

Variational, geometric and statistical methods for modeling brain anatomy and function

Olivier Faugeras, Geoffray Adde, Guillaume Charpiat, Christophe Chéd'Hotel, Maureen Clerc, Rachid Deriche, Gerardo Hermosillo, Renaud Keriven, Pierre Kornprobst, Jan Kybic, Christophe Lenglet, Lucero Lopez-Perez, Théo Papadopoulo, Jean-Philippe Pons, Florent Segonne, Bertrand Thirion, David Tschumperlé, Thierry Viéville, Nicolas Wotawa
Odysée Laboratory – ENPC/ENS/INRIA.
INRIA, 2004 Route des Lucioles, BP93, 06902 Sophia-Antipolis, France.

April 1, 2004

Abstract

We survey the recent activities of the Odysée Laboratory in the area of the application of mathematics to the design of models for studying brain anatomy and function. We start with the problem of reconstructing sources in MEG and EEG and discuss the variational approach we have developed for solving these inverse problems. This motivates the need for geometric models of the head. We present a method for automatically and accurately extracting surface meshes of several tissues of the head from anatomical MR images. Anatomical connectivity can be extracted from Diffusion Tensor Magnetic Resonance Images but, in the current state of the technology, it must be preceded by a robust estimation and regularization stage. We discuss our work based on variational principles and show how the results can be used to track fibers in the white matter as geodesics in some Riemannian space. We then go to the statistical modeling of fMRI signals from the viewpoint of their decomposition in a pseudo-deterministic and stochastic part which we then use to perform clustering of voxels in a way that is inspired by the theory of Support Vector Machines and in a way that is grounded in information theory. Multimodal image matching is discussed next in the framework of image statistics and Partial Differential Equations with an eye on registering fMRI to the anatomy. The paper ends with a discussion of a new theory of random shapes that may prove useful in building anatomical and functional atlases.

Keywords: MEG, EEG, fMRI, DT-MRI, Maxwell equations, inverse problems, segmentation, level sets, Riemannian spaces, Partial Differential Equations, Brownian motion, Diffusion tensor, Lie groups, Eikonal equation, tractography, mutual information, information theory, Kernel PCA, Information Bottleneck, retinotopy, Laplace-Beltrami operator, shape topologies, Hausdorff distance, mean shape, covariance of shapes.

INTRODUCTION

The Odysée laboratory is interested in developing a detailed understanding of the neural computations underlying human visual perception. This interest arises from several motivations. One is the desire to participate in the increase of basic knowledge regarding one of the most sophisticated sensory modality that supports action and reasoning, another one is the hope that this quest will eventually lead to breakthroughs in the way we interact with computers.

In order to model human visual perception it is necessary to observe humans while they perform the act of seeing. This can be achieved by using such imaging/measurement techniques as MRI, MEG, EEG in order to make qualitative and quantitative measurements of the changes in the state parameters of some volunteers' brains. These measurements can then be used to support the design and the test of neural mathematical and computational models of human visual perception. In this article we focus only on the part of the problem that consists in processing the data that are produced by some of the previous imaging modalities.

The theory of Partial Differential Equations (PDE) is central to the source reconstruction problem in MEG and EEG, and to our work with Diffusion Tensor Magnetic Resonance Imaging (DT-MRI); it is also the core of our curve and surface evolution work as applied to segmentation and warping. Differential geometry plays an important role because we are dealing with spaces with a natural Riemannian structure, e.g. the white matter through the diffusion tensor or the cortical surface. Many of these PDEs derive from energy functionals through the calculus of variation as in our work on multimodal image matching. Last but not least, statistics are crucial to correctly take into account the immense variability of the signals and the shapes that arise when one attempts to "look" at the brain.

SOURCE RECONSTRUCTION IN MEEG

MEG and EEG (commonly called MEEG) record noninvasively the electromagnetic field resulting from electrical activity inside the brain. Their high time resolution (of the order of 1 ms) make these two modalities very valuable for the functional analysis of the human brain. Their drawbacks are 1) a relatively poor spatial resolution (compared to fMRI) and 2) the need for solving a delicate inverse problem for localizing the electrical activity inside the brain.

This inverse problem is driven by a forward model, which computes the electromagnetic field outside the head from a known electrical activity profile inside the brain. The quality of the reconstruction greatly depends on the forward problem, whose accuracy must be controlled with great care. This problem obeys the rules of electromagnetic propagation under the quasi-static approximation (because at frequencies of interest and at spatial scales smaller than the head, inductive and capacitive effects can be ignored). In this case, the Maxwell equations relate the potential V to the sources J_p for tissues of conductivity σ through

$$\nabla \cdot (\sigma \nabla V) = \nabla \cdot J_p, \quad (1)$$

with a vanishing Neumann boundary condition on the scalp. The magnetic field can be computed from the potential, e.g. through the Biot-Savart equation.

Of course, a good geometrical and physical model of the head is crucial to solve properly this equation. Our group relies on two different – mesh based – strategies: the boundary element method (BEM) which describes the head as a set of nested surfaces delimiting domains each with a uniform, isotropic conductivity; the finite (volume) element method (FEM) which can assign a different conductivity to each tessell. The meshes describing the head are naturally subject dependent and must be computed beforehand from physiological data, via anatomic MRI segmentation. As the head is a very complex object, an accurate geometrical description signifies handling a huge number of geometrical elements, on which the solution is discretized, and the resulting computations can only be solved via iterative methods (rather than a direct matrix inversion) using leading-edge numerical methods.

In our work, we use a distributed source model for the inverse problem. This is notoriously ill-posed, due to the existence of “silent sources”. Consequently, some constraints on the solution must be added (minimum norm or minimum gradient solutions are often chosen). The inverse problem is thus solved by minimizing an energy term which is the sum of a data term and a regularizing term taking the constraints into account. The sources are iteratively updated by computing the gradient of the energy. The computationally efficient way to compute the energy gradient is to use the adjoint problem (Faugeras et al., 1999). Note that, since the relationship (1) between the source term J_p and the potential V is linear, the matrix which represents the adjoint problem is simply the transpose of the matrix

representing the forward problem.

The BEM method takes advantage of the harmonicity of a potential V satisfying (1) inside each compartment of homogeneous conductivity. It uses the representation theorem to reformulate the problem in terms of single-layer and double-layer potentials defined on the interfaces between conductivities. It requires only surface meshes, which are much easier to obtain than volume meshes. A drawback in terms of numerical complexity is that in the BEM the matrix representing the forward problem is dense and requires a lot of memory to be computed and stored. To alleviate this problem, we have introduced the Fast Multipole Method (FMM), an algorithm originally developed for N-body gravitational field computations (Clerc et al., 2002). The FMM is a multi-resolution approach which approximates the electromagnetic interaction between surface elements by performing multipole expansions at coarse resolutions. It avoids matrix storage, and significantly reduces the computational burden of the matrix-vector products: a matrix-vector product of dimension P is performed in $O(P \log P)$ instead of $O(P^2)$. The computational savings allow finer discretizations to be used, leading to more accurate forward and inverse problem solutions.

Accuracy of the forward problem is also a major concern in BEM, because the electrical sources are close to the discontinuities in conductivity, causing the potential to have very sharp variations, which are difficult to discretize. All the BEM methods used so far in MEEG have used the same integral formulation (Geselowitz, 1967). However, this integral formulation, based on a double-layer potential approach, is by no means the only one available. A thorough analysis of integral formulations deriving from the representation theorem has enabled us to propose a new formulation, combining single and double-layer potentials (Kybic et al., 2003). This approach has three main advantages compared to the previous BEM: it leads to symmetric matrices, it only couples elements from adjacent surfaces, and its accuracy outperforms all other surface methods (Adde et al., 2003), especially when the ratio of conductivities between neighboring layers is high, as occurs inside the head.

For the BEM inverse problem, the electrical sources are constrained to be orthogonal to a known surface inside the cortex. In this case, the problem is no longer ill-posed and the intensity field of the sources can be recovered, up to a constant, from boundary measurements (Amir, 1994). The field to be reconstructed is then simply the signed intensity of the source.

Figure 1 shows the reconstruction of electrical sources, using the BEM model, from measurements of a somatosensory evoked magnetic field (experimental protocol presented in (Munier et al., 2001) and data also used in (David and Garnero, 2002)).

The FEM method represents all the head related quantities (conductivities, potentials, sources) as piecewise linear functions on the elements of the mesh describing the head. Obviously, this restrictive model has a strong impact on the accuracy of the results compared to the BEM method (in which no re-




Figure 1. BEM reconstruction of somatosensory sources from MEG data (CTF Omega 151-channel, Hopital de la Salpetriere, Paris).

striction is made on the solution in the domains delimited by the surfaces). This has, however, several advantages: anisotropy of conductivities can easily be modeled for each volume element, the accuracy seems somewhat less sensitive to the closeness of the sources to the conductivity discontinuity interfaces, the matrices generated by the method are quite sparse leading to computation times that are much smaller than those of the BEM method. It remains that constructing accurate 3D meshes to model the head is somewhat difficult mainly because the cortex is very thin and thus a lot of triangles are needed to represent it correctly. The sources are restricted to belong to the *volume* of the gray matter (instead of a surface for the BEM method), but they can again be constrained in direction by forcing them to be aligned with the (common) normal to the interfaces of the grey matter.

GEOMETRIC MODELING

Anatomical MRI segmentation

We have designed a method to automatically and accurately extract surface meshes of several head tissues from anatomical MRI images. The input of our algorithm is a T_1 -weighted MRI image and the approximate intensities of the main head tissues: air, skin, cerebrospinal fluid (CSF), gray matter (GM) and white matter (WM). It robustly generates triangle meshes of the outer skin interface, of the brain contour and of the inner and outer interfaces of the cortex. In the future, we plan to extend it to the skull and to some subcortical structures of interest.

Our method guarantees some topological properties of the meshes, such as spherical topology, absence of self- or mutual intersections. These properties are crucial in some applications such as cortex unfolding or source reconstruction in MEG and EEG with the BEM or the FEM.

Our method is a successful combination of hidden Markov random field (HMRF) classification (Zhang et al., 2001) and of active contour segmentation with the topology preserving level set method (Han et al., 2003).

The former is a statistical approach to classify voxels into a small number of tissue classes chosen a priori. The tissue distribution is modeled by a Markov random field (MRF) encouraging neighboring voxels to have the same class labels, while the observed intensity of each tissue class is modeled by a Gaus-

sian distribution. The labels of the voxels are estimated with a maximum a posteriori (MAP) criterion. The problem translates into the minimization of an energy function but an exact minimization is computationnally unfeasible due to the huge number of unknowns. As a consequence, a greedy strategy yielding a suboptimal solution, such as the iterated conditional modes (ICM) algorithm, is preferred. The parameters of the model are the mean and the variance of each tissue class and a bias field accounting for the inhomogeneties in the RF field. In our method, this bias is taken as affine against intensities and smooth and non-parametric over space. An initial estimate of the tissue parameters is provided by the user. Iteratively, class labels are estimated by MAP, then the tissue parameters and the bias field are updated with the expectation-maximization (EM) algorithm. The output is a classification of image voxels, the mean/variance of each tissue class and a bias-corrected image.

The ability to automatically handle topology changes is a long claimed advantage of the level set method over explicit deformable models, but is often not desirable in biomedical image segmentation, where the topology of the target shape is prescribed by anatomical knowledge. A topology preserving variant of the level set method has been proposed in (Han et al., 2003) to overcome this problem: the level set function is evolved with a modified update procedure based on some concepts from digital topology, then the final mesh is obtained with a modified topology-consistent marching cubes algorithm. This method ensures that the resulting mesh has the same topology than the user-defined initial level set. We have extended this method in order to evolve several nested level sets while preventing mutual intersections. Contrarily to some methods for explicit deformable models based on repulsion forces, our method guarantees the absence of intersections and is computationnally cheaper than checking mesh-to-mesh intersection.

We apply successively the HMRF classification and active contour segmentation with the topology preserving level set method to benefit from the advantages of both methods while discarding their respective drawbacks. We first run the HMRF classification. This statistical approach is powerful as regards automatic parameter estimation but it is not sub-voxel accurate and disregards topology. Hence we feed the resulting output into an active contour segmentation task. We first fit a set of nested topology preserving level sets to the labeled image, in order to alleviate the sensitivity of active contour segmentation to the position of the initial contour. Then we drop the labels and we evolve the level sets according to the intensities of the bias-corrected image. Since the image inhomogeneities have been removed, the interfaces between the different tissues can be found easily and robustly by driving each level set with an adequate intensity threshold. Moreover, the thresholds that best separate the intensity distributions of the tissues can be computed from the tissue parameter estimated during the HMRF classification.

Our method is used routinely to obtain the surfaces that are needed in the BEM or FEM source localization techniques from

MEG and EEG. The cortical surface in Figure 1 was extracted with this segmentation tool. As an illustration of the technique, see Fig. 2.




Figure 2. The skin/air, CSF/GM and GM/WM interfaces of one of the authors' head.

Diffusion Tensor Estimation

The basic principles of Diffusion MRI and the formalism of the Diffusion Tensor have been introduced in (LeBihan et al., 1986; LeBihan and Breton, 1985; Basser et al., 1994). The key concept is that the random motion of water molecules, referred to as diffusion, reflects the structure of the underlying biological tissues at a microscopic scale, well beyond the usual image resolution. This opens the possibility of recovering a detailed geometric description of the anatomical connectivity between brain areas. We are attempting to get closer to this challenging goal.

The estimation of a 3×3 symmetric positive definite tensor at each voxel from diffusion weighted data uses the Stejskal-Tanner equation (Stejskal and Tanner, 1965). Many approaches have been derived to estimate the tensor. Minimal approaches, (Westin et al., 2002), are very sensitive to noise and outliers. Outlier-related artifacts can be combatted by using more measurements and robust estimators as in (Mangin et al., 2002). However the resulting tensors may not be positive definite. We have proposed in (Tschumperlé and Deriche, 2003) to incorporate such priors as tensor positivity and spatial regularity into a variational formulation of the estimation problem on the manifold of positive definite tensors.

Diffusion Tensor Regularization

We have studied the regularization of noisy diffusion tensor data. There exist two main classes of techniques. Non-spectral methods are based on a direct anisotropic smoothing of the diffusion weighted data (Vemuri et al., 2001) or consider each tensor as 6 scalar coordinates. This method suffers from the fact that the eigenvalues tend to diffuse faster than eigenvectors. Spectral methods process the eigen-elements of the tensors. Eigenvalue smoothing is typically performed by a vector-valued anisotropic PDE ((Tschumperlé and Deriche, 2001) and references therein) satisfying the maximum principle in order to preserve positiveness. However, these approaches are plagued by the fact that all vectors are defined up to a change of direction, resulting in a computational explosion (Coulon et al., 2001).

We recently proposed an alternative to the previous spec-

tral techniques, called the fast isospectral method (Chefd'hotel et al., 2004). It builds flows acting on a given sub-manifold of the linear space of matrix-valued functions and preserving some constraints. The constraints (orthogonality, invariance of eigenvalues...) can be expressed in terms of Lie groups and homogeneous spaces. Results of non-spectral smoothing and isospectral flow on diffusion tensors estimated in the genu of the corpus callosum are shown in Fig. 3.




Figure 3. Left to right: (a) Raw tensors in the genu of the corpus callosum and regularized fields by (b) a non-spectral method, (c) our isospectral flow. Data courtesy of CEA-SHFJ/Orsay. We thank J.F. Mangin and J.B. Poline for providing us with the data.

Fiber Tractography

The main idea most classical algorithms for brain connectivity mapping rely on ((Mori et al., 1999; Moseley et al., 1999; Pierpaoli et al., 1996) and references therein) is that water diffusion in many regions of the white matter is highly anisotropic and thus the orientation of the principal eigenvector is that of the predominant axonal direction. All these line propagation techniques however fail whenever they enter a region of low anisotropy, the estimate of the curve tangent becoming highly unreliable.

To overcome this problem we have introduced a physically motivated distance function in the white matter through stochastic processes and differential geometry. In our approach, the white matter is seen as a 3-manifold M and fibers become geodesics (Lenglet et al., 2004) of this manifold. There is a fascinating connection between the diffusion tensor, Brownian motion and the Eikonal equation. If C is the concentration of water molecules, it can be shown that it satisfies the equation

$$\frac{\partial C}{\partial t} = \nabla \cdot (\mathbf{D} \nabla C) = \mathcal{L}C \quad (2)$$

When \mathbf{D} is the identity, \mathcal{L} is the Laplace operator, the diffusion is isotropic. A Brownian motion in Euclidean space (e.g. \mathbb{R}^3) is entirely defined by its initial distribution and the conditional probability p of finding a molecule at x at time t given that it was at x_0 at time 0. For an unbounded anisotropic homogeneous medium, p is the minimal fundamental solution of (2).

These notions have their counterparts when moving from the Euclidean space to a Riemannian space M . In particular an isotropic diffusion on M is governed by a differential operator, the Laplace-Beltrami operator, which defines the geometry of M . In effect it is possible to consider an anisotropic diffusion in the Euclidean space (governed by the diffusion tensor \mathbf{D}) as an isotropic diffusion in \mathbb{R}^3 seen as a Riemannian space with

metric tensor $\mathbf{G} = (g_{ij})$ equal to \mathbf{D}^{-1} . It is therefore natural to consider the intrinsic distance d in the space of the brain white matter to any voxel x_0 as defined by this Riemannian metric. It verifies the intrinsic Eikonal equation

$$|\nabla d|^2 = \frac{\partial d}{\partial x_i} \frac{\partial d}{\partial x_j} g^{ij} = 1, \quad (3)$$

with $(g^{ij}) = \mathbf{G}^{-1}$. In (Lenglet et al., 2004), we propose a level-set formulation and the associated numerical scheme to solve (3). Computing neural fibers as geodesics in the region of the splenium of the corpus callosum yields the results presented in Fig. 4. The main advantage of this method over line propagation techniques is that it is not influenced by locally isotropic areas.




Figure 4. Inferred geodesics in the splenium of the corpus callosum (red: low anisotropy - blue: high anisotropy)

STATISTICAL MODELING

Superresolution in fMRI

In functional Magnetic Resonance Imaging (fMRI) a major goal is to maximize the image spatial resolution. The decrease in SNR induced by the decrease of voxel size can be obviated by the use of higher magnetic fields implying much higher equipment costs, an increased inhomogeneity and hence larger distortion artifacts in the images. To overcome these problems, one possible solution is to use superresolution techniques which allow us to generate a high resolution volume from a set of low resolution ones. These techniques have already been used in different image processing applications. For anatomical MR images, superresolution can be used in 2D FT MR imaging in the image space, i.e. in the slice direction, see (Greenspan et al., 2002; Peled and Yeshurun, 2001) and (Roullot et al., 2000) who proposed a superresolution approach for 3D volumes.

Our aim is to investigate how these techniques can be applied to process long fMRI sequence and improve the activation maps (Peeters et al., 2004; Kornprobst et al., 2003). The approach is based on two steps. The first is the acquisition protocol. Images are acquired at a low resolution using alternate shifts of the image slice stack over half-a-slice thickness, generating two separate slice-shifted overlapping volumes. The second

is a variational superresolution reconstruction technique which combines recent work on edge preserving PDEs (see (Aubert and Kornprobst, 2002) for a review) and convergence rate studies (Nikolova and Ng, 2001). Experiments on synthetic and real data (see Fig. 5) clearly establish the interest of using such techniques in this context.

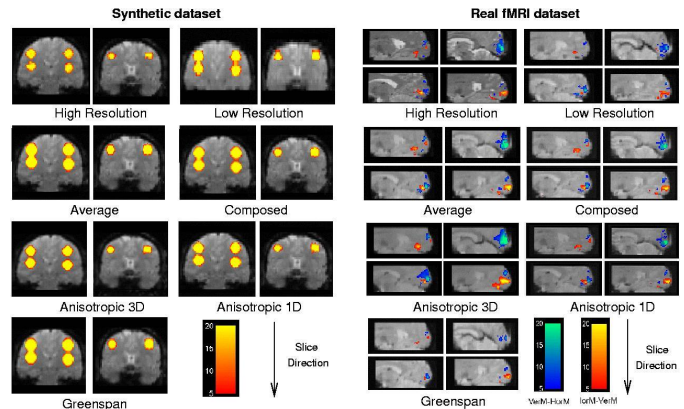


Figure 5. Comparisons between activated areas for different reconstructions in a synthetic case (Left) and a real visual experiment (Right). The latter shows the activated areas for the horizontal meridian (red-yellow) and the vertical meridian (blue-green). Real data were generated in the MR Research Center, Dept. of Radiology, K.U.Leuven, Medical School, Leuven, Belgium.

fMRI modeling

fMRI is a rich source of data that is hard to analyze without the help of models. This is because it is only an indirect measure of brain activity based on brain oxygenation and many confounders such as subject movements, respiratory and heart artifacts, temperature drift, machine noise, are known to corrupt the signal. The models can be weak as in the case of exploratory methods such as PCA, ICA or clustering, or strong as in the case of the General Linear Model (GLM) where the shape of the hemodynamic response is assumed to be known.

We have contributed to the latter by developing a class of flexible models that attempt to extract coherent, i.e. task-related or autocorrelated, patterns from each voxel time course (Thirion and Faugeras, 2003a). This is achieved by decomposing each voxel time course as the sum of a pseudo-deterministic term which captures all the information in the signal that is related to its past values or to the experimental paradigm and a stochastic part that is the non task-related part. The pseudo-deterministic part is obtained by applying the Minimum Description Length (MDL) method and can be very efficiently estimated. If we note $X^n(t)$ the fMRI time course at voxel n , this analysis results in the decomposition

$$X^n(t) = z^n(t) + \varepsilon^n(t),$$

z^n being the pseudo-deterministic part, ε^n the stochastic part. This univariate analysis is followed by a multivariate one that overcomes some of the difficulties encountered in PCA (resp.

ICA) (the assumption that structures of interest in the data are uncorrelated in the temporal and (resp. or) spatial domain, the fact that the experimental paradigm is not taken into account). This is achieved by defining a generalized covariance matrix κ whose eigenvectors represent the spatial modes of coherent dynamic activity in the data. κ is built from a kernel K_σ that depends upon the scale parameter σ . Two voxels X^i and X^j are represented by the signals z^i and z^j . The correlation $cc(z^i, z^j)$ can be thought of in terms of shared information between the time courses. σ penalizes the correlations the values of cc which are far from 1. To that effect we multiply the usual covariance matrix $K_\infty(z^i, z^j)$ by the function $\exp(\frac{cc(z^i, z^j) - 1}{\sigma})$. As described in (Thirion and Faugeras, 2003a), the number of eigenvectors kept in the final description can be determined by a variant of the MDL criterion.

We have also explored the use of clustering in the analysis of fMRI data. The problem of choosing the number of clusters can be addressed by the Information Bottleneck (IB) approach developed for vector quantization (Tishby et al., 1999) which deals explicitly with a tradeoff between quantization and data fidelity through an information theoretic formulation. Our approach assumes that $X^n(t)$ has been projected in the space spanned by R regressors g_r :

$$X^n(T) = \sum_{r=1}^R \gamma_r^n g_r(t) + \varepsilon^n(t),$$

for example by the GLM. The vectors γ^n are modelled as Gaussian random variables with known mean and covariance. The IB method can be formulated as follows. Given the set of voxels X , the set of interest Γ (the set of possible values for γ) and the normal densities $p(\Gamma|X = x)$, find the fuzzy clusters ξ that maximize compression while retaining most of the information on $p(\Gamma|X)$. This leads to the minimization with respect to ξ of the function $I(X, \xi) - \beta I(\xi, \Gamma)$, where $I(X, \xi)$ is the mutual information between the dataset X and its compressed representation, $I(\xi, \Gamma)$ is the mutual information between the compressed representation and the variable of interest, and β a positive scalar. The IB method finds its roots in statistical physics and β plays the role of an inverse temperature: a high value freezes the system into a hard clustering while a small one heats the system and ultimately fuses all clusters into a single one. The minimization can be done by an EM algorithm. The results are sensitive to the choice of β but this choice can be made in a principled manner as explained in (Thirion and Faugeras, 2003b). We illustrate the method with a synthetic example. We simulate one slice of fMRI data with $N = 1963$ voxels and 3 foci of 21 pixels. Independent Gaussian noise is added to all voxels so that the SNR is 0.5 in the activated areas. The simulated paradigm has two conditions, the simulated time courses and the spatial maps are shown in Fig. 6(a), (b). By keeping only the effects of interest, we obtain a 2-dimensional feature-space. The estimated feature values at each voxel are represented in Fig. 6(c). The IB method yields four clusters.

The corresponding pdfs $p(\gamma|\xi)$ are shown in Fig. 6(d) For com-

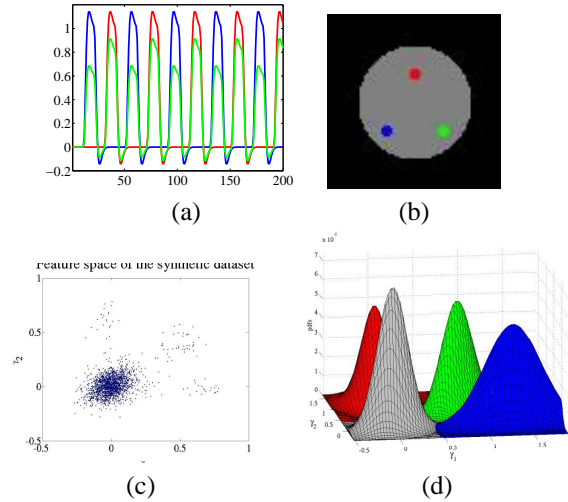


Figure 6. a): Simulated time courses for the three foci shown in b). c): Feature values showing one big cluster and two big ones. d) Pdfs produced by the IB method.

parison, we have applied a C-Means and a fuzzy-C-Means algorithm with an initial number of clusters equal to 4 and never obtained the results shown in Fig. 6(d). More statistical approaches for modeling fMRI data can be found in the thesis (Thirion, 2003) which is written in English.

Cortical filtering and applications to retinotopy

Delineating the cortical visual areas is a "calibration" experiment for any study of the human cortical visual system. Beyond estimating the areas borders, the retinotopic mapping process provides information linking the actual visual field and the cortical surface. Our method derives from previous work (Serenio et al., 1995). It is fast -the whole retinotopic map being acquired in about 15 mn of functional images acquisition-, semi-automatic and shows an excellent intra- and inter-subject, reproducibility, as shown in (Wotawa et al., 2003). The stimuli are a wedge, coding for polar angle maps and a ring, coding for eccentricity maps. The functional analysis scheme is based on a frequency analysis, which allows a continuous and accurate response phase estimation. The only requirement on the hemodynamic response is that it be linear with respect to the stimulation. The main steps are the computation of a statistical map to determine the voxels activated by our periodic stimuli, followed by a phase estimation, including a voxel-based hemodynamic delay estimation. The phase links the voxel activity to its preferred stimulus position. The results, values of eccentricity and polar angle at each suprathreshold voxel, are then projected on the cortical surface model derived from a high resolution anatomical image. The model is finally inflated to facilitate the visualization, see Fig. 7. Part of the processing for obtaining these results is spent on smoothing the fMRI data. The way this is done is important. The classical approach, for




Figure 7. Retinotopic polar angle and eccentricity maps projected on an inflated left hemisphere.

instance in SPM, is to smooth the whole volume of data by convolving it with a 3D Gaussian kernel. This may have two undesirable effects: mixing voxels from different anatomical tissues, affecting the analysis sensitivity and blending signals across sulci, reducing the spatial discrimination power. In order to avoid this, we have used such interfaces as the one shown in Fig. 2 and replaced the 3D filtering with a 2D filtering on the surface of the cortex, another example of the Laplace-Beltrami operator. The difference between the two methods is shown in Fig. 8.




Figure 8. Left: original activation; Middle: 3mm 3D isotropic smoothing, leading to false activation on the opposite bank of the sulcus; Right: 3mm Laplace-Beltrami smoothing.

Multimodal image matching

One of the problems that is often encountered in the analysis of fMRI data is that of registering it with the anatomy. This is difficult because, as shown in Fig. 9, there are significant geometric distortions between the two volumes and the intensities of corresponding areas rarely match (not visible in this figure). We discuss these two points next. The geometric distortion cannot in general be described by a simple global transformation rigid, i.e. rotation plus translation, or non-rigid, i.e. affine ((Ayache, 2003) and references therein). In order to model the distortion one is therefore led to call upon some general deformation flow represented by a largely arbitrary vector field \mathbf{h} . The different intensities in corresponding areas preclude the use of the usual sum of squared differences (SSD) as an error criterion to guide the search for the deformation flow. One must revert to more sophisticated, i.e. statistical, ways of measuring similarities in intensity distribution. Such measures as the cross-correlation, the correlation ratio, or the mutual infor-

mation have been successfully used in the literature (Viola and Wells III, 1997).

Our contribution has been to clearly pose the problem of the estimation of \mathbf{h} as that of minimizing an energy functional $\mathcal{I}(\mathbf{h})$ on a well-defined functional space \mathcal{F} . The energy functional is the sum of dissimilarity term $\mathcal{J}(\mathbf{h})$ and a regularization term $\mathcal{R}(\mathbf{h})$. The first term is based upon the idea of modeling the two images $I_1(\mathbf{x})$ and $I_2(\mathbf{x})$ as samples of two random processes and of estimating the joint probability density function (pdf) of the vector $(I_1(\mathbf{x}), I_2(\mathbf{h}(\mathbf{x})))$. From this pdf one can then compute any of the above statistical measures as functions of the field \mathbf{h} . The regularization criterion takes into account the idea that this field cannot vary arbitrarily and therefore enforces some regularity. This is done by choosing $\mathcal{R}(\mathbf{h})$ to be a function of the first order derivative of \mathbf{h} . The next step is to precisely define the functional space \mathbf{h} belongs to. It turns out that the form of the regularization term is determinant and implies that we work in the Sobolev space $\mathcal{F} = H_0^1 \cap H^2$. After showing that there exist minimizers of $\mathcal{I}(\mathbf{h})$ in \mathcal{F} , we turn the problem of finding them into one of solving a semilinear abstract initial value problem that can be written as

$$\frac{d\mathbf{h}}{dt} - A\mathbf{h}(t) = F(\mathbf{h}(t)), \quad t > 0, \quad \mathbf{h}(0) = \mathbf{h}_0 \in H. \quad (4)$$

In this equation, the time has been introduced to reflect the fact that we start from an initial deformation flow \mathbf{h}_0 and look for the corresponding stationary solution of (4). A is a spatial differential operator arising from the Euler-Lagrange equation of the regularization term $\mathcal{R}(\mathbf{h})$. The function F in the right-hand side arises from the dissimilarity term $\mathcal{J}(\mathbf{h})$. Because $\mathcal{J}(\mathbf{h})$ involves the pdf of the vector $(I_1(\mathbf{x}), I_2(\mathbf{h}(\mathbf{x})))$ F is non local, i.e. its value at \mathbf{x} depends upon the values of the current deformation flow $\mathbf{h}(t)$ at other points. This implies that (4) is not a PDE but a functional equation, in effect an ordinary differential equation (ODE) in the unknown \mathbf{h} that lives in the functional space \mathcal{F} . This makes the analysis of the registration problem significantly more difficult than in the case of the SSD criterion where (4) is a PDE.

Within this general framework, we have a) shown that there exist minimizers of $\mathcal{I}(\mathbf{h})$ in \mathcal{F} , b) computed the function F for the previous statistical criteria c) proved the existence and uniqueness of a solution of the problem (4) using the theory of analytical semigroups of operators and d) proved that the limit when $t \rightarrow \infty$ of the solution of (4) satisfies the Euler-Lagrange equations of $\mathcal{I}(\mathbf{h})$. The implementation of the method has been done in the context of a recent thesis (Hermosillo, 2002). The theoretical results are in (Faugeras and Hermosillo, 2004), with preliminary work reported in (Hermosillo et al., 2002). An example of the kind of results that can be achieved is shown in Fig. 9. The code is used routinely by Prof. Guy Orban's group in Leuven in their work on monkey fMRI (Fize et al., 2003).

Statistical shapes

The variability of anatomical structures from individual to individual is quite large within the human species. It is even




Figure 9. Left: The geometric distortion between anatomical MR (purple) and functional MR data (green). Right: The geometric distortion shown left has been mostly compensated for.

larger when one attempts to compare anatomy and function in different but closely related species, e.g. human and non-human primates. This variability hints at two main areas of mathematics, topology for defining meaningful shape metrics, and statistics for defining meaningful notions of shape variability, e.g. with respect to an average shape. The Odyssée laboratory has been actively pursuing both objectives. In the area of shape metrics we have been studying a set of shapes \mathcal{S} which are defined as subsets Ω of a set of \mathbb{R}^n (in practice $n = 2$ or 3) with a regular (i.e. C^2) boundary $\partial\Omega$ whose curvature is upper-bounded by a positive number $1/h_0$. h_0 is also a lower-bound on the pinch distance of this boundary and is lower than the distance between two pixels of the grid on which the shapes are defined. The question of measuring the similarity between two shapes in \mathcal{S} builds upon the seminal work of (Delfour and Zolésio, 2001) who have introduced new families of sets, complete metric topologies and compactness theorems. We prove in (Charpiat et al., 2004) that three of the most important metrics, including the Hausdorff distance, are topologically equivalent in \mathcal{S} . We also propose to use them to define a way to continuously deform, i.e. warp, a shape onto another one. The way to pose this problem is to define a "dissimilarity" measure $E(\Gamma_1, \Gamma_2)$ between two shapes Γ_1 and Γ_2 , to show that the gradient $\nabla E(\Gamma_1, \Gamma_2)$ can be defined in a reasonable manner, and to solve the following initial value problem

$$\frac{d\Gamma}{dt} = -\nabla E(\Gamma, \Gamma_2)\mathbf{n}, \quad \Gamma(0) = \Gamma_1 \in \mathcal{S}, \quad (5)$$

Note the similarity between (5) and (4). One difficulty with this approach is that the metric appears in the definition of E . Therefore the gradient ∇E is not well-defined because the metric is not differentiable. We have therefore constructed classes of smooth (i.e. whose gradient is well-defined) approximations of the metric based upon the idea of replacing the sup and min operators that arise in the definition of the distance function to a set and in the Hausdorff distance between two sets by averages taken on the boundaries of the shapes. We prove that these approximations can approximate the metric arbitrarily well and compute the gradient of the corresponding dissimilarity measures. This defines a warping algorithm between two

shapes that can be seen as an infinitesimal gradient descent in order to minimize E . We prove that there exist minimizers of $E(\Gamma, \Gamma_2)$. This approach can be seen as the opposite of that consisting in first building a Riemannian structure on the set of shapes, i.e. going from an infinitesimal metric structure to a global one. This is mostly dealt with in (Miller and Younes, 2001; Trounevé, 1998; Younes, 2003; Klassen et al., 2004). The problem with these approaches, beside that of having to deal with parametrizations of the shapes, a difficult problem that is avoided in ours, is that there exist global metric structures on the set of shapes (like those we have considered) which are useful and relevant to the problem of the comparison of shapes but that do not arise from an infinitesimal structure.

Equation (5) defines a generic "shape warper" that can be used to address the second objective above, i.e. the definition of the empirical mean and covariance of a set of shape examples. The empirical mean of N shapes $\Gamma_1, \dots, \Gamma_N$ is defined as any shape $\hat{\Gamma}$ that achieves a local minimum of the function $\mu : \mathcal{S} \rightarrow \mathbb{R}^+$ defined by

$$\Gamma \rightarrow \mu(\Gamma, \Gamma_1, \dots, \Gamma_N) = \frac{1}{N} \sum_{i=1, \dots, N} E^2(\Gamma, \Gamma_i),$$

and we prove that there exists at least one mean. An algorithm for the effective computation of a mean is proposed in (Charpiat et al., 2004) and an example of the mean of eight silhouettes of corpus callosum is shown in Fig. 10 (left). The empirical co-




Figure 10. Left: The mean of eight silhouettes of corpus callosum (middle, thick line). Right: From top to bottom, the first three principal modes of variation for the eight sample shapes. They are the solutions of equation (6) for $k = 1, 2, 3$.

variance of N shapes is slightly more difficult to define. From a mean $\hat{\Gamma}$ we compute the gradients $\nabla E(\hat{\Gamma}, \Gamma_i)$, $i = 1, \dots, N$. These are functions defined on $\hat{\Gamma}$ which we use to define a positive symmetric $N \times N$ matrix which supports our notion of empirical covariance. Its eigenvectors and eigenvalues are used to define the analog of the principal modes v_k , $k = 1, \dots, N$ of variation of the mean shape $\hat{\Gamma}$. The variability of the mean shape with respect to the k th mode is obtained by solving the following initial value problem

$$\frac{d\Gamma}{dt} = \pm v_k \mathbf{n}, \quad \Gamma(0) = \hat{\Gamma} \in \mathcal{S}. \quad (6)$$

As an example, the first mode of variation for the above eight sample shapes is shown in Fig. 10(right).

CONCLUSION

We have shown that some large pieces of fairly sophisticated mathematics are very useful for modeling the types of signals that are currently used for observing the brain "in vivo". We are presently exploring two fascinating areas. One is the combination of these modalities (MEEG, MRI) into a more robust and more accurate meta-sensor. Another one is the introduction of models of the activity of the assemblies of neurons that the sensory modalities are trying to measure in the processing of the signals they deliver.

ACKNOWLEDGEMENTS

The authors are grateful to S. Meunier of the Physiology and Physiopathology of Human Motricity Lab., La Salpêtrière Hospital, INSERM, Paris for permission to use the somatosensory MEG data. They thank L. Garnero, S. Baillet and B. Renault from the Cognitive Neuroscience and Brain Imaging Lab., CNRS, Paris for providing helpful suggestions and software tools for accessing the MEG data. They also thank Profs. P. Van Hecke and G. Orban from the MR Research Center, Dept. of Radiology, K.U. Leuven, Medical School, Leuven, Belgium for providing them with MR data.

References

- Adde, G., Clerc, M., Faugeras, O., Keriven, R., Kybic, J., and Papadopoulo, T. 2003. Symmetric BEM formulation for the M/EEG forward problem. In C. Taylor and J. A. Noble (eds.), *Information Processing in Medical Imaging*, volume 2732 of *LNCS*, pp. 524–535. Springer.
- Amir, A. 1994. Uniqueness of the generators of brain evoked potential maps. *IEEE Transactions on Biomedical Engineering* 41.
- Aubert, G. and Kornprobst, P. 2002. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147 of *Applied Mathematical Sciences*. Springer-Verlag.
- Ayache, N. 2003. Epidaure: a research project in medical image analysis, simulation and robotics at INRIA. *IEEE Trans. on Medical Imaging* 22:1185–1201.
- Basser, P., Mattiello, J., and LeBihan, D. 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance B*:247–254.
- Charpiat, G., Faugeras, O., and Keriven, R. 2004. Approximations of shape metrics and application to shape warping and empirical shape statistics. *Foundations of Computational Mathematics*. Accepted for publication.
- Chefd'hotel, C., Tschumperlé, D., Deriche, R., and Faugeras, O. 2004. Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision* 20:147–162.
- Clerc, M., Keriven, R., Faugeras, O., Kybic, J., and Papadopoulo, T. 2002. The fast multipole method for the direct E/MEG problem. In *Proceedings of ISBI*, Washington, D.C. IEEE, NIH.
- Coulon, O., Alexander, D., and Arridge, S. 2001. A regularization scheme for diffusion tensor magnetic resonance images. In *XVIIth International Conference on Information Processing in Medical Imaging*.
- David, O. and Garnero, L. 2002. Time-coherent expansion of MEG/EEG cortical sources. *NeuroImage* 17:1277–89.
- Delfour, M. and Zolésio, J.-P. 2001. *Shapes and geometries. Advances in Design and Control*. Siam.
- Faugeras, O., Clément, F., Deriche, R., Keriven, R., Papadopoulo, T., Roberts, J., Viéville, T., Devernay, F., Gomes, J., Hermosillo, G., Kornprobst, P., and Lingrand, D. 1999. The inverse EEG and MEG problems: The adjoint space approach I: The continuous case. Technical Report 3673, INRIA.
- Faugeras, O. and Hermosillo, G. 2004. Well-posedness of two non-rigid multimodal image registration methods. *Siam Journal of Applied Mathematics*. To appear.
- Fize, D., Vanduffel, W., Nelissen, K., Denys, K., Chef d'Hotel, C., Faugeras, O., and Orban, G. A. 2003. The retinotopic organization of primate dorsal V4 and surrounding areas: A functional magnetic resonance imaging study in awake monkeys. *Journal of Neuroscience* 23:7395–7406.
- Geselowitz, D. B. 1967. On bioelectric potentials in an homogeneous volume conductor. *Biophysics Journal* 7:1–11.
- Greenspan, H., Oz, G., Kiryati, N., and Peled, S. 2002. MRI inter-slice reconstruction using super-resolution. *Magn. Res. Imag.* 20:437–446.
- Han, X., Xu, C., and Prince, J. 2003. A topology preserving level set method for geometric deformable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25:755–768.
- Hermosillo, G. 2002. *Variational Methods for Multimodal Image Matching*. PhD thesis, INRIA, The document is accessible at <ftp://ftp-sop.inria.fr/robotvis/html/Papers/hermosillo:02.ps.gz>.
- Hermosillo, G., Chefd'hotel, C., and Faugeras, O. 2002. Variational methods for multimodal image matching. *The International Journal of Computer Vision* 50:329–343.
- Klassen, E., Srivastava, A., Mio, W., and Joshi, S. 2004. Analysis of planar shapes using geodesic paths on shape spaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26:372–383.
- Kornprobst, P., Peeters, R., Nikolova, M., Deriche, R., Ng, M., and Hecke, P. V. 2003. A superresolution framework for fmri sequences and its impact on resulting activation maps. In *Medical Image Computing and Computer-Assisted Intervention-MICCAI2003*, volume 2 of *Lecture Notes in Computer Science*, pp. 117–125. Springer-Verlag.
- Kybic, J., Clerc, M., Abboud, T., Faugeras, O., Keriven, R., and Papadopoulo, T. 2003. Integral formulations for the eeg

- problem. Technical Report 4735, INRIA.
- LeBihan, D. and Breton, E. 1985. Imagerie de diffusion *in vivo* par résonance magnétique nucléaire. *CR Académie des Sciences* pp. 1109–1112.
- LeBihan, D., Breton, E., Lallemand, D., Grenier, P., Cabanis, E., and Laval-Jeantet, M. 1986. MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology* pp. 401–407.
- Lenglet, C., Deriche, R., and Faugeras, O. 2004. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In T. Pajdla and J. Matas (eds.), Proceedings of the 8th European Conference on Computer Vision, Prague, Czech Republic. Springer-Verlag.
- Mangin, J., Poupon, C., Clark, C., and Le Bihan, D. and Bloch, I. 2002. Distortion correction and robust tensor estimation for mr diffusion imaging. *Med Image Anal* 6:191–8.
- Meunier, S., Garnero, L., Ducorps, A., Mazières, L., Lehéricy, S., Tezenas Du Montcel, S., Renault, B., and Vidailhet, M. 2001. Human brain mapping in dystonia reveals both endophenotypic traits and adaptative reorganization. *Annals of Neurology* 50:521–527.
- Miller, M. and Younes, L. 2001. Group actions, homeomorphisms, and matching : A general framework. *International Journal of Computer Vision* 41:61–84.
- Mori, S., Crain, B., Chacko, V., and Zijl, P. V. 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45:265–269.
- Moseley, M., Cohen, Y., Kucharczyk, J., Mintorovitch, J., Asgari, H., Wendland, M., Tsuruda, J., and Norman, D. 1999. Diffusion-weighted mr imaging of anisotropic water diffusion in cat central nervous system. *Radiology* 176:439–445.
- Nikolova, M. and Ng, M. 2001. Fast image reconstruction algorithms combining half-quadratic regularization and preconditioning. In Proceedings of the International Conference on Image Processing, volume 1, pp. 277–280. IEEE Signal Processing Society.
- Peeters, R., Kornprobst, P., Nikolova, M., Sunaert, S., Vieville, T., Malandain, G., Deriche, R., Faugeras, O., Ng, M., and Hecke, P. V. 2004. The use of superresolution techniques to reduce slice thickness in functional MRI. *International Journal of Imaging Systems and Technology (IJIST), Special issue on High Resolution Image Reconstruction* . To appear.
- Peled, S. and Yeshurun, Y. 2001. Superresolution in MRI : Application to human white matter fiber tract visualization by diffusion tensor imaging. *Magn. Reson. Med.* 45:29–35.
- Pierpaoli, C., Jezzard, P., Basser, P., Barnett, A., and Chiro, G. D. 1996. Diffusion tensor mr imaging of human brain. *Radiology* 201:637–648.
- Roullot, E., Herment, A., Bloch, I., Nikolova, M., and Mousseaux, E. 2000. Regularized reconstruction of 3D high-resolution magnetic resonance images from acquisitions of anisotropically degraded resolutions. In Proceedings of the International Conference on Image Processing, volume 1, pp. 350–353. IEEE Signal Processing Society.
- Sereno, M., Dale, A., Reppas, J., Kwong, K., Belliveau, J., Brady, T., Rosen, B., and Tootell, R. 1995. Borders of multiple visual areas in human revealed by functional magnetic resonance imaging. *Science* pp. 889–893.
- Stejskal, E. and Tanner, J. 1965. Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics* 42:288–292.
- Thirion, B. 2003. Analyse de données d’IRM fonctionnelle: statistiques, information et dynamique. PhD thesis, École Doctorale Télécom Paris.
- Thirion, B. and Faugeras, O. 2003a. Dynamical components analysis of fMRI data through kernel PCA. *NeuroImage* 20:34–49.
- Thirion, B. and Faugeras, O. 2003b. Feature detection in fMRI data: The information bottleneck approach. In MICCAI 2003, volume 2878 of LNCS, pp. 83–91. Springer-Verlag.
- Tishby, N., Pereira, F. C., and Bialek, W. 1999. The Information Bottleneck method. In Proc. of the 37-th Annual Allerton Conference on Communication, Control and Computing, pp. 368–377.
- Trouvé, A. 1998. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision* 28:213–21.
- Tschumperlé, D. and Deriche, R. 2001. Diffusion tensor regularization with constraints preservation. In IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Kauai Marriott, Hawaii.
- Tschumperlé, D. and Deriche, R. 2003. Variational frameworks for DT-MRI estimation, regularization and visualization. In Proceedings of the 9th International Conference on Computer Vision, Nice, France. IEEE Computer Society, IEEE Computer Society Press.
- Vemuri, B., Chen, Y., Rao, M., McGraw, T., Mareci, T., and Wang, Z. 2001. Fiber tract mapping from diffusion tensor mri. In 1st IEEE Workshop on Variational and Level Set Methods in Computer Vision (VLSM’01).
- Viola, P. and Wells III, W. M. 1997. Alignment by maximization of mutual information. *The International Journal of Computer Vision* 24:137–154.
- Westin, C., Maier, S., Mamata, H., Nabavi, A., Jolesz, F., and Kikinis, R. 2002. Processing and visualization for diffusion tensor MRI. In In proceedings of Medical Image Analysis, volume 6(2), pp. 93–108.
- Wotawa, N., Thirion, B., Castet, E., Anton, J.-L., and Faugeras, O. 2003. Efficient human retinotopic mapping using fMRI. In T. Paus, E. Bullmore, and J. D. Cohen (eds.), NeuroImage (HBM’03), New York, USA. Academic Press.
- Younes, L. 2003. Invariance, déformations et reconnaissance de formes. Mathématiques et Applications. Springer-Verlag.
- Zhang, Y., Brady, M., and Smith, S. 2001. Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging* 20.