

Log-Euclidean Metrics for Fast and Simple Calculus on Diffusion Tensors¹.

Vincent Arsigny ², Pierre Fillard ³, Xavier Pennec ⁴ and Nicholas Ayache ⁵.

June 7, 2006

¹This is a preprint version of an article to appear in Magnetic Resonance in Medicine, published by Wiley-Liss Inc. All rights reserved, including copyright. Running heading: **Log-Euclidean metrics on diffusion tensors.** Total word count: 6400.

²Corresponding Author. INRIA Sophia - Epidaure Research Project, BP 93, 06902 Sophia Antipolis Cedex, France. Tel: (+) 33 4 92 38 71 59, Fax: (+) 33 4 92 38 76 69 , e-mail: Vincent.Arsigny@Polytechnique.org.

³INRIA Sophia-Antipolis, Epidaure Project. E-mail: Pierre.Fillard@Sophia.inria.fr

⁴INRIA Sophia-Antipolis, Epidaure Project. E-mail: Xavier.Pennec@Sophia.inria.fr

⁵INRIA Sophia-Antipolis, Epidaure Project. E-mail: Nicholas.Ayache@Sophia.inria.fr

Abstract

Diffusion tensor imaging (DT-MRI or DTI) is an emerging imaging modality whose importance has been growing considerably. However, the processing of this type of data (i.e. symmetric positive-definite matrices), called ‘tensors’ here, has proved difficult in recent years. Usual Euclidean operations on matrices suffer from many defects on tensors, which have led to the use of many *ad hoc* methods. Recently, affine-invariant Riemannian metrics have been proposed as a rigorous and general framework in which these defects are corrected. These metrics have excellent theoretical properties and provide powerful processing tools, but also lead in practice to complex and slow algorithms. To remedy this limitation, a new family of Riemannian metrics called Log-Euclidean is proposed in this article. They also have excellent theoretical properties and yield similar results in practice, but with much simpler and faster computations. This new approach is based on a novel vector space structure for tensors. In this framework, Riemannian computations can be converted into Euclidean ones once tensors have been transformed into their matrix logarithms. Theoretical aspects are presented and the Euclidean, affine-invariant and Log-Euclidean frameworks are compared experimentally. The comparison is carried out on interpolation and regularization tasks on synthetic and clinical 3D DTI data.

Key words: DT-MRI, Riemannian metrics, vector space, interpolation, regularization.

INTRODUCTION

Diffusion tensor imaging (DT-MRI or DTI or equivalently DT imaging) (1) is an emerging imaging modality whose importance has been growing considerably. In particular, most attempts to reconstruct non-invasively the connectivity of the brain are based on DTI (see (2–7) and references within for classical fiber tracking algorithms). Other applications of DT-MRI also include the study of diseases such as stroke, multiple sclerosis, dyslexia and schizophrenia (8).

The diffusion tensor is a simple and powerful model used to analyze the content of Diffusion-Weighted images (DW-MRIs). It is based on the assumption that the motion of water molecules can be well approximated by a Brownian motion in each voxel of the image. This Brownian motion is entirely characterized by a symmetric and positive-definite matrix, called the ‘diffusion tensor’ (1). In this article, we restrict the term ‘tensor’ to mean a symmetric and positive-definite matrix.

With the increasing use of DT-MRI, there has been a growing need to generalize to the tensor case many usual vector processing tools. In particular, regularization techniques are required to denoise them. Furthermore, classical tasks like interpolation also need to be generalized to resample DT images, for example to work with isotropic voxels, as recommended in (6). It would also be very valuable to generalize to tensors classical vector statistical tools, in order to analyze the variability of tensors or model the noise that corrupts them. Previous attempts to do so are only partially satisfactory: for example, it was proposed in (9) to define a Gaussian distribution on tensors as a Gaussian distribution on symmetric matrices, without taking into account the positive-definiteness constraint. This becomes problematic with Gaussians whose covariance is large: in this case, non-positive eigenvalues do appear with a significant probability.

Many *ad hoc* approaches have already been proposed in the literature to process tensors (see (10, 11) and references within). But in order to fully generalize to tensors the usual PDEs or statistical tools used on scalars or vectors, one needs to define a consistent operational framework. The framework of Riemannian metrics (12, 13) has recently emerged as particularly adapted to this task (14–17).

The Defects of Euclidean Calculus

The simplest Riemannian structures are the Euclidean ones. Let \mathbf{S}_1 and \mathbf{S}_2 be two tensors. An example of Euclidean structure is given by the so-called ‘Frobenius distance’: $\text{dist}^2(\mathbf{S}_1, \mathbf{S}_2) = (\text{Trace}((\mathbf{S}_1 - \mathbf{S}_2)^2))$. This straightforward metric leads *a priori* to simple computations. Unfortunately, though Euclidean distances are well-adapted to general square matrices, they are unsatisfactory for tensors, which are very specific matrices. Typically, symmetric matrices with null or negative eigenvalues appear on clinical data as soon as we perform on tensors Euclidean operations which are non-convex. Example of such situations are the estimation of tensors from diffusion-weighted images, the regularization of tensors fields, etc. The noise in the data is at the source of this problem. To avoid obtaining non-positive eigenvalues, which are difficult to interpret physically, it has been proposed to regularize only *features* extracted from tensors, like first eigenvectors (18) or orientations (11). This is only partly satisfactory, since such approaches do not take into account *all* the information carried by tensors.

After a diffusion time Δ , we know with a confidence say of 95% that a water molecule is located within a region called a *confidence region*, which is the multidimensional equivalent of a confidence in-

terval. The larger the volume of these regions, the larger is the dispersion of the random displacement of water molecules. In the case of Brownian motion, the random displacement is Gaussian, and confidence regions are therefore ellipsoids. The volumes of these ellipsoids are proportional to the square root of the *determinant* of the covariance matrix of the displacement. In DT-MRI, this covariance matrix is equal to the diffusion tensor multiplied by 2Δ (1). The value of the determinant of the diffusion tensor is therefore a direct measure of the dispersion of the local diffusion process. But the Euclidean averaging of tensors generally leads to a *tensor swelling effect* (11,19,20): the determinant (and thus the dispersion) of the Euclidean mean of tensors can be larger than the determinants of the original tensors! Introducing more dispersion in computations amounts to introducing more diffusion, which is physically unrealistic.

Riemannian Metrics

To fully circumvent these difficulties, affine-invariant Riemannian metrics have been recently proposed for tensors by several teams. The application of these metrics to the averaging of tensors and the definition of a Riemannian anisotropy measure were presented (15, 21). The generalization of principal component analysis (PCA) to tensors was given in (17). The affine-invariant statistical framework and its application to the segmentation of DT-MRI was presented in (16). PDEs within the affine-invariant framework were studied in (14) with applications to the interpolation, extrapolation and regularization of tensor fields.

With affine-invariant metrics, symmetric matrices with negative and null eigenvalues are at an infinite distance from any tensor and the swelling effect disappears. Practically, this prevents the appearance of non-positive eigenvalues, which is particularly difficult to avoid in Euclidean algorithms. But the price paid for this success is a high computational burden, essentially due to the curvature induced on the tensor space. This substantial computational cost can be seen directly from the formula giving the distance between two tensors \mathbf{S}_1 and \mathbf{S}_2 (14):

$$\text{dist}(\mathbf{S}_1, \mathbf{S}_2) = \left\| \log \left(\mathbf{S}_1^{-\frac{1}{2}} \cdot \mathbf{S}_2 \cdot \mathbf{S}_1^{-\frac{1}{2}} \right) \right\|, \quad [1]$$

where $\|\cdot\|$ is a Euclidean norm on symmetric matrices. In general, affine-invariant computations involve an intensive use of matrix inverses, square roots, logarithms and exponentials.

We present in this article a new Riemannian framework to fully overcome these computational limitations while preserving excellent theoretical properties. Moreover, we obtain this result without any unnecessary complexity, since all computations on tensors are converted into computations on vectors. This framework is based on a new family of metrics named *Log-Euclidean*, which are particularly simple to use. They result in classical Euclidean computations in the domain of matrix logarithms. In the next section, we present the theory of Log-Euclidean metrics (more details on this theory can be found in a research report, see (22)). In the Methods section, we describe the adaptation of classical processing tools to the Log-Euclidean framework for interpolation and regularization tasks. We also present a highly useful tool for the visualization of difference between tensors: the *absolute value* of a symmetric matrix. Then, we show that the affine-invariant and Log-Euclidean frameworks perform better than the Euclidean one for the interpolation and regularization of our synthetic and clinical 3D DT-MRI data. Affine-invariant and Log-Euclidean results are very similar, but computations are simpler and experimentally much faster in the Log-Euclidean than in the affine-invariant framework.

THEORY

Matrix Exponential, Logarithm and Powers

The notions of matrix logarithm and exponential are central in the theoretical framework presented here. For any matrix \mathbf{M} , its exponential is given by: $\exp(\mathbf{M}) = \sum_{k=0}^{\infty} \mathbf{M}^k / k!$. As in the scalar case, the *matrix logarithm* is defined as the inverse of the exponential. One should note that for general matrices, neither the uniqueness nor the existence of a logarithm is guaranteed for a given invertible matrix (23,24). However, the important point here is that the logarithm of a tensor is well-defined and is a symmetric matrix. Conversely, the exponential of any symmetric matrix yields a tensor. This means that under the matrix exponentiation operation, there is a one-to-one correspondence between symmetric matrices and tensors.

This one-to-one correspondence can be seen quite intuitively thanks to the simple spectral decomposition of these matrices. Indeed, the matrix logarithm \mathbf{L} of a tensor \mathbf{S} can be calculated in three steps:

1. perform a diagonalization of \mathbf{S} , which provides a rotation matrix \mathbf{R} and a diagonal matrix \mathbf{D} with the eigenvalues of \mathbf{S} in its diagonal, with the equality: $\mathbf{S} = \mathbf{R}^T \cdot \mathbf{D} \cdot \mathbf{R}$.
2. transform each diagonal element of \mathbf{D} (which is necessarily positive, since it is an eigenvalue of \mathbf{S}) into its natural logarithm in order to obtain a new diagonal matrix $\tilde{\mathbf{D}}$.
3. recompose $\tilde{\mathbf{D}}$ and \mathbf{R} to obtain the logarithm with the formula $\mathbf{L} = \log(\mathbf{S}) = \mathbf{R}^T \cdot \tilde{\mathbf{D}} \cdot \mathbf{R}$.

Conversely, the matrix exponential \mathbf{S} is obtained by replacing the natural logarithm with the scalar exponential. One also generalizes of the notion of powers (and in particular square roots) to tensors by replacing their eigenvalues by the corresponding scalar power (for example by their square roots).

Definition of Log-Euclidean Metrics

Based on the specific properties of the matrix exponential and logarithm on tensors that we presented above, we can now define a novel *vector space* structure on tensors. This is quite a surprising result: in the sense of this new algebraic structure, *tensors* can be also looked upon as *vectors*! As will be shown in the rest of this article, this novel viewpoint provides a particularly powerful and simple-to-use framework to process tensors.

Since there is a one-to-one mapping between the tensor space and the *vector* space of symmetric matrices, one can transfer to tensors the standard algebraic operations (addition “+” and scalar multiplication “.”) with the matrix exponential. This defines on tensors a *logarithmic multiplication* \odot and a *logarithmic scalar multiplication* \otimes , given by:

$$\begin{cases} \mathbf{S}_1 \odot \mathbf{S}_2 \stackrel{def}{=} \exp(\log(\mathbf{S}_1) + \log(\mathbf{S}_2)) \\ \lambda \otimes \mathbf{S} \stackrel{def}{=} \exp(\lambda \cdot \log(\mathbf{S})) = \mathbf{S}^\lambda. \end{cases}$$

The logarithmic multiplication is *commutative* and coincides with matrix multiplication whenever the two tensors \mathbf{S}_1 and \mathbf{S}_2 commute in the matrix sense. With \odot and \otimes , the tensor space has by construction a *vector*

space structure, which is *not* the usual structure directly derived from addition and scalar multiplication on matrices.

When one considers only the multiplication \odot on the tensor space, one has a *Lie group* structure (13), i.e. a space which is both a smooth manifold and a group in which multiplication and inversion are *smooth mappings*. This type of mathematical tool is for example particularly useful in theoretical physics (25). Here, the smoothness of \odot comes from the fact that both the exponential and the logarithm mappings are smooth (22). Among Riemannian metrics in Lie groups, the most convenient in practice are *bi-invariant metrics*, i.e. metrics that are invariant by multiplication and inversion. When they exist, these metrics are used in differential geometry to generalize to Lie groups a notion of mean which is completely consistent with multiplication and inversion. This approach applies particularly well in the case of the group of rotations (26–28). However, such metrics do not always exist, as in the case of the groups of Euclidean motions (29, 30) and affine transformations. It is remarkable that bi-invariant metrics exist in our tensor Lie group. Moreover, they are particularly simple. Their existence simply results from the *commutativity* of logarithmic multiplication between tensors. We have named such metrics *Log-Euclidean metrics*, since they correspond to Euclidean metrics in the domain of logarithms. From a Euclidean norm $\|\cdot\|$ on symmetric matrices, they can be written:

$$\text{dist}(\mathbf{S}_1, \mathbf{S}_2) = \|\log(\mathbf{S}_1) - \log(\mathbf{S}_2)\|. \quad [2]$$

From Eq. [2], it is clear that Log-Euclidean metrics are also Euclidean distances for the vector space structure we defined earlier. We did not define them directly from the latter algebraic structure to emphasize the fact that they are also *Riemannian metrics*, like affine-invariant metrics.

As one can see, the Log-Euclidean distance is much simpler than the equivalent affine-invariant distance given by Eq. [1], where matrix multiplications, square roots and inverses are used before taking the norm of the logarithm. The greater simplicity of Log-Euclidean metrics can also be seen from Log-Euclidean geodesics in the tensor space. In the Log-Euclidean case, the shortest path $\gamma_{\text{LE}}(t)$ going from the tensor \mathbf{S}_1 at time 0 to the tensor \mathbf{S}_2 at time 1 is a straight line in the domain of logarithms. This geodesic is given by:

$$\gamma_{\text{LE}}(t) = \exp((1-t)\log(\mathbf{S}_1) + t\log(\mathbf{S}_2)).$$

Its affine-invariant equivalent $\gamma_{\text{Aff}}(t)$ involves the use of square roots and inverses and takes the following form:

$$\gamma_{\text{Aff}}(t) = \mathbf{S}_1^{\frac{1}{2}} \cdot \exp\left(t \log\left(\mathbf{S}_1^{-\frac{1}{2}} \cdot \mathbf{S}_2 \cdot \mathbf{S}_1^{-\frac{1}{2}}\right)\right) \cdot \mathbf{S}_1^{\frac{1}{2}}.$$

Contrary to the classical Euclidean framework on tensors, one can see from Eq. [2] that symmetric matrices with null or negative eigenvalues are at an infinite distance from any tensor and therefore will not appear in practical computations. The same property holds for affine-invariant metrics (14).

Invariance Properties of Log-Euclidean Metrics

Log-Euclidean metrics satisfy a number of *invariance properties*, i.e. are left unchanged by several operations on tensors. First, distances are not changed by inversion, since taking the inverse of a system of

matrices only results in the multiplication by -1 of their logarithms, which does not change the value of the distance given by Eq. [2]. Also, Log-Euclidean metrics are by construction invariant with respect to any logarithmic multiplication, i.e. are invariant by any translation in the domain of logarithms. However, there is more. Although Log-Euclidean metrics do not yield full affine-invariance as the affine-invariant metrics defined in (14), a number of them are invariant by *similarity* (orthogonal transformation and scaling) (22). This means that computations on tensors using these metrics will be invariant with respect to a change of coordinates obtained by a similarity. In this work, we use the simplest similarity-invariant Log-Euclidean metric, which is given by:

$$\text{dist}(\mathbf{S}_1, \mathbf{S}_2) = \left(\text{Trace} \left(\{ \log(\mathbf{S}_1) - \log(\mathbf{S}_2) \}^2 \right) \right)^{\frac{1}{2}} .$$

Log-Euclidean Computations on Tensors

From a practical point of view, one would like operations such as averaging, filtering, etc. to be as simple as possible. In the affine-invariant case, such operations rely on an intensive use of matrix exponentials, logarithms, inverses and square roots. In our case, the space of tensors with a Log-Euclidean metric is in fact *isomorphic* (the algebraic structure of vector space is conserved) and *isometric* (distances are conserved) with the corresponding Euclidean space of symmetric matrices. As a consequence, the Riemannian framework for statistics and analysis is extremely simplified. To illustrate this, let us recall the notion of *Fréchet mean* (12, 31), which is the Riemannian equivalent of the Euclidean (or arithmetic) mean. Given a Riemannian metric, the associated Fréchet mean of N tensors $\mathbf{S}_1, \dots, \mathbf{S}_N$ with arbitrary positive weights w_1, \dots, w_N is defined as the point $\mathbb{E}(\mathbf{S}_1, \dots, \mathbf{S}_N)$ minimizing the following *metric dispersion*:

$$\mathbb{E}(\mathbf{S}_1, \dots, \mathbf{S}_N) = \arg \min_{\mathbf{S}} \sum_{i=1}^N w_i \text{dist}^2(\mathbf{S}, \mathbf{S}_i),$$

where $\text{dist}(\cdot, \cdot)$ is the distance associated to the metric. The Log-Euclidean Fréchet mean is a direct generalization of the geometric mean of positive numbers and is given *explicitly* by:

$$\mathbb{E}_{\text{LE}}(\mathbf{S}_1, \dots, \mathbf{S}_N) = \exp \left(\sum_{i=1}^N w_i \log(\mathbf{S}_i) \right) . \quad [3]$$

The closed form given by Eq. [3] makes the computation of Log-Euclidean means straightforward. On the contrary, there is no closed form for affine-invariant means $\mathbb{E}_{\text{Aff}}(\mathbf{S}_1, \dots, \mathbf{S}_N)$ as soon as $N > 2$ (21). The affine-invariant is only *implicitly* defined through the following barycentric equation:

$$\sum_{i=1}^N w_i \log \left(\mathbb{E}_{\text{Aff}}(\mathbf{S}_1, \dots, \mathbf{S}_N)^{-1/2} \cdot \mathbf{S}_i \cdot \mathbb{E}_{\text{Aff}}(\mathbf{S}_1, \dots, \mathbf{S}_N)^{-1/2} \right) = 0. \quad [4]$$

In the literature, this equation is solved iteratively, for instance using a Gauss-Newton method as detailed in (14, 16, 17). This optimization method has the advantage of having quite a fast convergence speed, like all Newton methods.

Contrary to the affine-invariant case, the processing of tensors in the Log-Euclidean framework is simply Euclidean in the logarithmic domain. Tensors can be transformed first into symmetric matrices (i.e. vectors)

using the matrix logarithm. Then, to simplify even more computations, these matrices with 6 degrees of freedom can be represented by 6D vectors in the following way:

$$\log(\mathbf{S}) \simeq \vec{\mathbf{S}} = \left(\log(\mathbf{S})_{1,1}, \log(\mathbf{S})_{2,2}, \log(\mathbf{S})_{3,3}, \sqrt{2} \cdot \log(\mathbf{S})_{1,2}, \sqrt{2} \cdot \log(\mathbf{S})_{1,3}, \sqrt{2} \cdot \log(\mathbf{S})_{2,3} \right)^T,$$

where $\log(\mathbf{S})_{i,j}$ is the coefficient of $\log(\mathbf{S})$ placed in the (i, j) position. With this representation, the classical Euclidean norm between such 6D vectors is equal to a Log-Euclidean similarity-invariant distance between the tensors they represent. Note that this is true only for the particular similarity-invariant distance used in this work. To deal with another Log-Euclidean distance, one should adapt the 6D vector representation to the metric by changing adequately the relative weights of the matrix coefficients.

Once tensors have been transformed into symmetric matrices or 6D vectors, classical vector processing tools can be used *directly* on these 6D representations. Finally, results obtained on logarithms are mapped back to the tensor domain with the exponential. Hence, vector statistical tools or PDEs are readily generalized to tensors in this framework.

Comparison of the Affine-Invariant and Log-Euclidean Frameworks

As will be shown experimentally in the Results section, Log-Euclidean computations provide results very similar to their affine-invariant equivalent, presented in (14). The reason behind this is the following: the two families of metrics provide two different generalizations to tensors of the geometric mean of positive numbers. By this we mean that the determinants of both Log-Euclidean and affine-invariant means of tensors are exactly equal to the *scalar* geometric mean of the determinants of the data (22). This explains the absence of swelling effect in both cases, since the interpolation of tensors along geodesics yields in both cases the same monotonic interpolation of determinants.

The two Riemannian means are even identical in a number of cases, in particular when averaged tensors commute in the sense of matrix multiplication. Yet, the two means are different in general, as shown theoretically in (22) (the trace of the Log-Euclidean mean is always larger (or equal) than the trace of the affine-invariant mean) and experimentally in the Results section. More precisely, Log-Euclidean means are generally more anisotropic than their affine-invariant equivalent. We observed that this resemblance between the two means extends to general computations which involve averaging, such as regularization procedures, as is shown in the Results section.

METHODS

Interpolation

Voxels in clinical DT images are often quite anisotropic. Algorithms tracking white matter tracts can be biased by this anisotropy, and it is therefore recommended (e.g. see (6)) to use isotropic voxels. A preliminary resampling step with an adequate interpolation method is therefore important for such algorithms. Adequate interpolation methods are also required to generalize to the tensor case usual registration techniques used on scalar or vector images. The framework of Riemannian metrics allows a direct generalization of classical

resampling methods, by re-interpreting them as computing weighted means of the original data. Then the idea is to replace the Euclidean mean by its Riemannian counterpart, i.e. the Fréchet mean. See (14) for a more detailed discussion of this topic. This way one can generalize the classical linear, bilinear and trilinear interpolations to tensors with a Riemannian metric. For both metrics mentioned in this work, this entails in one case using directly Eq. [3] and in the other case iteratively solving Eq. [4].

Regularization

DT images are corrupted by noise, and regularizing them can be a crucial preliminary step for DTI-based algorithms that reconstruct the white matter connectivity. As shown in (14), Riemannian metrics provide a general framework to regularize to tensors usual vector regularization tools.

Practically, an anisotropic regularization is very valuable, since it allows a substantial reduction of the noise level while sharp contours and structures are mostly preserved. We focus here on a simple and typical Riemannian criterion for the anisotropic regularization of tensor fields, which is based on Φ -functions (11, 32). In this context, the regularization is obtained by the minimization of a Φ -functional $Reg(\mathbf{S})$ given by:

$$Reg(\mathbf{S}) = \int_{\Omega} \Phi(\|\nabla\mathbf{S}\|_{\mathbf{S}(\mathbf{x})}(\mathbf{x})) dx,$$

where Ω is the spatial domain of the image and $\Phi(s)$ a function penalizing large values of the norm of the *spatial gradient* $\nabla\mathbf{S}$ of the tensor field $\mathbf{S}(\mathbf{x})$. The spatial gradient is defined here as $\nabla\mathbf{S} = (\frac{\partial\mathbf{S}}{\partial x_1}, \frac{\partial\mathbf{S}}{\partial x_2}, \frac{\partial\mathbf{S}}{\partial x_3})$, where x_1, x_2 and x_3 are the three spatial coordinates, and where $\frac{\partial\mathbf{S}}{\partial x_i}$ is the matrix describing how $\mathbf{S}(\mathbf{x})$ linearly varies near \mathbf{x} in the i^{th} spatial direction. Note that $\frac{\partial\mathbf{S}}{\partial x_i}$ is only symmetric and not necessarily positive definite because it is given by an infinitesimal *difference* between two tensors, which is a non-convex operation. For more details on how spatial gradients can be practically computed, see (14) Section 5.

Here, we use the classical function $\Phi(s) = 2\sqrt{1 + s^2/\kappa^2} - 2$ (11). We would like to emphasize that contrary to the Euclidean case, the norm of $\nabla\mathbf{S}$ *depends explicitly on the current point* $\mathbf{S}(\mathbf{x})$ (see (14, 22) for more details) and is given by:

$$\|\nabla\mathbf{S}\|_{\mathbf{S}(\mathbf{x})}^2 = \sum_{i=1}^3 \left\| \frac{\partial\mathbf{S}}{\partial x_i}(\mathbf{x}) \right\|_{\mathbf{S}(\mathbf{x})}^2.$$

In general, this dependence on the current point leads to complex resolution methods. Thus, in the affine-invariant case, these methods rely on an intensive use of matrix inverses, square roots, exponentials and logarithms (14). However, in the Log-Euclidean framework the general Riemannian formulation is extremely simplified. The reason is that the dependence on the current tensor disappears on the logarithms of tensors (22), so that the norm of the gradient is given by:

$$\|\nabla\mathbf{S}\|_{\mathbf{S}(\mathbf{x})} = (\langle \nabla\mathbf{S}, \nabla\mathbf{S} \rangle_{\mathbf{S}(\mathbf{x})})^{\frac{1}{2}} = \|\nabla \log(\mathbf{S})\|_{\mathbf{Id}},$$

where \mathbf{Id} is the identity matrix. This means that only the scalar product at the identity needs to be used.

The transformation of tensors into their matrix logarithms transforms Riemannian computations at $\mathbf{S}(\mathbf{x})$ into Euclidean computations at \mathbf{Id} . As a consequence, *the energy functional can be minimized directly on the vector field of logarithms*. The regularized tensor field is given in a final step by the matrix exponential of regularized logarithms.

In the regularization experiments of this article, the minimization method used is a first-order gradient descent with a fixed time step dt . We use an explicit finite difference scheme on logarithms in the Log-Euclidean case (see (33) for details about numerical schemes and others aspects of the implementation) and the geodesic marching scheme described in (14) in the affine-invariant case. In the Euclidean framework, we also use affine-invariant geodesic marching rather than a classical explicit scheme to limit the appearance of non-positive eigenvalues, proceeding similarly as in (11). Homogeneous Neumann boundary conditions are used, parameters were empirically chosen to be $\kappa = 0.05$, $dt = 0.1$, and 100 iterations are performed in the results shown in Fig. 5 and 50 iterations for those shown in Fig. 6.

Absolute Value of a Symmetric Matrix

When several variants of an algorithm are used to process tensors images, visualization tools are quite valuable to inspect the results. A simple solution is to visualize an image of the norm of the (Euclidean) difference between tensors. Regrettably, all information about *orientation* is lost in this case.

To visualize simultaneously the magnitude *and* the orientation of differences, one can use the *absolute value* of a symmetric matrix. Similarly to the exponential or square root, it is defined as the symmetric positive semi-definite matrix obtained by replacing the eigenvalues of the original matrix by their absolute values. Thus, this absolute value retains all the information about the magnitude and the orientation of any symmetric matrix, and can still be visualized directly with the usual ellipsoid representation. As a consequence, this mathematical tool is very useful to visualize the difference between two tensors, as can be seen in the Results section. We first introduced this tool in (34).

Materials

The experiments in this study are carried out partly on synthetic tensor images, and partly on a clinical DTI volume. The clinical scan of the brain was acquired with a 1.5-T MR imaging system (Siemens Sonata) with actively shielded magnetic field gradients (G maximum, 40 mT/m). A sagittal spin-echo single shot echo-planar parallel Grappa diffusion-weighted imaging sequence with acceleration factor two and six non collinear gradient directions was applied with two b values (b=0 and $1000s.mm^{-2}$). Field of view: 24.0×24.0 cm; image matrix: 128×128 voxels; 30 sections with a thickness of 4mm; nominal voxel size: $1.875 \times 1.875 \times 4mm^3$. TR/TE= 4600/73 ms. The gradient directions used were as follows: $[(1/\sqrt{2}, 0, 1/\sqrt{2}); (-1/\sqrt{2}, 0, 1/\sqrt{2}); (0, 1/\sqrt{2}, 1/\sqrt{2}); (0, 1/\sqrt{2}, -1/\sqrt{2}); (1/\sqrt{2}, 1/\sqrt{2}, 0); (-1/\sqrt{2}, 1/\sqrt{2}, 0)]$ providing the best accuracy in tensor components when six directions are used (35). The acquisition time of diffusion-weighted imaging was 5 minutes and 35 seconds. Image analysis was performed on a voxel-by-voxel basis by using dedicated software (DPTools, <http://fmritools.hd.free.fr>). Before performing the tensor estimation, an unwarping algorithm was applied to the DTI data set to reduce distortions related to

eddy currents induced by the large diffusion-sensitizing gradients. This algorithm relies on a three-parameter distortion model including scale, shear, and linear translation in the phase-encoding direction (36). The optimal parameters were assessed independently for each section relative to the T2-weighted corresponding image by the maximization of the mutual information. However, due to the low signal-to-noise ratio in these images, part of the distortions remained. The tensors were estimated using the method described in (33), with a small regularization. The parameters of this estimation were set to $\lambda = 0.25$ and $\kappa = 0.1$. 50 iterations were used.



Figure 1: **Geodesic interpolation of two tensors.** **Left:** interpolated tensors. **Right:** graphs of the determinants of the interpolated tensors. **Top:** linear interpolation on coefficients. **Middle:** affine-invariant interpolation. **Bottom:** Log-Euclidean interpolation. The coloring of ellipsoids is based on the direction of dominant eigenvectors, and was only added to enhance the contrast of tensor images. Note the characteristic swelling effect observed in the Euclidean case due to a parabolic interpolation of determinants. This effect is not present in both Riemannian frameworks since determinants are monotonically interpolated. Note also that Log-Euclidean means are more anisotropic their affine-invariant counterparts.

RESULTS

Interpolation

Results of the (geodesic) linear interpolation of two synthetic tensors are presented in Fig. 1. One can clearly see the swelling effect characteristic of the Euclidean interpolation, which has no physical interpretation. On the contrary, a monotonic (and identical) interpolation of determinants is obtained in both Riemannian frameworks. The larger anisotropy in Log-Euclidean means is also clearly visible in this figure.

Fig. 2 shows the results obtained for the bilinear interpolation of four synthetic tensors with three methods: Euclidean (linear interpolation of coefficients), affine-invariant and Log-Euclidean. Again, there is a pronounced swelling effect in the Euclidean case, which does not appear in both Riemannian cases.

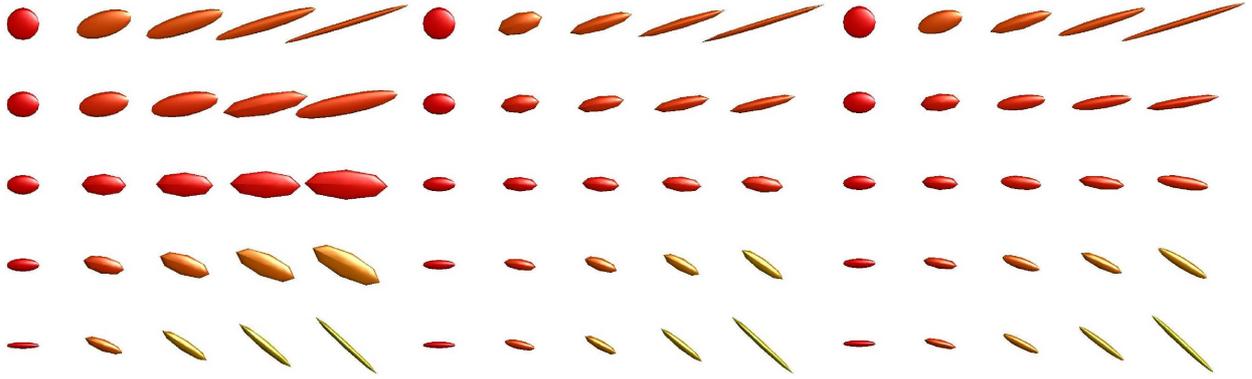


Figure 2: **Bilinear interpolation of 4 tensors at the corners of a grid.** **Left:** Euclidean reconstruction. **Middle:** affine-invariant reconstruction. **Right:** Log-Euclidean interpolation. Note the characteristic swelling effect observed in the Euclidean case, which is not present in both Riemannian frameworks. Note also that Log-Euclidean means are slightly more anisotropic than their affine-invariant counterparts.

Also, there is a slightly larger anisotropy in Log-Euclidean means. One should note that the computation of the affine-invariant mean here is iterative, since the number of averaged tensor is greater than 2 (we use the Gauss-Newton method described in (14)), whereas the closed form given by Eq. [3] is used directly in the Log-Euclidean case. This has a large impact on computation times: 0.003s (Euclidean), 0.009s (Log-Euclidean) and 1s (affine-invariant) for a 5×5 grid on a Pentium M 2 GHz. Computations were carried out in the MatlabTM framework, which explains the poor computational performance. A C++ implementation would yield much lower computation times, but the ratio would be comparable. This clearly demonstrate that Log-Euclidean metrics combine greater simplicity and performance, as compared to affine-invariant metrics, at least in terms of interpolation tasks.

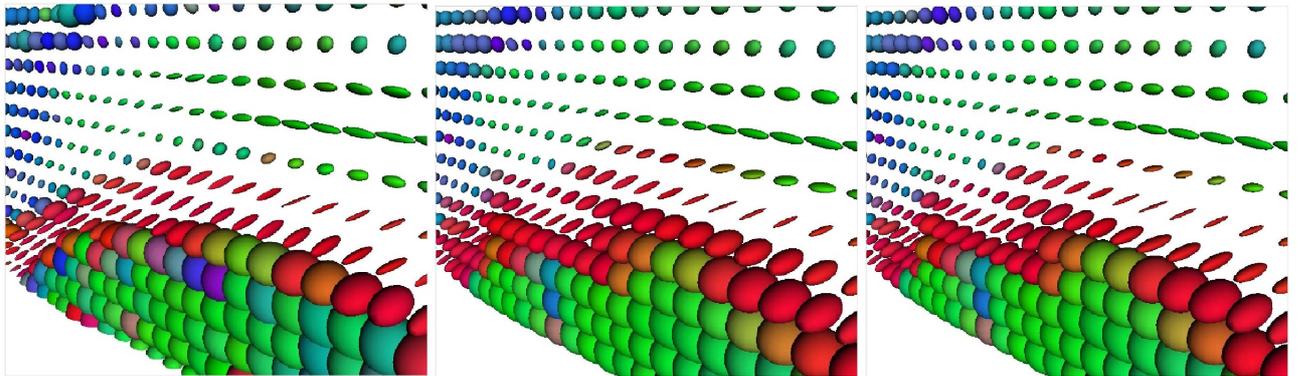


Figure 3: **Bilinear interpolation in a real DTI slice.** **Left:** Original DTI slice, before down-sampling. **Middle:** Euclidean interpolation. **Right:** Log-Euclidean interpolation. Half the columns and lines of the original DTI slice were removed before reconstruction with a bilinear interpolation. The slice is taken in the mid-sagittal plane and displayed in perspective. Again, the coloring of ellipsoids is based on the direction of dominant eigenvectors, and was only added to enhance the contrast of tensor images. Note how the tensors corresponding to the Corpus Callosum (in red, above the large and round tensors corresponding to a part of the ventricles) are better reconstructed (more anisotropic) in the Log-Euclidean case.

From a numerical point of view, one should note that the computation of Log-Euclidean means is more *stable* than in the affine-invariant case. On synthetic examples, we noticed that for large anisotropies (for instance with the dominant eigenvalue larger than 500 times the smallest), large numerical instabilities appear, essentially due to limited numerical accuracy of the logarithm computations (even with double precision). This can complicate greatly the computation of affine-invariant means. In the case of our clinical DTI data, this type of phenomenon also occurs, although to a lesser degree. We observed that the computation of the affine-invariant mean can in this case be 5 to 10 times longer than usual at times, when the averaged data presents a substantial inhomogeneity. On the contrary, the computation of Log-Euclidean means is much more stable since the logarithm and exponential are taken only once and thus even very large anisotropies can be dealt with. Of course, on clinical DT images anisotropies are not so pronounced and drastic instabilities will not appear. But for the processing of other types of tensors with much higher anisotropies, this could be crucial.

| Similarity Measure | Euclidean interpol. | Affine-invariant interpol. | Log-Euclidean interpol. |
|-----------------------------|---------------------|----------------------------|-------------------------|
| Mean Euclidean Error | 0.2659 | 0.2614 | 0.2611 |
| Mean Affine-invariant Error | 0.2703 | 0.2586 | 0.2584 |
| Log-Euclidean Error | 0.2694 | 0.2577 | 0.2575 |

Table 1: **Mean reconstruction errors for the clinical slice reconstruction experiment.** The three interpolation results are quite close. However, both Riemannian frameworks perform slightly better than the Euclidean one, independently of the similarity measure considered. This is essentially due to the better Riemannian reconstruction of the Corpus Callosum.

To compare the Euclidean and Riemannian bilinear interpolations on clinical data, we have reconstructed by bilinear interpolation a down-sampled DTI slice. One column out of two and one line out of two were removed. The slice was chosen in the mid-sagittal plane where strong variations are present in the DT image. The results in Fig. 3 show that the tensors corresponding to the corpus callosum are better reconstructed in the Log-Euclidean case. Affine-invariant results are very similar to Log-Euclidean ones and not shown here. In other regions, the differences between the interpolations are much smaller. The mean reconstruction errors for all three frameworks are shown in Tab. 1. We assessed the reconstruction errors with three similarity measures: with our Euclidean, Log-Euclidean and affine-invariant metrics, we computed the mean distance between original and reconstructed tensors. As can be seen in this table, Log-Euclidean and affine-invariant results are quantitatively slightly better than Euclidean results, independently of the similarity measure considered. This is essentially due to the better reconstruction of the Corpus Callosum in both Riemannian cases.

Regularization

To compare the Euclidean, affine-invariant and Log-Euclidean frameworks, let us begin with a simple example where we restored a noisy synthetic image of tensors. The eigenvalues of the original tensors were

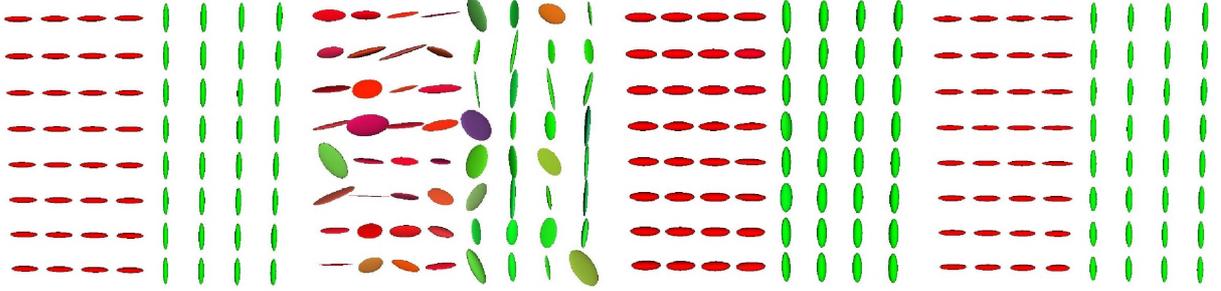


Figure 4: **Regularization of a synthetic DTI slice.** **Left:** original synthetic data. **Middle Left:** noisy data. **Middle Right:** Euclidean regularization. **Right:** Log-Euclidean regularization. The original data is correctly reconstructed in the Log-Euclidean case, as opposed to the Euclidean case where the result is marred by the swelling effect.

set to $(2, 1, 1)$. We added some isotropic Gaussian white noise of variance 0.5 on the b_0 image and each of the 6 synthetic diffusion-weighted images, and tensors were estimated with the method presented in (33) with parameters $\lambda = 0.25$ and $\kappa = 0.1$ (the regularization was small). Results are shown in Fig. 4: surprisingly, although no anisotropic filtering other than the one described in the Methods Section was used, the boundaries between the two regions are kept perfectly distinct, thanks to the strong gradients in this area. Furthermore, the impact of the Euclidean swelling effect is clearly visible. On the contrary, both Riemannian frameworks yield very good results, the only extremely small difference being as predicted slightly more anisotropy for Log-Euclidean results. Affine-invariant results are not shown here because they are very close to the Log-Euclidean ones. Like in the interpolation reconstruction experiment, we assessed the reconstruction errors with the Euclidean, Log-Euclidean and affine-invariant metrics. For each metric, we computed the mean distance between original and reconstructed tensors. The quantitative results are shown in Tab. 2: as expected affine-invariant and Log-Euclidean results are close and yield much better results than in the Euclidean case, regardless of the similarity measure used.

| Similarity Measure | Euclidean regul. | Affine-invariant regul. | Log-Euclidean regul. |
|-----------------------------|------------------|-------------------------|----------------------|
| Mean Euclidean Error | 0.228 | 0.080 | 0.051 |
| Mean Affine-invariant Error | 0.533 | 0.142 | 0.119 |
| Log-Euclidean Error | 0.532 | 0.135 | 0.111 |

Table 2: **Mean reconstruction errors for the synthetic regularization experiment.** Both Riemannian results are much better than the Euclidean one, independently of the similarity measure considered. This is due to the absence of swelling effect in both Riemannian cases.

Let us now turn to a clinical DTI volume, which presents a substantial level of noise. A quantitative evaluation or validation of the restoration results presented here remains to be done, and this general problem will be the subject of future work. However, as shown in Fig. 5, both Riemannian results are qualitatively satisfactory: the smoothing is done without blurring the edges in both Riemannian cases, contrary to the Euclidean results which are marred by a pronounced swelling effect, especially in the regions

of high anisotropy. Also note that to a lesser degree, this swelling effect is present in regions with much less anisotropy, in fact almost everywhere except in the ventricles. The affine-invariant and Log-Euclidean results are very similar to each other, with only slightly more anisotropy in the Log-Euclidean case.

To highlight this similarity, we display in Fig. 5 the *absolute values* of the (Euclidean) differences between affine-invariant and Log-Euclidean results. The definition of the absolute value of a symmetric matrix is given in the Methods section, and this mathematical tool is much useful to visualize the difference between two tensors. We can see in Fig. 5 that the differences are mainly concentrated along the dominant directions of diffusion, which is explained by the larger anisotropy in Log-Euclidean means. However, this relative difference is very small, of the order of less than 1%.

A regularization of DT images should not only correctly regularize the determinants of tensors, but also adequately regularize other scalar measures associated to tensors. In Fig. 6, the effect of the Log-Euclidean regularization on the fractional anisotropy (FA) and on the norm of the gradient are shown. In this experiment, only half of the regularization used to obtain the results of Fig. 5 is kept. As one can see, the regularization, which is performed directly on the tensors, induces a regularization of the FA and gradient norm. Qualitatively, major anisotropic structures have been preserved, including for example the internal capsule, while the noise has been substantially reduced.

As in the case of interpolation, the simpler Log-Euclidean computations are also significantly faster: our current implementation in C++ requires for 100 iterations 30 minutes in the Log-Euclidean case instead of 122 minutes for affine-invariant results on a Pentium Xeon 2.8 GHz with 1 Go of RAM. Our implementation has not been optimized yet and will be improved in the near future. Consequently, the values given here are only upper bounds of what can be achieved. However, the difference in computation times is typical and Log-Euclidean computations can even be 6 or 7 times faster than their affine-invariant equivalent (22).

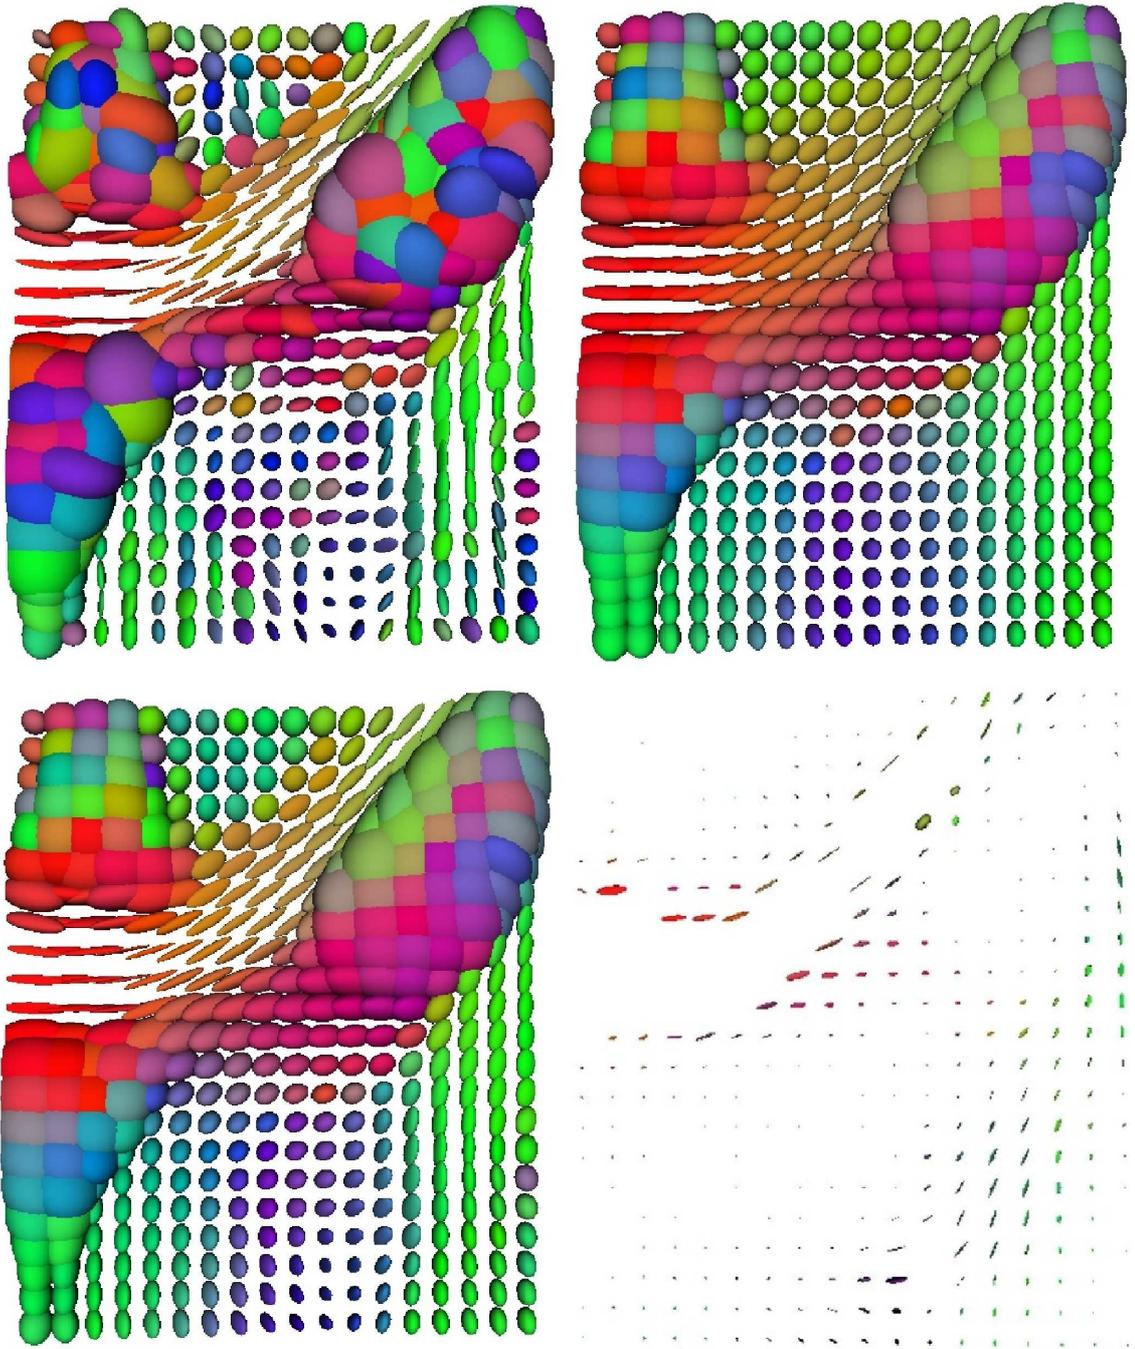


Figure 5: **Regularization of a clinical DTI volume (3D).** **Top Left:** close-up on a slice containing part of the left ventricle and nearby. **Top Right:** Euclidean regularization. **Bottom Left:** Log-Euclidean regularization. **Bottom Right:** highly magnified view ($\times 100$) of the absolute value of the difference between Log-Euclidean and affine-invariant results. The absolute value of tensors is taken to allow the simultaneous visualization of the amplitude *and* orientation of the differences. See the Methods section for a definition of the absolute value. Note that there is *no* tensor swelling in the Riemannian cases. On the contrary, in the Euclidean case, a swelling effect occurs almost everywhere (except maybe in the ventricles), in particular in regions of high anisotropy. Last but not least, the difference between Log-Euclidean and affine-invariant results is very small. Log-Euclidean results are only slightly more anisotropic than their affine-invariant counterparts.

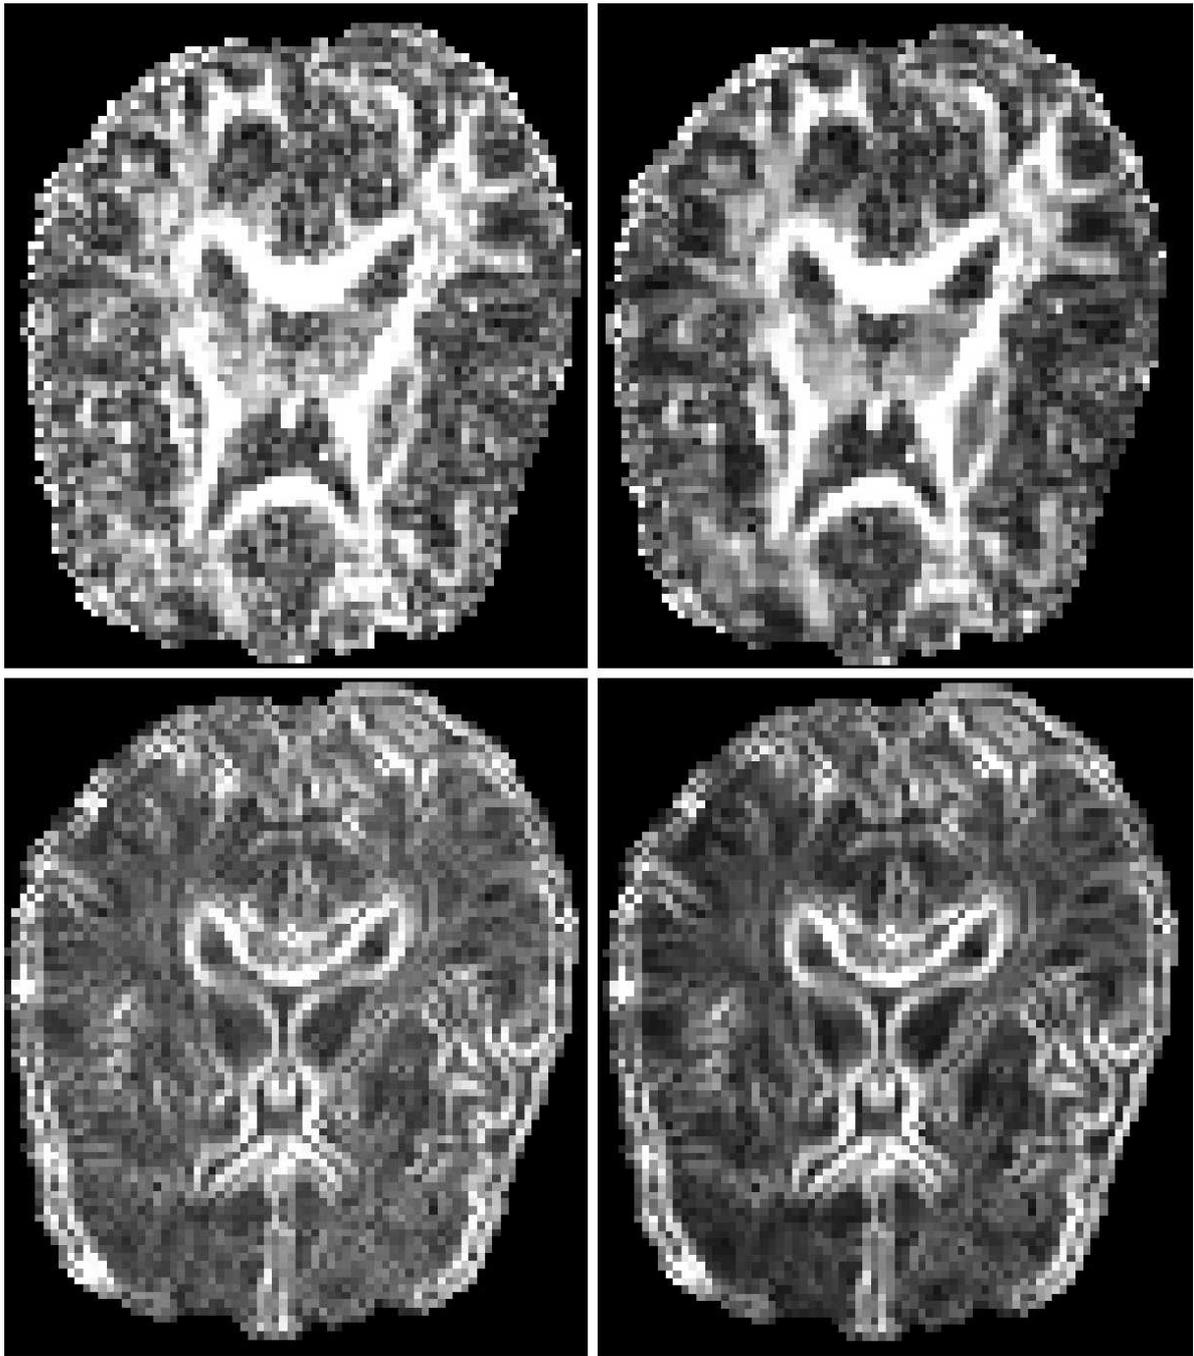


Figure 6: **Log-Euclidean regularization of a clinical DTI volume (3D): typical effect on FA and gradient.** **Top left:** FA before Log-Euclidean regularization. **Top right:** FA after regularization. **Bottom left:** Log-Euclidean norm of the gradient before regularization. **Bottom right:** Log-Euclidean norm of the gradient after regularization. The effect of the Log-Euclidean regularization on scalar measures like FA and the norm of the gradient is qualitatively satisfactory: the noise has been reduced while most structures are preserved.

DISCUSSION AND CONCLUSIONS

The Defects of Euclidean Calculus

As shown in the Results section, Log-Euclidean metrics correct the defects of the classical Euclidean framework (20): the positive-definiteness is preserved and determinants are monotonically (geometrically, in fact) interpolated along geodesics. Log-Euclidean results are very similar to those obtained in the affine-invariant framework, only recently introduced for diffusion tensor calculus (14–17). This is not surprising: we have shown that the two families of metrics are very close, since their respective Fréchet means are *both* generalizations to tensors of the geometric mean of positive numbers. Yet, these two metrics are different, and it is striking that this similarity in results is obtained with much simpler *and* faster algorithms in the Log-Euclidean case. This comes from the fact that all Log-Euclidean computations on tensors are equivalent to Euclidean computations on the logarithms of tensors, which are simple vectors.

Of course, this large simplification is obtained at the cost of affine-invariance, which is replaced by similarity-invariance for a number of Log-Euclidean metrics, like the one used in this study. This means that affine-invariant results cannot be biased by the coordinate system chosen, whereas Log-Euclidean results potentially can. However, invariance by similarity is already a strong property, since it guarantees that computations are not biased neither by the spatial orientation nor by the spatial scale chosen. Moreover, the very large similarity between the Log-Euclidean and affine-invariant results on typical clinical DT images show that this loss of invariance does not result in any significant loss of quality. One would have to change the system of coordinates very anisotropically, for instance rescaling one coordinate by a factor of 20 and leaving the other two unchanged, to substantially bias Log-Euclidean results. But such situations do not occur in medical imaging, where the usual changes of coordinates (e.g. changing current coordinates to Talairach coordinates) are not anisotropic enough to induce such a bias.

In terms of regularization, the Log-Euclidean framework also has the advantage of taking into account simultaneously *all* the information carried by tensors, like the affine-invariant one. This is not the case in methods based on the regularization of *features* extracted from tensors, like their dominant direction of diffusion (18) or their orientation (11). An alternative representation of tensors are Cholesky factors, which are used in (37). However, with this representation, tensors can leave the set of positive definite matrices during iterated computations, and the positive-definiteness is not easily maintained, as mentioned in (37). Also, it is unclear how the smoothing of Cholesky factors affect tensors, whereas the smoothing of tensor logarithms can be interpreted as a *geometric* regularization of tensors which geometrically smoothes determinants.

In this article, we have presented results obtained only with one particular Log-Euclidean metric, inspired from the classical Frobenius norm on matrices. The relevance of this particular choice will be investigated in future work. This is necessary, because it has been shown (38) that the choice of Euclidean metric on tensors can substantially influence the registration of DT images. This should also be the case in the Log-Euclidean framework.

Last but not least, in this work, we have assumed that diffusion tensors are positive-definite. This assumption is consistent with the choice of Brownian motion to model the motion of water molecules. It

could be argued that our framework does not apply to diffusion tensors which have been estimated without taking into account this constraint, and can therefore have non-positive eigenvalues. But these non-positive eigenvalues are difficult to interpret from a physical point of view, and are essentially due to the noise corrupting DW-MRIs! The problem lies therefore in the estimation method and not in our framework. Non-positive eigenvalues can be avoided for example by using a simultaneous estimation and smoothing of tensors, which relies on spatial correlations between tensors to reduce the amount of noise. In this work, we have used the method described in (33), which was inspired by the approach developed in (37).

Conclusions and Perspectives

In this work, we have presented a particularly simple and efficient Riemannian framework for diffusion tensor calculus. Based on Log-Euclidean metrics on the tensor space, this framework transforms Riemannian computations on tensors into Euclidean computations on vectors in the domain of matrix logarithms. As a consequence, classical statistical tools and PDEs usually reserved to vectors are simply and efficiently generalized to tensors in the Log-Euclidean framework.

In this article, we only focus on two important tasks: the interpolation and the regularization of tensors. But this metric approach can be effectively used in all situations where diffusion tensors are processed. Indeed, efficient Log-Euclidean extrapolation techniques are presented in (22), as well as the Log-Euclidean statistical framework for tensors. In this framework, for instance, a Gaussian distribution of random tensors is given by the exponential of a classical Gaussian in the vector space of symmetric matrices. Another important task is the estimation of tensors from DW-MRIs. Adapting ideas from (37) to the Log-Euclidean framework, we have completed a joint estimation and regularization of diffusion tensors directly from the Stejskal-Tanner equations (33). This joint estimation and smoothing is largely facilitated by the Log-Euclidean framework because all computations are carried out in a vector space.

In future work, we will study in further detail the restoration of noisy DT images. In particular, we plan to quantify the impact of the regularization on the tracking of fibers in the white matter of the human nervous system. We also intend to use these new tools to model and reconstruct better the anatomical variability of the human brain with tensors as we began to do in (34). Last but not least, the generalization of our approach to more sophisticated models of diffusion like generalized diffusion tensors (39) or Q-balls (40) is a challenging task we plan to investigate.

ACKNOWLEDGMENTS

The authors thank Denis Ducreux, MD, Kremlin-Bicêtre Hospital (France), for the DT-MRI data he kindly provided for this study.

A patent is pending for the general Log-Euclidean processing framework on tensors (French filing number 0503483, 7th of April, 2005).

References

1. Basser PJ, Mattiello J, Le Bihan D. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 1994;66:259–267.
2. Mori S, Kaufmann WE, Davatzikos C, Stieltjes B, Amodei L, Fredericksen K, Pearlson GD, Mehlem ER, Solaiyappan M, Raymond GV, Moser HW, van Zijl PC. Imaging cortical association tracts in the human brain using diffusion-tensor-based axonal tracking. *Magnetique Resonance in Medecine* 2002; 47:215–223.
3. Lenglet C, Deriche R, Faugeras O. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In T Pajdla, J Matas, eds., *Proc. of the 8th European Conference on Computer Vision, LNCS*. Springer, 2004; 127–140.
4. Fillard P, Gilmore J, Piven J, Lin W, Gerig G. Quantitative Analysis of White Matter Fiber Properties along Geodesic Paths. volume 2879 of *LNCS*. Springer, 2003; 16–23.
5. Vemuri BC, Chen Y, Rao M, McGraw T, Wang Z, Mareci T. Fiber tract mapping from diffusion tensor MRI. In *Proceedings of the IEEE Workshop on Variational and Level Set Methods (VLSM'01)*. IEEE, 2001; 81–88.
6. Basser PJ, Pajevic S, Pierpaoli C, Duda J, Aldroubi A. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 2000;44:625–632.
7. Poupon C, Clark CA, Frouin V, Regis J, Bloch I, LeBihan D, Mangin JF. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *Neuroimage* 2000;12(2):184–95.
8. Le Bihan D, Mangin JF, Poupon C, Clark CA, Pappata S, Molko N, Chabriat H. Diffusion tensor imaging: Concepts and applications. *Journal of Magnetic Resonance Imaging* 2001;13:534–546.
9. Basser PJ, Pajevic S. A normal distribution for tensor-valued random variables: Applications to diffusion tensor MRI. *IEEE Transactions on Medical Imaging* 2003;22(7):785–794.
10. Westin CF, Maier SE, Mamata H, Nabavi A, Jolesz FA, Kikinis R. Processing and visualization of diffusion tensor MRI. *Medical Image Analysis* 2002;6:93–108.
11. Ched'hotel C, Tschumperlé D, Deriche R, Faugeras O. Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision* 2004;20(1-2):147–162.
12. Pennec X. Probabilities and statistics on Riemannian manifolds: Basic tools for geometric measurements. In A Cetin, L Akarun, A Ertuzun, M Gurcan, Y Yardimci, eds., *Proc. of Nonlinear Signal and Image Processing (NSIP'99)*, volume 1. June 20-23, Antalya, Turkey: IEEE-EURASIP, 1999; 194–198.
13. Gallot S, Hulin D, Lafontaine J. *Riemannian Geometry*. Springer-Verlag, 2nd edition edition, 1993.

14. Pennec X, Fillard P, Ayache N. A Riemannian framework for tensor computing. *International Journal of Computer Vision* 2006;66(1):41–66. A preliminary version appeared as INRIA Research Report 5255, July 2004.
15. Batchelor PG, Moakher M, Atkinson D, Calamante F, Connelly A. A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine* 2005;53:221–225.
16. Lenglet C, Rousson M, Deriche R, Faugeras O. Statistics on multivariate normal distributions: A geometric approach and its application to diffusion tensor MRI. Research Report RR-5242 and RR-5243, INRIA, 2004.
17. Fletcher PT, Joshi SC. 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. Springer, 2004; 87–98.
18. Coulon O, Alexander D, Arridge S. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis* 2004;8(1):47–67.
19. Feddern C, Weickert J, Burgeth B, Welk M. Curvature-driven PDE methods for matrix-valued images. Technical Report 104, Department of Mathematics, Saarland University, Saarbrücken, Germany, 2004.
20. Tschumperlé D, Deriche R. Diffusion tensor regularization with constraints preservation. In *Conference on Computer Vision and Pattern Recognition (CVPR), volume I. Kauai, Hawaii, 2001*; 948–953.
21. Moakher M. A differential geometry approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications* 2005;26:735–747.
22. Arsigny V, Fillard P, Pennec X, Ayache N. Fast and simple computations on tensors with Log-Euclidean metrics. Research Report RR-5584, INRIA, May 2005.
23. Culver WJ. On the existence and uniqueness of the real logarithm of a matrix. *Proceedings of the American Mathematical Society* 1966;17(5):1146–1151.
24. Bourbaki N. *Elements of Mathematics: Lie Groups and Lie Algebra, Chapters 1-3*. Springer-Verlag, 2nd printing edition, 1989.
25. Tarantola A. *Elements for Physics - Quantities, Qualities, and Intrinsic Theories*. Springer Verlag, 2006.
26. Pennec X. *L'incertitude dans les problèmes de reconnaissance et de recalage – Applications en imagerie médicale et biologie moléculaire*. Thèse de sciences (PhD thesis), Ecole Polytechnique, Palaiseau (France), 1996.
27. Pennec X. Computing the mean of geometric features - application to the mean rotation. Research Report RR-3371, INRIA, 1998.

28. Moakher M. Means and averaging in the group of rotations. *SIAM J Matrix Anal Appl* 2002;24(1):1–16.
29. Woods RP. Characterizing volume and surface deformations in an atlas framework: theory, applications, and implementation. *Neuroimage* 2003;18(3):769–88.
30. Pennec X. Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision* 2006; To appear. A preliminary version is available as INRIA RR-5093, January 2004.
31. Jones DK, Griffin LD, Alexander DC, Catani M, Horsfield MA, Howard R, Williams SCR. Spatial normalization and averaging of diffusion tensor MRI data sets. *NeuroImage* 2002;17:592–617.
32. Tschumperlé D, Deriche R. Vector-valued image regularization with PDE's : A common framework for different applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2005; 27(4):506–517.
33. Fillard P, Arsigny V, Pennec X, Ayache N. Joint estimation and smoothing of clinical DT-MRI with a Log-Euclidean metric. Research Report RR-5607, INRIA, Sophia-Antipolis, France, 2005.
34. Fillard P, Arsigny V, Pennec X, Thompson PM, Ayache N. Extrapolation of sparse tensor fields: Application to the modeling of brain variability. In G Christensen, M Sonka, eds., *Proc. of Information Processing in Medical Imaging 2005 (IPMI'05)*, volume 3565 of *LNCS*. Glenwood springs, Colorado, USA: Springer, 2005; 27–38.
35. Basser P, Pierpaoli C. A simplified method to measure the diffusion tensor from seven MR images. *Magnetique Resonance in Medecine* 1998;39:928–934.
36. Haselgrove JC, Moore JR. Correction of distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medecine* 1996;36:960–964.
37. Wang Z, Vemuri BC, Chen Y, Mareci TH. A constrained variational principle for simultaneous smoothing and estimation of the diffusion tensors from complex DWI data. *IEEE TMI* 2004;23(8):930–939.
38. Zhang H, Yushkevich PA, Gee JC. Towards diffusion profile image registration. In *ISBI. 2004*; 324–327.
39. Özarıslan E, Mareci TH. Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetique Resonance in Medecine* 2003;50:955–965.
40. Tuch DS. Q-ball imaging. *Magnetic Resonance in Medecine* 2004;52:1358–1372.