Riemannian Elasticity: a statistical regularization framework for non-linear registration

X. Pennec[1], R. Stefanescu[1], V. Arsigny[1], P. Fillard[1], and N. Ayache[1]

INRIA Sophia - Projet Epidaure, 2004 Route des Lucioles BP 93
06902 Sophia Antipolis Cedex, France
Xavier.Pennec@sophia.inria.fr

Abstract. In inter-subject registration, one often lack a good model of the transformation variability to choose the optimal regularization. Some works attempt to model the variability in a statistical way, but the re-introduction in a registration algorithm is not easy. In this paper, we interpret the elastic energy as the distance of the Green-St Venant strain tensor to the identity, which reflects the deviation of the local deformation from a rigid transformation. By changing the Euclidean metric for a more suitable Riemannian one, we define a consistent statistical framework to quantify the amount of deformation. In particular, the mean and the covariance matrix of the strain tensor can be consistently and efficiently computed from a population of non-linear transformations. These statistics are then used as parameters in a Mahalanobis distance to measure the statistical deviation from the observed variability, giving a new regularization criterion that we called the statistical Riemannian elasticity. This new criterion is able to handle anisotropic deformations and is inverse-consistent. Preliminary results show that it can be quite easily implemented in a non-rigid registration algorithms.

1 Introduction

Most non-linear image registration algorithms optimize a criterion including an image intensity similarity and a regularization term. In inter-subject registration, the main problem is not really the intensity similarity measure but rather the regularization criterion. Some authors used physical models like elasticity or fluid models [1, 2]. For efficiency reasons, other authors proposed to use non-physical but efficient regularization methods like Gaussian filtering [3, 4], recently extended to non-stationary but still isotropic diffusion in order to take into account some anatomical information about the tissue types [5, 6]. However, since we do not have in general a model of the deformation of organs across subjects, no regularization criterion is obviously more justified than the others. We could think of building a model of the developing organ: inverting the model from the first subject to a sufficiently early stage and growing toward the second subject image would allow to relate the two anatomies. However, such a computational model is out of reach now, and most of the existing work in the literature rather try to capture the organ variability from a statistical point of view on a representative population of subjects (see e.g. [7–9]). Although the image databases are now large enough to be representative of the organ variability, the problem remains of how to use this information to better guide inter-subject registration.

We propose in this paper an integrated framework to compute the statistics on deformations and reintroduce them in the registration procedure. The basic idea is to interpret the elastic energy as a distance in the space of positive definite symmetric matrices (tensors). By changing the classical Euclidean metric for a more suitable one, we define a natural framework for computing statistics on the strain tensor. Taking them into account in a statistical distance lead to the Riemannian elasticity energy. Notice that we do not enter the fluid vs elastic registration debate as the energy we propose can be used either on the deformation field itself or on its temporal derivative (fluid-like method) [4].

In the sequel, we first recall how to optimize the elastic energy in a registration algorithm. Then, we define in Section 3 the Riemannian elasticity energy as the Mahalanobis distance on the logarithmic strain tensor. To better exemplify its properties, we investigate in Section 4 the isotropic Riemannian Elasticity, which is close to the classical elasticity energy while being inverse-consistent. Preliminary experiments show in Sec. 5 that the Riemannian elasticity framework can be implemented quite effectively and yields promising results.

2 Non-linear elastic regularization

Let $\Phi(x)$ be a non-linear space transformation with a positive Jacobian everywhere. We denote by $\partial_\alpha \Phi$ the directional derivatives of the transformation along the spaces axis α (we assume an orthonormal basis). The general registration method is to optimize an energy of the type: $C(\Phi) = Sim(\text{Images}, \Phi) + Reg(\Phi)$. Starting from an initial transformation Φ_0, a first order gradient descent methods computes the gradient of the energy $\nabla C(\Phi)$, and update the transformation using: $\Phi_{t+1} = \Phi_t - \eta\, \nabla C(\Phi_t)$. From a computational points of view, this Lagrangian framework can be advantageously changed into a Eulerian framework to better conserve the diffeomorphic nature of the mappings [6]. In the following, we only focus on the computation of the gradient of the regularization.

2.1 Elastic deformations

In continuum mechanics [10], one characterizes the deformation of an infinitesimal volume element in the Lagrangian framework using the Cauchy-Green tensor $\Sigma = \nabla\Phi^{\mathrm{T}}\, \nabla\Phi = \sum_\alpha \partial_\alpha \Phi\, \partial_\alpha \Phi^{\mathrm{T}}$. This symmetric matrix is positive definite and measures the local amount of non-rigidity. Let $\nabla\Phi = V\ S\ R^{\mathrm{T}}$ be a singular value decomposition of the transformation Jacobian (R and V are two rotation matrices and S is the diagonal matrix of the positive singular values). The Cauchy-Green tensor $\Sigma = R\, S^2\, R^{\mathrm{T}}$ is equal to the identity if and only if the transformation is locally a rigid transformation. Eigenvalues between 0 and 1 indicate a local compression of the material along the associated eigenvector, while a value above 1 indicates an expansion. To quantify the deformation, one usually prefers the related Green-St Venant strain tensor $E = \frac{1}{2}(\Sigma - \mathrm{Id})$, whose eigenvalues are null for no deformation. Assuming an isotropic material and a linear Hooks law to relate strain and stress tensors, one can show that the motion equations derive from the *St Venant-Kirchoff elasticity* energy [10]:

$$Reg_{SVKE}(\Phi) = \int \mu \mathrm{Tr}(E^2) + \frac{\lambda}{2}\mathrm{Tr}(E)^2 = \int \frac{\mu}{4}\mathrm{Tr}\left((\Sigma - \mathrm{Id})^2\right) + \frac{\lambda}{8}\mathrm{Tr}(\Sigma - \mathrm{Id})^2$$

where λ, μ are the Lamé coefficients. To minimize this energy in a registration algorithm, we need its gradient. Since $\partial_u \Sigma = \sum_\alpha \left(\partial_\alpha \Phi\, \partial_\alpha u^{\mathrm{T}} + \partial_\alpha u\, \partial_\alpha \Phi^{\mathrm{T}}\right)$, the derivative of the elastic energy in the direction (i.e. displacement field) u is:

$$\begin{aligned} \partial_u Reg_{SVKE}(\Phi) &= \textstyle\int \frac{\mu}{2}\mathrm{Tr}((\Sigma - \mathrm{Id})\, \partial_u \Sigma) + \frac{\lambda}{4}\mathrm{Tr}(\Sigma - \mathrm{Id})\, \mathrm{Tr}(\partial_u \Sigma) \\ &= \textstyle\sum_\alpha \int \langle\, \mu\, (\Sigma - \mathrm{Id})\, \partial_\alpha \Phi \mid \partial_\alpha u \rangle + \frac{\lambda}{2}\mathrm{Tr}(\Sigma - \mathrm{Id})\ \langle\, \partial_\alpha \Phi \mid \partial_\alpha u \rangle \end{aligned}$$

Using an integration by part with homogeneous Neumann boundary conditions [4], we have $\int \langle\, v \mid \partial_\alpha u \rangle = -\int \langle\, \partial_\alpha v \mid u \rangle$, so that the gradient is finally:

$$\nabla Reg_{SVKE}(\Phi) = -\textstyle\sum_\alpha \partial_\alpha \big(Z \partial_\alpha \Phi\big) \quad \text{with} \quad Z = \mu(\Sigma - \mathrm{Id}) + \frac{\lambda}{2}\mathrm{Tr}(\Sigma - \mathrm{Id})\,\mathrm{Id} \quad (1)$$

2.2 Practical implementation

In practice, a simple implementation is the following. First, one computes the image of the gradient of the transformation, for instance using finite differences. This operation is not computationally expensive, but requires to access the value of the transformation field at neighboring points, which can be time consuming due to systematic memory page faults in large images. Then, we process these 3 vectors completely locally to compute 3 new vectors $v_\alpha = Z(\partial_\alpha \Phi)$. This operation is computationally more expensive but is memory efficient as the resulting vectors can replace the old directional derivatives. Finally, the gradient of the criterion $\nabla E = \sum_\alpha \partial_\alpha v_\alpha$ may be computed using finite differences on the resulting image. Once again, this is not computationally expensive, but it requires intensive memory accesses.

3 Riemannian Elasticity

In the standard Elasticity theory, the deviation of the positive definite symmetric matrix Σ (the strain tensor) from the identity (the rigidity) is measured using the Euclidean matrix distance $\mathrm{dist}^2_{Eucl}(\Sigma, \mathrm{Id}) = \mathrm{Tr}((\Sigma - \mathrm{Id})^2)$. However, it has been argued in recent works that the Euclidean metric is not a good metric for the tensor space because positive definite symmetric matrices only constitute a cone in the Euclidean matrix space. Thus, the tensor space is not complete (null or negative eigenvalues are at a finite distance). For instance, an expansion of a factor $\sqrt{2}$ in each direction (leading to $\Sigma = 2\,\mathrm{Id}$) is at the same Euclidean distance from the identity than the "black hole" transformation $\Phi(x) = 0$ (which has a non physical null strain tensor). In non-linear registration, this asymmetry of the regularization leads to different results if we look for the forward or the backward transformation: this is the inverse-consistency problem [11].

3.1 A Log-Euclidean metric on the strain tensor

To solve the problems of the Euclidean tensor computing, affine-invariant Riemannian metrics were recently proposed [12–15]. Using these metrics, symmetric matrices with null eigenvalues are basically at an infinite distance from any tensor, and the notion of mean value corresponds to a geometric mean, even if it has to be computed iteratively. More recently, [16] proposed the so-called Log-Euclidean metrics, which exhibit the same properties while being much easier to compute. As these metrics simply consist in taking a standard Euclidean metric after a (matrix) logarithm, we rely on this one for the current article. However, the Riemannian Elasticity principle can be generalized to any Riemannian metric on the tensor space without any restriction.

In this framework, the deviation between the tensor Σ and the identity is the tangent vector $\log(\Sigma)$. Interestingly, this tensor is known in continuum mechanics as the logarithmic or Hencky strain tensor and is used for modeling very large deformations. It is considered as the natural strain tensor for many materials, but its use was hampered for a long time because of its computational complexity [17]. For registration, the basic idea is to replace the elastic energy with a regularization that measures the amount of logarithmic strain by taking the Riemannian distance between Σ and Id. This give the *Riemannian elasticity*:

$$Reg_{RE}(\Phi) = \frac{1}{4} \int \mathrm{dist}^2_{Log}(\Sigma, \mathrm{Id}) = \frac{1}{4} \int \|\log(\Sigma) - \log(\mathrm{Id})\|_2^2 = \frac{1}{4} \int \mathrm{Tr}\left(\log(\Sigma)^2\right)$$

It is worth noticing that the logarithmic distance is inverse-consistent, i.e. that $\mathrm{Tr}\left(\log(\Sigma(\Phi(x)))^2\right) = \mathrm{Tr}\left(\log(\Sigma(\Phi^{(-1)}(y)))^2\right)$ if $y = \Phi(x)$. This comes from the fact that $\nabla(\Phi^{(-1)})(y) = (\nabla\Phi(x))^{(-1)}$. In particular, a scaling of a factor 2 is now at the same distance from the identity than a scaling of 0.5, and the "black hole" transformation is at an infinite distance from any acceptable deformation.

3.2 Incorporating deformation statistics

To incorporate statistics in this framework, we consider the strain tensor as a random variable in the Riemannian space of tensors. In the context of inter-subject or atlas-to-image registration, this statistical point of view is particularly well adapted since we do not know a priori the deformability of the material. Starting from a population of transformations $\Phi_i(x)$, we define the *a priori* deformability $\bar{\Sigma}(x)$ as the Riemannian mean of deformation tensors $\Sigma_i(x) = \nabla\Phi_i^{\mathrm{T}}\, \nabla\Phi_i$. A related idea was suggested directly on the Jacobian matrix of the transformation $\nabla\Phi$ in [18], but using a general matrix instead of a symmetric one raises important computational and theoretical problems. With the Log-Euclidean metric on strain tensors, the statistics are quite simple since we have a closed form for the mean value:

$$\bar{\Sigma}(x) = \exp(\bar{W}(x)) \quad \text{with} \quad \bar{W}(x) = \frac{1}{N} \sum_i \log(\Sigma_i(x))$$

Going one step further, we can compute the covariance matrix of the random process $\mathrm{Cov}(\Sigma_i(x))$. Let us decompose the symmetric tensor $W = \log(\Sigma)$ into a vector $\mathrm{Vect}(W)^{\mathrm{T}} = (w_{11}, w_{22}, w_{33}, \sqrt{2}w_{12}, \sqrt{2}w_{13}, \sqrt{2}w_{23})$ that gathers all the tensor components in an orthonormal basis. In this coordinate system, we can define the covariance matrix $\mathrm{Cov} = \frac{1}{N}\sum \mathrm{Vect}(W_i - \bar{W})\,\mathrm{Vect}(W_i - \bar{W})^{\mathrm{T}}$.

To take into account these first and second order moments of the random deformation process, a well known and simple tool is the Mahalanobis distance, so that we finally define the *statistical Riemannian elasticity (SRE)* energy as:

$$Reg_{SRE}(\Phi) = \frac{1}{4}\int \mu^2_{(\bar{W},Cov)}(\log(\Sigma(x))) = \frac{1}{4}\int \mathrm{Vect}(W-\bar{W})\mathrm{Cov}^{(\text{-}1)}\mathrm{Vect}(W-\bar{W})^{\mathrm{T}}$$

As we are using a Mahalanobis distance, this least-squares criterion can be seen as the log-likelihood of a Gaussian process on strain tensor fields: we are implicitly modeling the a-priori probability of the deformation. In a registration framework, this point of view is particularly interesting as it opens the way to use Bayesian estimation methods for non-linear registration.

4 Isotropic Riemannian elasticity

With the general statistical Riemannian elasticity, we can take into account the anisotropic properties of the material, as they could be revealed by the statistics. However, in order to better explain the properties of this new tool, we focus in the following on isotropic covariances matrices. Seen as a quadratic form, the covariance is isotropic if $\mu^2(W) = \mu^2(R\,W\,R^{\mathrm{T}})$ for any rotation R. This means that it only depends on the eigenvalues of W, or equivalently on the matrix invariants $\mathrm{Tr}(W)$, $\mathrm{Tr}(W^2)$ and $\mathrm{Tr}(W^3)$. However, as the form is quadratic in W, we are left only with $\mathrm{Tr}(W)^2$ and $\mathrm{Tr}(W^2)$ that can be weighted arbitrarily, e.g. by μ and $\lambda/2$. Finally, the *isotropic Riemannian elasticity (IRE)* energy has the form:

$$Reg_{IRE}(\Phi) = \int \frac{\mu}{4}\mathrm{Tr}\left((\log(\Sigma) - \bar{W})^2\right) + \frac{\lambda}{8}\mathrm{Tr}(\log(\Sigma) - \bar{W})^2$$

For a null mean $\bar{W}$, we retrieve the classical form of isotropic elastic energy with Lamé coefficients, but with the logarithmic strain tensor. This form was expected as the St Venant-Kirchoff energy was also derived for isotropic materials.

4.1 Optimizing the Riemannian elasticity

To use the logarithmic elasticity energy as a regularization criterion in the registration framework we summarized in Section 2, we have to compute its gradient. Let as assume that $\bar{W} = 0$. Thanks to the properties of the differential of the log [19], we have $\mathrm{Tr}(\partial_V \log(\Sigma)) = \mathrm{Tr}(\Sigma^{(\text{-}1)}\,V)$ and $\langle\, \partial_V \log(\Sigma) \mid W\,\rangle = \langle\, \partial_W \log(\Sigma) \mid V\,\rangle$. Thus, using $V = \partial_u \Sigma = \sum_\alpha \left(\partial_\alpha u\, \partial_\alpha \Phi^{\mathrm{T}} + \partial_\alpha \Phi\, \partial_\alpha u^{\mathrm{T}}\right)$ and $W = \log(\Sigma)$, we can write the directional derivative of the criterion:

$$\begin{aligned}
\partial_u Reg_{IRE}(\Phi) &= \int \frac{\mu}{2}\langle\, W \mid \partial_V \log(\Sigma)\,\rangle + \frac{\lambda}{4}\mathrm{Tr}(W)\,\mathrm{Tr}(\partial_V \log(\Sigma)) \\
&= \int \frac{\mu}{2}\langle\, \partial_W \log(\Sigma) \mid V\,\rangle + \frac{\lambda}{4}\mathrm{Tr}(W)\,\mathrm{Tr}(\Sigma^{(\text{-}1)}\,V) \\
&= \sum_\alpha \int \mu\,\langle\, \partial_W \log(\Sigma)\,\partial_\alpha \Phi \mid \partial_\alpha u\rangle + \frac{\lambda}{2}\mathrm{Tr}(W)\,\langle\, \Sigma^{(\text{-}1)}\,\partial_\alpha \Phi \mid \partial_\alpha u\rangle
\end{aligned}$$

Integrating once again by part with homogeneous Neumann boundary conditions, we end up with the gradient:

$$\nabla Reg_{IRE}(\Phi) = -\sum_\alpha \partial_\alpha(Z\,\partial_\alpha\Phi) \quad \text{with} \quad Z = \mu\,\partial_W \log(\Sigma) + \tfrac{\lambda}{2}\mathrm{Tr}(W)\,\Sigma^{(-1)} \quad (2)$$

Notice the similarity with the gradient of the standard elasticity (eq. 1). For a non null mean deformation $\bar{W}(x)$, we just have to replace W by $W - \bar{W}$ in the above formula. One can even show that the same formula still holds for the general statistical Riemannian elasticity with $Z = \partial_X \log(\Sigma)$ where X is the symmetric matrix defined by $\mathrm{Vect}(X) = \mathrm{Cov}^{(-1)}\,\mathrm{Vect}(\log(\Sigma))$.

4.2 Practical implementation

Thus, we can optimize the logarithmic elasticity exactly like we did in Section 2.2 for the Euclidean elasticity. The only additional cost is the computation of the tensor Z, which implies the computation of the logarithm $W = \log(\Sigma)$ and its directional derivative $\partial_W \log(\Sigma)$. This cost would probably be prohibitive if we had to rely on numerical approximation methods. Fortunately, we were able to compute an explicit and very simple closed-form expression that only requires the diagonalization $\Sigma = R\,D\,R^{\mathrm{T}}$ [19]:

$$[R^{\mathrm{T}}\,\partial_V \log(\Sigma)\,R]_{ij} = [R^{\mathrm{T}}\,V\,R]_{ij}\,\lambda_{ij} \quad \text{with} \quad \lambda_{ij} = (\log(d_i) - \log(d_j))/(d_i - d_j)$$

Notice that formula is computationally well posed since $\lambda_{ij} = \frac{1}{d}(1 + \frac{1}{12}\varepsilon^2 d^2 + \frac{1}{80}\varepsilon^4 d^4 + O(\varepsilon^6))$ with $d = (d_i + d_j)/2$ and $\varepsilon = d_i - d_j$.

5 Experiments

To evaluate the potential of the Riemannian elasticity as a regularization criterion in non-rigid registration, we implemented the following basic gradient descent algorithm: at each iteration, the algorithm computes the gradient of the similarity criterion (we chose the local correlation coefficient to take image biases into account), and adds the derivative of the Euclidean or the Riemannian elastic energies according Sections 2.1 and 4.1. Then, a fraction η of this gradient is added to the current displacement field. This loop is embedded in a multiscale pyramid to capture larger deformations. At each pyramid step, iterations are stopped when the evolution of the transformation is too small (typically a hundred iterations by level). The whole algorithm is implemented in C++, and parallelized using Message Passing Interface (MPI) library on a cluster of PCs.

We tested the algorithm on clinical T1 weighted MR images of the brain of Parkinsonian patients (see Fig. 1). The ROI including the full head has 186 x 124 x 216 voxels of size 0.94 x 1.30 x 0.94 mm. Images were first affine registered and intensity corrected. We used a fraction $\eta = 5.10^{-4}$ for the gradient descent and standard values $\mu = \lambda = 0.2$ for both Euclidean and isotropic Riemannian elastic energies. The algorithm took about 1h for the Euclidean elasticity and 3h for the isotropic Riemannian regularization on a cluster of 12 AMD bi-Opteron PC

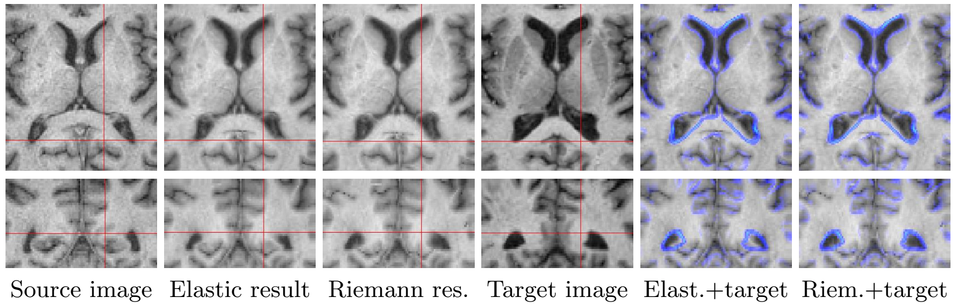

Fig. 1. Experimental comparison of registration with the Euclidean and the Riemannian elasticity regularization. From left to right, we displayed corresponding zoom of axial and coronal slices of: the source image, the elastically deformed source image, the Riemannian elastic result, the target image and the elastic and Riemannian results with the contours of the target image superimposed. Euclidean and Riemannian results are globally equivalent. One can only notice a slightly larger and better deformation of the right ventricle (near the crossing of the axes) with the Riemannian elasticity.

at 2 Ghz, connected by a gigabit Ethernet Network. These computations times show that our basic implementation of the Riemannian Elasticity is only 3 times slower than the Euclidean one. The diagonalization of the symmetric matrices being performed using a standard Jacobi method, we could easily design a much more efficient computation in dimensions 2 and 3. In terms of deformation, the results are quite similar for both methods *in the absence of any a priori statistical information*. However, we expect to show in the near future that taking into account statistical information about the expected deformability improves the results both in terms of accuracy and robustness.

6 Conclusion

We proposed in this paper an integrated framework to compute the statistics on deformations and re-introduce them as constraints in non-linear registration algorithms. This framework is based on the interpretation of the elastic energy as a Euclidean distance between the Cauchy-Green strain tensor and the identity (i.e. the local rigidity). By providing the space of tensors with a more suitable Riemannian metric, namely a Log-Euclidean one, we can define proper statistics on deformations, like the mean and the covariance matrix. Taking these measurements into account in a statistical (i.e. a Mahalanobis) distance, we end-up with the statistical Riemannian elasticity regularization criterion. This criterion can also be viewed as the log-likelihood of the deformation probability, which opens the way to Bayesian deformable image registration algorithms.

Riemannian elasticity gives a natural framework to measure statistics on inter-subject deformations. We demonstrate with the isotropic version that it is also an effective regularization criterion for non-linear registration algorithms. There is of course room for a lot of improvements that we plan to tackle in the near future. We are currently computing the deformation statistics (mean and covariance of the logarithmic strain tensor) on a database of brain images to assess their impact on the registration results. We also plan to evaluate carefully how the implementation influences the theoretical inverse-consistency property of the Riemannian elasticity, as this feature may turn out to be very useful for fine measurements of volume changes.

References

1. R. Bajcsy and S. Kovačič. Multiresolution elastic matching. *Computer Vision, Graphics and Image Processing*, 46:1–21, 1989.
2. G. E. Christensen, S. C. Joshi, and M. I. Miller. Volumetric transformation of brain anatomy. *IEEE Trans. Med. Imaging*, 16(6):864–877, 1997.
3. J.-P. Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis*, 2(3), 1998.
4. Jan Modersitzki. *Numerical Methods for Image Registration*. Numerical Mathematics and Scientific Computations. Oxford University Press, 2004.
5. H. Lester, S.R. Arridge, K.M. Jansons, L. Lemieux, J.V. Hajnal, and A. Oatridge. Non-linear registration with the variable viscosity fluid algorithm. In *Proc of IPMI;99*, pages 238–251, 1999.
6. R. Stefanescu, X. Pennec, and N. Ayache. Grid powered nonlinear image registration with locally adaptive regularization. *Med. Image Anal.*, 8(3):325–342, 2004.
7. 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.
8. D. Rueckert, A.F. Frangi, and J.A. Schnabel. Automatic construction of 3D statistical deformation models of the brain using non-rigid registration. *IEEE TMI*, 22:1014–1025, 2003.
9. P. Fillard, V. Arsigny, X. Pennec, P. Thompson, and N. Ayache. Extrapolation of sparse tensor fields: Application to the modeling of brain variability. In *Proc. of IPMI'05*, LNCS, Glenwood springs, CO, USA, July 2005. Springer.
10. P.G. Ciarlet. *Mathematical elasticity Vol. 1: Three-dimensionnal elasticity*. Elsevier Science B.V., 1988.
11. GE Christensen and HJ. Johnson. Consistent image registration. *IEEE Trans Med Imaging*, 20(7):568–82, July 2001.
12. 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.
13. 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.
14. Ch. 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.
15. X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *IJCV*, 65(1), October 2005. Also as INRIA Research Report 5255.
16. V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Fast and simple calculus on tensors in the Log-Euclidean framework. In *Proc. of MICCAI'05*, LNCS, Palm Springs, California, October 2005. Springer Verlag. To Appear.
17. A.D. Freed. Natural strain. *J. of Eng. Materials & Technology*, 117:379–385, 1995.
18. R.P. Woods. Characterizing volume and surface deformations in an atlas framework: theory, applications, and implementation. *NeuroImage*, 18(3):769–788, 2003.
19. V. Arsigny, X. Pennec, P. Fillard, and N. Ayache. Dispositif perfectionné de traitement ou de production d'images de tenseurs. Patent filing 0503483, April 2005.