

Extrapolation of Sparse Tensor Fields: Application to the Modeling of Brain Variability.

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

¹ INRIA Sophia - Projet Epidaure, 2004 Route des Lucioles BP 93
06902 Sophia Antipolis Cedex, France

{Pierre.Fillard, Vincent.Arsigny, Xavier.Pennec,
Nicholas.Ayache}@Sophia.Inria.fr

² Lab of Neuro Imaging, UCLA School of Medicine
4238 Reed Neurology, 710 Westwood Plaza, Los Angeles, CA, USA
thompson@loni.ucla.edu

Abstract. Modeling the variability of brain structures is a fundamental problem in the neurosciences. In this paper, we start from a dataset of precisely delineated anatomical structures in the cerebral cortex: a set of 72 sulcal lines in each of 98 healthy human subjects. We propose an original method to compute the average sulcal curves, which constitute the mean anatomy in this context. The second order moment of the sulcal distribution 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 PDEs to tensor fields. As a result, we obtain a dense 3D variability map which proves to be in accordance with previously published results on smaller samples subjects. Moreover, leave one (sulcus) out tests show that our model is globally able to recover the missing information when there is a consistent neighboring variability. Last but not least, we propose innovative methods to analyze the asymmetry of brain variability. As expected, the greatest asymmetries are found in regions that includes the primary language areas. Interestingly, such an asymmetry in anatomical variance could explain why there may be greater power to detect group activation in one hemisphere than the other in fMRI studies.

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. 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, and could be used in Bayesian approaches for nonlinear registration as well (which adjust for anatomical variations across subjects). Finally, neuroscientists are interested in identifying the causes of brain variability at a genetic or environmental level. 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. 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.

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 (linear) differences are factored out. Ideally, one could measure or 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 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*³. Thus, our goal is to compute the field of variability tensors within the brain, from information that may be sparsely distributed.

However, working with tensors is not so easy as the underlying space is a manifold that is not a vector space. As tensors constitute a convex half-cone in the vector space of matrices, many operations (like computing the mean) are stable. Nonetheless, this 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.

In Sec. 2.1 we present a consistent Riemannian framework to compute with tensors. Then, we show in Sec. 2.2 how to extend these tools to implement harmonic diffusion PDEs and extrapolate tensors that are sparsely distributed in space. Solving these PDEs is computer intensive, so in Sec. 2.3 we provide a practical but efficient initialization by extending the radial basis functions (RBF) concept to tensors. In Sec. 3, we consider low dimensional but anatomically very readily defined and delineated features (sulcal lines) as a way to obtain meaningful brain variability tensors. We show in Sec. 3.1 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 in Sec. 3.2 a parametric tensor model to these data. Then, we come back to our original goal in 3.3 by extrapolating this sparse tensor model to the whole brain. The validity of our extrapolated model is analyzed in Sec. 3.4. In Sec. 3.5

³ We follow here the terminology used in the Diffusion Tensor Imaging field: by tensor we mean a positive definite symmetric matrix.

we generalize our statistical model to examine the correlation of the variations observed at symmetric points in the brain.

2 A Mathematical Framework to Extrapolate Tensors

Much 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, [1] anisotropically restores the principal direction of the tensors, while [2] independently restores the eigenvalues and eigenvectors. This last approach requires an additional reorientation step of the eigenvectors due to the non-uniqueness of the decomposition.

More recently, differential geometric approaches have been developed to generalize the PCA to tensor data [3], for statistical segmentation of tensor images [4], for computing a geometric mean and an intrinsic anisotropy index [5], or as the basis for a full framework for Riemannian tensor calculus [6]. In this last work, 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 [7], and turns out to be the basis of all the previous differential geometric approaches.

2.1 A Riemannian Framework for Tensor Calculus

The invariant metric provides a new framework in which the limitations of Euclidean calculus are fully overcome: it endows the tensor space with a very regular structure where matrices with null or negative eigenvalues are at an infinite distance from any positive definite matrix. Moreover, the geodesic between two tensors is uniquely defined, leading to interesting properties such as the existence and uniqueness of the (geometric) mean [6].

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. 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 [6]:

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

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 as $Z = \Sigma_1^{-1/2} \log_{\Sigma_1}(\Sigma_2) \Sigma_1^{-1/2} = \log(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2})$ in the tangent space at the identity (i.e. Z is a symmetric but not necessarily positive matrix). The distance between the two tensors is simply given by:

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

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* $\Sigma_{t+1} = \exp_{\Sigma_t}(-\varepsilon \nabla C(\Sigma_t))$.

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 [6] a least square 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 homogeneous result. The simplest criterion is the harmonic regularization: $\text{Reg}(\Sigma) = \frac{1}{2} \int_{\Omega} \|\nabla \Sigma(x)\|_{\Sigma}^2$. We showed in [6] 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:

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

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 paper we develop a RBF interpolation.

2.3 Extending RBFs to Extrapolate Tensors

RBFs provide a family of methods to extrapolate sparsely defined observations [8]. The extrapolated scalar field is expressed as a linear combination of translated versions of a single radial function called 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, we find the coefficients such that $y(x_i) = y_i$, by solving the linear system $\forall j : y_j = \sum_i \lambda_i h(x_j - x_i)$. To ensure a unique solution exists for any set of measurements at any set of spatial positions, the symmetric matrix H generated by $[H]_{i,j} = h(x_i - x_j)$ must always be *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 $h(x) = 1 / (1 + (\|x\|^2 / \alpha^2)^\gamma)$. This way, we can choose a meaningful asymptotic value for the interpolation: the global mean of the tensors.

3 Modeling Brain Variability from Sulcal Lines

To model the spatial pattern of variability in brain structure, 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 sulcal lines. Over 70 sulcal curves consistently appear in all normal individuals and allow a consistent subdivision of the cortex into major lobes and gyri [9]. 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.

We use a dataset of sulcal lines, or *sulci*, manually delineated in 98 subjects by expert neuroanatomists according to a precise protocol⁴. We included the maximal subset of all sulcal curves that consistently appear in all normal subjects (72 in total), with formal rules governing the handling of branching patterns, breaks in sulci, and doubling of specific sulci (e.g. the cingulate). By repeated training on test sets of brain images, the maximum allowed inter- and intra-rater error (reliability) was ensured to be 2mm everywhere, and in most regions less than 1mm, far less than the intersubject anatomical variance. Delineations were made in 3D on cortical surfaces extracted from MR images linearly aligned to the ICBM stereotactic space, 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.

⁴ Reference undisclosed for the review

3.1 Learning Local Variability from a Sulcal Lines Dataset

Statistical models have frequently been constructed for objects such as open curves, closed curves, surfaces [10–12]. In each of these examples, the *aperture problem occurs*: without more specific anatomical knowledge, it is not possible to recover the exact correspondences between instances of a surface or a curve. Point-to-point correspondences of two instances of a sulcus are intrinsically subject to error, with greater tangential than normal uncertainty. In the absence of further information, 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. In this continuous formulation, the number of degrees of freedom can be adjusted to increase robustness to noise while avoiding resampling problems, as in [13]. Typically, we reduce the number of control points to one third of the original sampling points.

Many criteria have been proposed in the literature to evaluate the appropriateness of one-to-one correspondences between geometric objects. They usually invoke local differential characteristics such as the tangent space, curvature [14], the local Frenet frame for a curve on a surface, regional shape information [15]. In our case, the variability is so large (see Fig. 1), that using such refined measures is meaningless. Therefore, we simply use the total variance of curve models as a criterion. Minimizing this variance greatly reduces the variability due to inadequate correspondences. Practically, we alternately improve the correspondences between the mean curves 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 after a few iterations.

For each of the 72 sulci, we now have the mean curve position $\bar{c}(t)$, parameterized by B-splines, and a one-to-one mapping that gives the corresponding point $c_i(t)$ in each instance. The variability tensor $\Sigma(t)$ along the mean sulcus is: $\Sigma(t) = \sum_{i=1}^n [c_i(t) - \bar{c}(t)][c_i(t) - \bar{c}(t)]^T / (n - 1)$. An example set of covariance tensors estimated along the Sylvian Fissure is shown in Fig. 1. Variability is greater at the extremities of the sulci. These points should be landmarks as they are precisely identifiable by neuro-anatomists. We believe that the main part of their variability is due to a bias when we estimate the position of the end points of the mean curve. To remain consistent, we chose in this paper to remove this information from our model, and focussed only on the interior part of the sulci.

3.2 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 specific points along the mean sulcus, and interpolating in between them. We use interpolation along the geodesic joining 2 tensors, because it preserves the monotonic evolution of the determinant. This is crucial when interpolating two probability densities and is in general not possible 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 $\Sigma(t_i)$ be a set of N covariance tensors defined at parameter t_i along a mean sulcus. Riemannian interpolation between them gives the tensor: $\tilde{\Sigma}(t) = \exp_{\Sigma(t_i)}[(t - t_i)/(t_{i+1} - t_i) \log_{\Sigma(t_i)}(\Sigma(t_{i+1}))]$ for $t_i \leq t < t_{i+1}$. As we are working only on the interior of the sulcus, t takes its values between t_2 and $t_N - 1$, so that the *interpolated variability* $\tilde{\Sigma}(t)$ is always defined. The tensor picking operation consists of finding the optimal t_i such that the least-square error between the observed and interpolated variability tensors is minimized: $C(\Sigma) = \int_{t_1}^{t_N} \text{dist}^2(\Sigma(t), \tilde{\Sigma}(t)) dt$. To minimize this criterion, N points (i.e. tensors) are uniformly chosen along the mean curve. Then, an exhaustive search for the optimal point positions is done. If the criterion value at the optimal set is below a given threshold (0.7 in our experiments), the tensors are picked, otherwise the number of chosen tensors N is increased and the search is repeated.

Results of this operation are presented in Fig. 2 (middle panel): by choosing tensors at adequate positions, one can accurately reconstruct the full variability of each sulcus using 4 to 10 matrices, depending on its length and shape. The variability of all the major sulci can be represented by about 310 variability tensors out of 2000 initially.

3.3 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. Fig. 2 presents the result of the extrapolation of our 310 tensors on 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. (1)).

The spatial pattern of variability agrees with established neuroanatomical data. For instance, [16] computed the variability of the cortex surface in an independent normal sample (15 controls) using a non-linear surface registration algorithm. Fig. 3 compares his variability map with ours. Our model of variability presents the same high values in the temporo-parietal cortex (red area, marked “b” in Fig. 3) and low values in the superior frontal gyrus (marked “a” in Fig. 3),




Fig. 1. Sulcal variability. Left: The Sylvian Fissure mean curve (in red) with traces from 98 healthy normal individuals (in green and yellow). Right: 50 covariance matrices (1 σ ellipsoids) are overlaid on the mean sulcus. Note that the very first and last tensors are larger than the interior ones.

Broca’s language area, and the lower limits of the primary sensorimotor cortices in the central and precentral gyri. Phylogenetical older areas (e.g. orbitofrontal cortex), and primary cortices that myelinate earliest (e.g. primary somatosensory and auditory cortex) exhibit least variability. The planum parietale (marked “b” in Fig. 3) 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.

3.4 Validation of the Variability Model

Validating 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




Fig. 2. Accessing the full brain variability step by step. The color bar is the same as in Fig. 3. **Left:** Covariance matrices calculated along the mean sulci. **Middle:** Matrices selected by the tensor picking operation. **Right:** Result of the extrapolation.




Fig. 3. 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 [16]. **Right:** Extrapolation of our simplified sulci variability model to the full brain (the display is restricted to the cortex). Note the similarity in the superior frontal gyrus (a) and the temporo-parietal cortex [shown in red colors (b)].

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 fits 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 can effectively find missing data (the measures are independent) and somehow predict the variability of missing structures in our datasets.

Intra-Sulcus Variability Recovery We computed the “difference” or error vector between the observed variability tensor and the reconstructed one with our interpolation and extrapolation methods. We found that the mean errors were not significantly different from zero (p-value of 0.25 at the Hotelling’s test). Second, we found a standard deviation of $\sigma_{ref} = 0.15$ 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, which misleads the reconstruction. After removing these “outliers”, the error distributions after interpolation and extrapolation are comparable.

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 Main Body and the Inferior Temporal Sulcus. Fig. 4 displays the reconstructed sulci after extrapolation with and without their variability 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. Such behavior was expected due to the aperture problem and is confirmed by the “partial” error that is much lower than the standard

Sulcus	Sylvian Fiss.	Sup. Temporal	Inf. Temporal.
Interpolation	0.12 - 0.10*	0.17 - 0.15*	0.17 - 0.14*
Extrapolation	0.18 - 0.13*	0.21 - 0.17*	0.17 - 0.15*
Extrapolation w/o sulcus	0.43 - 0.27*	0.37 - 0.32*	0.27 - 0.22*

Table 1. RMSe of reconstruction of 3 sulci with the interpolation, extrapolation and leave one-sulcus out extrapolation methods. * indicates the “partial error” (Sec. 3.4).

error. By contrast, the variability of some sulci like the Central Sulcus 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.

3.5 Analysis of the Asymmetry of Brain Variability

The study of asymmetry in brain variability is of great interest for neuroscientists [17], and measures of structural and functional lateralization are of interest in mapping brain development, and disorders such as dyslexia 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 to detect group activation in functional brain imaging studies, due to structural variance asymmetries.

We therefore measured the symmetry/asymmetry of brain variability using our extrapolation model. The main principle is to compute the distance between the variability tensor at one point and the (symmetrized) tensor at the symmetric point in the brain. To define the symmetric point, we may simply use the 3D




Fig. 4. Result of the “leave one sulcus out” test. **Left:** Positions of the 3 tested sulci in the ICBM305 space. **Middle:** variability of each sulcus after extrapolation of the complete model. The color bar is the same as in Fig. 3. **Right:** extrapolated variability from the neighboring sulci only. **Top:** the Sylvian Fissure, **middle:** the Superior Temporal Sulcus main body, **bottom:** the Inferior Temporal Sulcus.



Fig. 5. Maps of the asymmetry of the brain variability. Red to purple colors indicate a significant asymmetry. **Left:** Asymmetry of the 3D extrapolation w.r.t. the mid-sagittal plane. **Middle:** Difference vectors between left-right variability tensors. **Right:** Extrapolation to the 3D volume of the “asymmetry vectors” of the previous figure.

position that is symmetric w.r.t. the mid-sagittal plane in the stereotaxic space (ICBM 305). In that case, we compute a dense asymmetry map from the extrapolated tensor values at each 3D point of a hemisphere (Fig. 5, left). We may also retrieve the corresponding points between each left and right mean sulcus by mapping the left sulci into the right hemisphere and computing the correspondences using the matching algorithm of Sec. 3.1. In that case, we end up with an error tensor measuring the asymmetry along each sulcus (Fig. 5 middle). This error tensor is finally extrapolated to the full brain using once again the framework previously developed (Fig. 5, right). A very interesting feature is that the regions with greatest asymmetries in variability include the 2 main language areas, Broca’s speech area (see pink colors in the inferior frontal cortex) and Wernicke’s language comprehension area (yellow and red colors surrounding the posterior Sylvian fissure). As expected, these areas vary more on the left hemisphere which is dominant for language. The greater left hemisphere variation may be attributable to the greater volumes of structures such as the planum temporale in the left hemisphere. 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).

4 Discussion

This paper applies 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. We applied this methodology to model the profile of brain variability, where tensors are measured along sulcal lines that are consistently anatomical landmarks.

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. For the model validation, we need to compare to other sources of information, like the 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 obtain a good fit in some areas, but we require complementary measures in other areas.

These results are also interesting neuroscientifically. Variance and the asymmetry of variability are greatest in language areas, which have fundamentally different developmental programs in each brain hemisphere, leading to volumetric and functional asymmetries (e.g. left hemisphere language dominance). This variance asymmetry was also 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. The modeling of variance is practically valuable for understanding the 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.

References

1. O. Coulon, D. Alexander, and S. Arridge. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67, 2004.
2. D. Tschumperlé and R. Deriche. Orthonormal vector sets regularization with pde's and applications. *Int. J. of Computer Vision (IJCV)*, 50(3):237–252, 2002.
3. P.T. Fletcher and S.C. Joshi. 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, 2004.
4. C. Lenglet, M. Rousson, R. Deriche, and O. Faugeras. Statistics on multivariate normal distributions: A geometric approach and its application to diffusion tensor MRI. Research Report 5242, INRIA, 2004.
5. P. Batchelor, M. Moakher, D. Atkinson, F. Calamante, and A. Connelly. A rigorous framework for diffusion tensor calculus. *Mag. Res. in Med.*, 53:221–225, 2005.
6. Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *International Journal of Computer Vision*, 2005. To appear (accepted for publication).
7. L.T. Skovgaard. A riemannian geometry of the multivariate normal model. *Scand. J. Statistics*, 11:211–223, 1984.
8. X. Sun. Conditional positive definiteness and complete monotonicity. In *Approximation Theory VIII*, volume 2, pages 211–234. World Scientific, 1995.
9. J.-F. Mangin, V. Frouin, J. Régis, I. Bloch, P. Belin, and Y. Samson. Towards better management of cortical anatomy in multi-modal multi-individual brain studies. *Physica Medica*, 12(Supplement 1):103–107, June 1996.

10. T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Comput. Vis. Image Underst.*, 61(1):38–59, 1995.
11. A. Trounev and L. Younes. Diffeomorphic matching problems in one dimension: Designing and minimizing matching functionals. In *Proc of ECCV'00, part I*, LNCS 1842, pages 573–587, 2000.
12. R. R. Paulsen and K. B. Hilger. Shape modelling using markov random field restoration of point correspondences. In *IPMI'03*, LNCS 2732, pages 1–12, 2003.
13. A. Baumberg and D. Hogg. Learning flexible models from image sequences. In *Proc of ECCV'94 (vol. 1), Stockholm, Sweden*, LNCS 800, pages 299–308, 1994.
14. M. Bakircioglu, U. Grenander, N. Khaneja, and M.I. Miller. Curve matching on brain surfaces using induced Frenet distance metrics. *Human Brain Mapping*, 6(5):329–331, 1998.
15. A. Pitiot, H. Delingette, A. Toga, and P. Thompson. Learning object correspondences with the observed transport shape measure. In *Proc. of IPMI'03*, LNCS 2732, pages 25–37. Springer Verlag, 2003.
16. P.M. Thompson, M.S. Mega, K.L. Narr, E.R. Sowell, R.E. Blanton, and A.W. Toga. Brain image analysis and atlas construction. In M. Fitzpatrick and M. Sonka, editors, *Handbook of Medical Image Proc. and Analysis*, chapter 17. SPIE, 2000.
17. A.W. Toga and P.M. Thompson. Mapping brain asymmetry. *Nature Reviews Neuroscience*, 4(1):37–48, January 2003.