating our 366 tensors on a discrete grid of size $91 \times 109 \times 91$ and with a spacing of $2 \times 2 \times 2 \text{ mm}^3$ (ICBM 305 space). We used the parameter values $\alpha = 20$ and $\gamma = 0.95$ for the RBF interpolation and $\sigma = 2$ for the discretization of the data attachment term in the extrapolation (Eq. (3)). We derived 2 scalar measures from the extrapolation. The first one (Fig. 7 left column) is a variability map given by the 3D root mean square (r.m.s.) of the trace of each covariance matrix: $rms = \sqrt{\text{trace}(\Sigma(x))}$. One can see highly variable regions (such as the parietal cortex and Broca's area) with hot colors, and more stable areas (such as the primary sensorimotor cortex) with cold colors. Note that the maximum allowed inter and intra-rater error for the manual tracings of the curves was ensured to be less than 2mm everywhere, less than the inter-patient variability that is observed (the smallest value is about 3mm). The second map shows the principal direction of each tensor (i.e., the eigenvector associated to the largest eigenvalue), whose coordinates are mapped on the RGB sphere, as presented in Fig. 7 right column. This map confirms the anatomical intuition that there are sets of sulci in certain cortical sectors that tend to vary in a consistent way in the top view, the principal direction of variation is lateral-to-medial for the superior frontal and parietal sulci, but the central and precentral sulci tend to vary more along an anterior-posterior direction. The temporal lobe sulci also tend to be consistent in varying with the same principal direction.

The spatial pattern of variability agrees with established neuroanatomical data. For instance, [Thompson et al., 2000] computed the variability of the cortex surface in an independent nor-




Fig. 6. Accessing the full map of cortical surface variability step by step. **Top:** Covariance matrices calculated along mean sulci. **Middle:** Matrices selected by the tensor picking operation. **Bottom:** Result of the extrapolation.

mal sample (15 controls) using a non-linear surface registration algorithm. Fig. 8 compares his variability map with ours. Our model of variability presents the same high values in the temporo-




Fig. 7. Left column: Variability map derived from the dense variability tensor field obtained by extrapolation. The color codes for the 3D rms variability (mm). Hot colors mean high variations among subjects. Right column: Maps showing the main direction of variability. The color codes for the main direction of the variability tensor at each point. Red: left-right oriented tensor, Green: posterior-anterior oriented, Blue: inferior-superior oriented.

parietal cortex (red and purple area, marked “A” in Fig. 8) and low values in the superior frontal gyrus (marked “B” in Fig. 8), Broca’s area, and the lower limits of the primary sensorimotor cor-

tices in the central and precentral gyri. Phylogenetically older areas (e.g. orbitofrontal cortex), and primary cortices that myelinate earliest during development (e.g., primary somatosensory and auditory cortex) exhibit least variability. The planum parietale (marked “A” in Fig. 8) consistently shows the highest variance of any cortical area, consistent with the complex pattern of secondary fissures surrounding the supramarginal and angular gyri (the perisylvian language cortex). It is also reasonable that the temporo-parietal areas around the Sylvian fissures are the most variable: they specialize and develop in different ways in each hemisphere, and are also the most asymmetric in terms of gyral patterning and volumes [Toga and Thompson, 2003].




Fig. 8. Comparison of two independent models of brain variability. The scalar value mapped on the mean cortex is the trace of the tensors (the variance). **Left:** Cortical variability map from [Thompson et al., 2000]. **Right:** Extrapolation of our simplified sulci variability model to the full brain. Note the similarity in the temporo-parietal cortex [shown in red colors (A)] and the superior frontal gyrus (B).

3.5 Evaluation of the Variability Model

Evaluating our extrapolated variability model is a tough issue. Obviously, using the information given by the sulci is not enough to infer the variability of the full brain, particularly within the brain (e.g. in the white matter, ventricles and deep gray matter nuclei). Moreover, we have no ground truth in these areas to validate the predicted variability. Thus, we restrict the evaluation of the predictive power of our model to the places where we have enough data: on the cortex. The first idea is to see how well our interpolation and extrapolation models fit the observed variability along each sulcus. This yields a root mean square error (RMSe) assessing the fidelity of the approximation. Then, we can perform a “leave one sulcus out” test to see if a group of sulci can correctly predict the variability of another sulcus in their neighborhood. This would mean that the model could effectively find missing data (i.e., the measures are dependent) and somehow predict the variability of missing structures in our datasets.

Intra-Sulcus Variability Recovery At each sampling point x , we computed the “difference” or error vector (Eq. 1) between the observed variability tensor (the reference) and the reconstructed one, either with our interpolation or with our extrapolation method. For both strategies, we found that the mean error was not significantly different from zero (p-value of 0.25 with Hotelling’s test). Second, we found a standard deviation $\sigma_{ref} = \sqrt{1/N \sum_{i=1}^N \text{trace}(Z_i^2)}$ of 0.15 for the interpolation error. This value gives us a lower bound on the range of the reconstruction errors. The slightly higher value of 0.21 for the extrapolation error could be attributed

to the aperture problem: in regions with orthogonal sulci, the normal component of one tensor influences the tangential part of its perpendicular neighbors and vice versa. One could think of solving the aperture problem in these regions by retrieving the tangential component of the variability from the normal one of surrounding sulci. However, in this study, we chose to consider such variability tensors as outliers and to remove them.

To further illustrate the prediction power of our extrapolation method, we compared the results of the extrapolation using 2000, 1000, 366 and 144 tensors, respectively. The first value corresponds to retaining all tensors. The second value is obtained by taking every other tensor, i.e. one tensor out of every two. The third value is the number of tensors retained by the tensor picking operation. The last value is the minimum number of tensors that the tensor picking operation can produce (2 measures for each of the 72 sulci, i.e. 144 tensors). Indeed, at least two tensor values per sulcus are required to be able to perform a linear interpolation between them. Figure 9 summarizes the experience results. One notices that even with very few tensors (366 compared to 2000 initially), the model is able to recover a correct estimation of the variability in almost all areas. Local errors arise when the correlation of variability between neighboring sulci is too low (see regions with hot colors in Fig. 9, right column).

To evaluate further our model, we now perform the "leave one sulcus out" test.

Leave One Sulcus Out This test removes one sulcus and its variability tensors from the model and extrapolates the rest of the data to the full brain. Then, the prediction error made on this specific sulcus is compared to the interpolation and extrapolation errors. As the measures are independent, an error below $3\sigma_{ref}$ is not significant and shows that our extrapolation model recovers the missing variability information up to the intrinsic reconstruction uncertainty. However, a RMSe larger than $3\sigma_{ref}$ means that we do not recover a comparable variability *in at least one direction*. We know that an uncertainty in the tangent of the mean sulcus could be induced by the aperture problem. To remove this effect, we "project" the error vector onto the plane perpendicular to the tangent of the mean sulcus. Thus, the error component in this direction is zeroed out. We will call this error the "partial error".

This test is performed on 3 sulci: the Sylvian Fissure, the Superior Temporal Sulcus and the Inferior Temporal Sulcus. Fig. 10 displays the variability of the reconstructed sulci after extrapolation with and without their tensors, while Table 1 summarizes the global RMSe statistics. The prediction error with missing sulci is globally 2 to 3 times larger than that incurred by interpolating or extrapolating the full model, but the difference is not high enough to be significant. However, errors are locally significant. In some places, like for the Sylvian Fissure, the prediction errors occur primarily in the tangential direction to the mean sulcus. Indeed, three main sulci (the Pre-Central, Central and Post-Central Sulcus), are orthogonally adjacent to the horizontal component of the Sylvian Fissure (Fig. 11), even though they do not actually merge with it, thus influencing the estimates of the tangential component of the variability as discussed in Sec. 3.5. Such behavior is confirmed by the "partial" error that is much lower than the standard one. By contrast, the variability of some sulci like the Central Sulcus (Fig. 11) cannot be correctly recovered from neighboring sulci: the error is not only due to the aperture problem but spatial correlations between adjacent sulci may be lower in some brain regions, making variations more difficult to predict.

In conclusion, our model is able to recover intra-sulcus variability from the selected tensors, and to predict the variability in regions that are locally correlated. Nevertheless, our evaluation method is limited for two reasons. First, the aperture problem will cause an underestimation of the variability along the direction of the mean curve, and this feature is in fact observed in regions with orthogonal sulci. However, there are not so many such regions, so this effect




Fig. 9. Decreasing the number of picked tensors. Left: Tensors retained for the model. **Middle:** Results of the extrapolation. **Right:** Error (Riemannian distance) between the extrapolated and initial covariance matrices field.

remains largely unnoticed. Second, the variability of some sulci is not correlated with that of their neighbors (this is the case for the Central Sulcus). These sulci carry some independent variability information, and should definitely be part of any brain variability model. One could consider including features as constraints for nonlinear registration of brain data when they contribute the greatest information on anatomical variation, i.e. information that cannot readily be predicted or interpolated from information on adjacent structures.




Fig. 10. Results of the “leave one sulcus out” test. **Top:** Positions of the 3 tested sulci in the ICBM305 space. **Bottom left:** variability of each sulcus after extrapolation of the complete model (from top to bottom: the Sylvian Fissure, the Superior Temporal Sulcus, the Inferior Temporal Sulcus). The color bar is the same as in Fig. 8. **Bottom right:** extrapolated variability from the neighboring sulci only.

Sulcus	Sylvian Fiss.	Sup. Temporal	Inf. Temporal
Interpolation	0.17 - 0.15*	0.17 - 0.15*	0.17 - 0.14*
Extrapolation	0.19 - 0.16*	0.21 - 0.19*	0.20 - 0.17*
Extrapolation w/o sulcus	0.56 - 0.34*	0.42 - 0.32*	0.38 - 0.32*

Table 1. RMSe of reconstruction of 3 sulci with the interpolation, extrapolation and leave one-sulcus out extrapolation methods. * indicates the “partial error” (Sec. 3.5).

4 Analysis of the Asymmetry of Brain Variability

The study of asymmetry of brain variability is of great interest for neuroscience [Toga and Thompson, 2003], and measures of structural and functional lateralization are of interest in mapping brain development, and disorders such as dyslexia, epilepsy, and schizophrenia. The two brain hemispheres develop according to slightly different genetic programs, and the right hemisphere is torqued forward relative to the left, with greatest volume asymmetries in the planum temporale and language cortex surrounding the Sylvian fissures (typically larger in the left hemisphere). If the types of variation in the two hemispheres could be differentiated, their genetic basis would be easier to investigate. It could also help understand whether there is an asymmetry in the power




Fig. 11. Illustration of orthogonal sulci. The Central, Pre-Central and Post-Central sulci have an orientation that is somewhat orthogonal to the Sylvian Fissure. This causes a difficulty in predicting the variability of the Sylvian Fissure from the other 3 somewhat orthogonal sulci. Another difficulty (less critical) is the prediction of the Central Sulcus from Post- and Pre-central sulci, because the Central Sulcus has a rather low variability compared to its immediate neighbors. Thus, the prediction will be difficult due to the lack of correlation in this region.

to detect consistent task-related or group-specific activation patterns in functional brain imaging studies, due to structural variance asymmetries.

We measured the symmetry/asymmetry of brain variability using our extrapolation model. The principle is to compute the distance between the variability tensor at one point and the (symmetrized) tensor at the symmetric point in the brain with Eq. 2. To define the symmetric point, a first geometric method is to simply use the 3D position that is symmetric with respect to the mid-sagittal plane in the stereotaxic space (mid-sagittal symmetry). In that case, we compute a dense asymmetry map from the extrapolated tensor values at each 3D point of a hemisphere (Fig. 12, left).

Another anatomical possibility is to measure brain asymmetry on sulcal lines, and extrapolate those measures to the whole brain (sulcal symmetry). First, all curves are mapped into a common hemisphere (e.g. the left hemisphere). Then, for each sulcus, a global mean is computed, as well as the left and right means (obtained by taking only the left (or respectively, the right) instances). Second, we compute the correspondences between this global mean and the left and right means. Thus, we define a common reference curve, the global mean, to compare left and right variability tensors. This prevents us from introducing a bias in the results, such as might happen if we had chosen either the left or right mean as the reference curve. Finally, difference vectors (Eq. 2) between left and right tensors are measured along the reference curve and extrapolated to the full brain using our framework. We end up with another dense asymmetry map, whose color is proportional to the distance between left-right tensors (Fig. 12 right).

A very interesting feature is that the regions with greatest asymmetries in variability include the one of main language areas, Broca’s speech area (labeled A in Fig. 12) as well as the parietal cortex, which exhibits the greatest gross anatomical variation of any cortical area (labeled B in Fig. 12). As expected, these areas are more structurally variable in the left hemisphere which is typically dominant for language in normal right-handers. One surprise is that the tips of the Sylvian fissures do not show the greatest difference in variability between the hemispheres, as these are regions with highly variably branching patterns that have been studied extensively in the language and dyslexia literature. Also as expected, the primary sensorimotor areas (central and pre-central gyri) are relatively symmetric in their variance, as the absolute variability is lower, as is their degree of hemispheric specialization (i.e. they perform analogous functions in each hemisphere, but innervate opposite sides of the body).




Fig. 12. Maps of brain variability asymmetry. **Left:** The mid-sagittal symmetry. **Right:** The sulcal symmetry.

5 Discussion

This article presents an original methodology to model the profile of structural variability in cortical landmarks. First, we began with a dataset of 72 expertly delineated sulcal lines in 98 subjects, and proposed an original mean sulcal landmark computation strategy. The approach consisted of alternately optimizing the position of the mean curve, computing the correspondences by dynamic programming and refining each subject’s affine transform based on the obtained matching. Then, variability tensors are measured along sulcal lines. Second, we applied a powerful Riemannian framework for tensor computing to extrapolate sparsely distributed tensors. We extend a RBF extrapolation method combined with diffusion PDEs. While the RBF provides a good initialization, the diffusion with attachment to the measures results in a smooth tensor field that stays close to the observed measures and converges in a few iterations. This agrees with the notion that there is a natural hierarchy of variability in the brain, with the variations of structures within a lobe or set of gyri being predictable from their immediate neighbors with relatively high accuracy. The resulting framework is very flexible in a sense that other sources of information can be easily incorporated to improve the accuracy of the variability model.

The results are also interesting neuroscientifically. The amplitude and asymmetry of variability are greatest in the most phylogenetically recent developments in the cortex: the frontal lobe, the dorsolateral prefrontal cortex, and in the more dorsal areas of the parietal association cortices. The language areas, in particular, have fundamentally different developmental programs in each brain hemisphere, leading to volumetric and functional asymmetries (e.g. left hemisphere language dominance). This variance asymmetry is seen in Broca’s area, which is specialized in the left hemisphere for producing speech, but is less commonly associated with structural asymmetries. Lower variance was seen in cortical regions subserving primary brain functions (e.g., touch, motor function, hearing) and these areas are the earliest to mature in utero. It would be interesting to hypothesize that the areas of the brain that emerged most recently in human evolution are the same ones that have greatest differences in variation patterns between the hemispheres, reflecting the drive towards hemispheric specialization of function in higher primates and man. The modeling of variance is practically valuable for understanding and discovering genetic and disease related factors that affect brain structure, which are currently hard to identify given the extremely complex patterns of variation in normal anatomy. For example, many neurodevelopmental disorders are associated with subtle variations in the patterning of the cortex, and new computational anatomy techniques are making these features easier to identify (e.g., increased cortical complexity in Williams syndrome [Thompson et al., 2005]).

When modeling variability, the main weakness is the unknown variability along the direction tangent to the mean sulci (aperture problem). We intend to tackle this point by first improving our sulcal matching algorithm to safely use the landmark information at the ends of sulci, and second by removing the data attachment term in the direction of the sulcal tangent. Doing this, the neighboring information could diffuse freely in that direction and hopefully reconstruct a globally coherent variability. Other approaches to this problem include the use of data from other imaging modalities, such as functional MRI and MEG, which further guide the matching of anatomy across subjects, for purpose of detecting systematic activation patterns or group or task-related differences. For the model validation, we need to compare to other sources of information, like sulcal ribbons [Mangin et al., 2004b], variability obtained from the matching of surfaces (e.g., ventricles or basal ganglia), fiber pathways mapped from DTI, or of full 3D images. As these sources of information are also subject to an aperture problem (we mainly retrieve the deformation in the direction of the gradient of the image), we expect to observe a good agreement in some areas, and complementary measures in other areas. Future applications of this work include the detection of genetic and demographic factors that contribute to brain variance, abnormality detection in individuals and groups, and improved nonlinear registration techniques that draw on tensor-valued statistical information regarding brain variation.

Acknowledgments. Paul M. Thompson is supported by the U.S. National Institutes of Health grants AG021431, LM05639, EB01651, and RR019771.

Part of this research is funded by the INRIA program “associated teams” (Brain Atlas). Website: <http://www-sop.inria.fr/epidaure/Collaborations/UCLA/>.

References

- [Arsigny et al., 2006a] Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. (2006a). Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*. In press.
- [Arsigny et al., 2006b] Arsigny, V., Fillard, P., Pennec, X., and Ayache, N. (2006b). Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421.
- [Ashburner and Friston, 2005] Ashburner, J. and Friston, K. J. (2005). Unified segmentation. *Neuroimage*, 26(3):839–851.

- [Bakircioglu et al., 1998] Bakircioglu, M., Grenander, U., Khaneja, N., and Miller, M. (1998). Curve matching on brain surfaces using induced Frenet distance metrics. *HBM*, 6(5):329–331.
- [Barillot et al., 1998] Barillot, C., Schwartz, D., le Goualher, G., Morandi, X., and Gibaud, B. (1998). Merging 3D cortical sulci models with 3D representations of meg data. In Paus, T., Gjedde, A., and Evans, A., editors, *Functional Mapping of the Human Brain*, volume 7(4), pages 796–796.
- [Batchelor et al., 2005] Batchelor, P., Moakher, M., Atkinson, D., Calamante, F., and Connelly, A. (2005). A rigorous framework for diffusion tensor calculus. *Mag. Res. in Med.*, 53:221–225.
- [Baumberg and Hogg, 1994] Baumberg, A. and Hogg, D. (1994). Learning flexible models from image sequences. In *Proc of ECCV'94 (vol. 1), Stockholm, Sweden*, LNCS 800, pages 299–308.
- [Cannon et al., 2005] Cannon, T., Hennah, W., van Erp, T., Thompson, P., Lonnqvist, J., Huttunen, M., Gasperoni, T., Tuulio-Henriksson, A., Pirkola, T., Toga, A. W., Kaprio, J., Mazziotta, J., and Peltunen, L. (2005). Association of disc1/trax haplotypes with schizophrenia, reduced prefrontal gray matter, and impaired short- and long-term memory. *Arch Gen Psychiatry*, 62:1205–1213.
- [Collins et al., 1995] Collins, D., Holmes, C., Peters, T., and Evans, A. (1995). Automatic 3d model-based neuroanatomical segmentation. *Human Brain Mapping*, 3(3):190–208.
- [Cootes et al., 1995] Cootes, T. F., Taylor, C. J., Cooper, D. H., and Graham, J. (1995). Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59.
- [Corouge et al., 2003] Corouge, I., Hellier, P., Gibaud, B., and Barillot, C. (2003). Interindividual functional mapping: a nonlinear local approach. *NeuroImage*, 19(4):1337–1348.
- [Coulon et al., 2004] Coulon, O., Alexander, D., and Arridge, S. (2004). Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67.
- [Fillard et al., 2005] Fillard, P., Arsigny, V., Pennec, X., Thompson, P. M., and Ayache, N. (2005). Extrapolation of sparse tensor fields: Application to the modeling of brain variability. In Christensen, G. and Sonka, M., editors, *Proc. of Information Processing in Medical Imaging 2005 (IPMI'05)*, volume 3565 of LNCS, pages 27–38, Glenwood springs, Colorado, USA. Springer.
- [Fischl et al., 2002] Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., van der Kouwe, A., Killiany, R., Kennedy, D., and A. Montillo, S. K., Makris, N., Rosen, B., and Dale, A. M. (2002). Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33:341–355.
- [Fletcher and Joshi, 2004] Fletcher, P. and Joshi, S. (2004). Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *Proc. of CVAMIA and MMBIA Workshops, Prague, Czech Republic, May 15, 2004*, LNCS 3117, pages 87–98. Springer.
- [Gee and R.K.Bajcsy, 1998] Gee, J. and R.K.Bajcsy (1998). *Brain Warping*, chapter Elastic matching: continuum mechanical and probabilistic analysis. Academic Press. In press.
- [Geschwind et al., 2002] Geschwind, D. H., Miller, B. L., DeCarli, C., and Carmelli, D. (2002). Heritability of lobar brain volumes in twins supports genetic models of cerebral laterality and handedness. *PNAS*, 99:3176–3181.
- [Goualher et al., 2000] Goualher, G. L., Argenti, A. M., Duyme, M., Baaré, W. F. C., Pol, H. E. H., Boomsma, D. I., Zouaoui, A., Barillot, C., and Evans, A. C. (2000). Statistical sulcal shape comparisons: Application to the detection of genetic encoding of the central sulcus shape. *Neuroimage*, 11(5):564–574.
- [Guéziec and Ayache, 1994] Guéziec, A. and Ayache, N. (1994). Smoothing and matching of 3-D space curves. *The International Journal of Computer Vision*, 12(1):79–104.
- [Guimond et al., 2000] Guimond, A., Meunier, J., and Thirion, J.-P. (2000). Average brain models: A convergence study. *Computer Vision and Image Understanding*, 77(2):192–210.
- [Kindlmann, 2004] Kindlmann, G. (2004). Superquadric tensor glyphs. In *IEEE TCVG Symposium on Visualization*, pages 147–154.
- [Kochunov et al., 2002] Kochunov, P., Lancaster, J., Thompson, P., Toga, A., Brewer, P., Hardies, J., and Fox, P. (2002). An optimized individual brain target in the Talairach coordinate system. *Neuroimage*, 17(2):922–927.
- [Lenglet et al., 2006] Lenglet, C., Rousson, M., and Deriche, R. (2006). Statistics on the manifold of multivariate normal distributions: Theory and application to diffusion tensor MRI processing. *Journal of Mathematical Imaging and Vision*. In press.
- [MacDonald et al., 1998] MacDonald, D., Avis, D., and Evans, A. (1998). Proximity constraints in deformable models for cortical surface identification. In *MICCAI*, pages 650–659.
- [Mangin et al., 2004a] Mangin, J.-F., Riviere, D., Cachia, A., Duchesnay, E., Cointepas, Y., Papadopoulos-Orfanos, D., Collins, D., Evans, A., and Regis, J. (2004a). Object-based morphometry of the cerebral cortex. *IEEE Transactions on Medical Imaging*, 23(8):968–982.
- [Mangin et al., 2004b] Mangin, J.-F., Riviere, D., Cachia, A., Duchesnay, E., Cointepas, Y., Papadopoulos-Orfanos, D., Scifo, P., Ochiai, T., Brunelle, F., and Regis, J. (2004b). A framework to study the cortical folding patterns. *NeuroImage*, 23(Supplement 1):S129–S138.
- [Paulsen and Hilger, 2003] Paulsen, R. R. and Hilger, K. B. (2003). Shape modelling using markov random field restoration of point correspondences. In *IPMI'03*, LNCS 2732, pages 1–12.

- [Pennec et al., 2006] Pennec, X., Fillard, P., and Ayache, N. (2006). A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66.
- [Pitiot et al., 2004] Pitiot, A., Delingette, H., Thompson, P., and Ayache, N. (2004). Expert knowledge guided segmentation system for brain MRI. *NeuroImage*, 23(supplement 1):S85–S96. Special Issue: Mathematics in Brain Imaging.
- [Pitiot et al., 2003] Pitiot, A., Delingette, H., Toga, A., and Thompson, P. (2003). Learning object correspondences with the observed transport shape measure. In *Proc. of IPMI'03*, LNCS 2732, pages 25–37. Springer Verlag.
- [Skovgaard, 1984] Skovgaard, L. (1984). A Riemannian geometry of the multivariate normal model. *Scand. J. Statistics*, 11:211–223.
- [Sowell et al., 2004] Sowell, E., Thompson, P., and Toga, A. (2004). Mapping changes in the human cortex throughout the span of life. *The Neuroscientist*, 10(4):372–92.
- [Sun, 1995] Sun, X. (1995). Conditional positive definiteness and complete monotonicity. In *Approximation Theory VIII*, volume 2, pages 211–234. World Scientific.
- [Thompson et al., 2001a] Thompson, P., Cannon, T., Narr, K., van Erp, T., Poutanen, V., Huttunen, M., Lönqvist, J., Standertskjöld-Nordenstam, C., Kaprio, J., Khaledy, M., Dail, R., Zoumalan, C., and Toga, A. (2001a). Genetic influences on brain structure. *Nature Neuroscience*, 4(12):1253–1258.
- [Thompson et al., 2004] Thompson, P., Hayashi, K., Sowell, E., Gogtay, N., Giedd, J., Rapoport, J., de Zubicaray, G., Janke, A., Rose, S., Semple, J., Doddrell, D., Wang, Y., van Erp, T., Cannon, T., and Toga, A. (2004). Mapping cortical change in alzheimer's disease, brain development, and schizophrenia. *Neuroimage Special Issue on Mathematics in Brain Imaging*.
- [Thompson et al., 2005] Thompson, P., Lee, A., Dutton, R., Geaga, J., Hayashi, K., Eckert, M., Bellugi, U., Galaburda, A., Korenberg, J., Mills, D., Toga, A., and Reiss, A. (2005). Abnormal cortical complexity and thickness profiles mapped in williams syndrome. *Journal of Neuroscience*, 25(16):4146–4158.
- [Thompson et al., 2000] Thompson, P., Mega, M., Narr, K., Sowell, E., Blanton, R., and Toga, A. (2000). Brain image analysis and atlas construction. In Fitzpatrick, M. and Sonka, M., editors, *Handbook of Medical Image Proc. and Analysis*, chapter 17. SPIE.
- [Thompson et al., 2001b] Thompson, P., Mega, M., Woods, R., Zoumalan, C., Lindshield, C., Blanton, R., Mousai, J., Holmes, C., Cummings, J., and Toga, A. (2001b). Cortical change in alzheimer's disease detected with a disease-specific population-based brain atlas. *Cerebral Cortex*, 11(1):1–16.
- [Toga and Thompson, 2003] Toga, A. and Thompson, P. (2003). Mapping brain asymmetry. *Nature Reviews Neuroscience*, 4(1):37–48.
- [Troune and Younes, 2000] Troune, A. and Younes, L. (2000). Diffeomorphic matching problems in one dimension: Designing and minimizing matching functionals. In *Proc of ECCV'00, part I*, LNCS 1842, pages 573–587.
- [Tschumperlé and Deriche, 2002] Tschumperlé, D. and Deriche, R. (2002). Orthonormal vector sets regularization with pde's and applications. *Int. J. of Computer Vision (IJCV)*, 50(3):237–252.
- [Wang et al., 2005] Wang, Y., Chiang, M., and Thompson, P. (2005). Automated surface matching using mutual information applied to riemann surface structures. In *Medical Image Computing and Computer Assisted Interventions (MICCAI)*, number 2 in lncs, pages 666–674, Palm Springs, CA.
- [Woods, 2003] Woods, R. P. (2003). Characterizing volume and surface deformations in an atlas framework: theory, applications, and implementation. *Neuroimage*, 18(3):769–788.