

Fast and Simple Calculus on Tensors in the Log-Euclidean Framework

Vincent Arsigny¹, Pierre Fillard¹, Xavier Pennec¹, and Nicholas Ayache¹

INRIA Sophia - Projet Epidaure, BP 93, 06902 Sophia Antipolis Cedex, France
{Vincent.Arsigny, Pierre.Fillard, Xavier.Pennec,
Nicholas.Ayache}@Sophia.Inria.fr

Abstract. Computations on tensors have become common with the use of DT-MRI. But the classical Euclidean framework has many defects, and affine-invariant Riemannian metrics have been proposed to correct them. These metrics have excellent theoretical properties but lead to complex and slow algorithms. To remedy this limitation, we propose new metrics called Log-Euclidean. They also have excellent theoretical properties and yield similar results in practice, but with much simpler and faster computations. Indeed, Log-Euclidean computations are Euclidean computations in the domain of matrix logarithms. Theoretical aspects are presented and experimental results for multilinear interpolation and regularization of tensor fields are shown on synthetic and real DTI data.

1 Introduction: Calculus on Tensors

Tensors, i.e. symmetric positive-definite matrices in medical imaging, appear in many contexts: Diffusion Tensor MRI (DT-MRI or DTI) [2], modeling of anatomical variability [7], etc. They are also a general tool in image analysis, especially for segmentation, motion and texture analysis (see [1] for references on this subject). Many approaches have been proposed in the literature to process tensors [13, 15, 14, 4, 5]. But in order to carry out general computations on these objects, one needs a consistent operational framework. This is necessary to completely generalize to tensors statistical tools and Partial Differential Equations (PDEs). The framework of Riemannian metrics has recently emerged as particularly adapted to this task [11].

One can directly use a Euclidean structure on square matrices to define a metric on the tensor space, for instance with the following distance: $\text{dist}^2(S_1, S_2) = (\text{Trace}((S_1 - S_2)^2))$. This is straightforward and leads *a priori* to simple computations. But this framework is practically and theoretically unsatisfactory for three main reasons. First, symmetric matrices with null or negative eigenvalues appear during Euclidean computations. And from a physical point of view, in DTI, a diffusion exactly equal to zero is impossible: above 0°Kelvin, water molecules *will* move in *all* directions. Even worse, a negative diffusion is meaningless. This occurs during iterated Euclidean computations, for instance during the estimation of tensors from diffusion-weighted images, the regularization of tensors fields, etc. To avoid going out of the tensor space, it has been proposed

to regularize only *features* extracted from tensors, like first eigenvectors [5] or orientations [4]. The regularization is propagated to tensors in a second step. This is not completely satisfactory, since it would be preferable to regularize tensors *directly* in order to take into account *all* the information they carry.

Second, a tensor corresponds typically to a *covariance matrix*. The value of its determinant is a direct measure of the dispersion of the associated multivariate Gaussian. The reason is that the volume of associated confidence regions are proportional to the square root of the covariance determinant. But the Euclidean averaging of tensors leads very often to a *tensor swelling effect*: the determinant (and thus the dispersion) of the Euclidean mean can be larger than the original determinants! In DTI, diffusion tensors are assumed to be covariance matrices of the local Brownian motion of water molecules. Introducing more dispersion in computations amounts to introducing more diffusion, which is physically unacceptable. For illustrations of this effect, see [4].

Third, the Euclidean metric is unsatisfactory in terms of *symmetry* with respect to matrix inversion. The Euclidean mean for tensors is an *arithmetic* mean which does not yield the identity for a tensor and its matrix inverse. When tensors model variability, one would rather have in many cases a *geometric* mean.

To fully circumvent these difficulties, affine-invariant Riemannian metrics have been recently proposed for tensors in [12, 8–10]. With them, negative and null eigenvalues are at an infinite distance, the swelling effect disappears and the symmetry with respect to inversion is respected. The price paid for this success is a high computational burden, essentially due to the curvature induced on the tensor space. Practically, this yields slow and hard to implement algorithms.

We propose a new Riemannian framework to fully overcome these computational limitations while preserving excellent theoretical properties. It is based on new metrics named *Log-Euclidean*, which are particularly simple to use. They result in classical Euclidean computations in the domain of matrix logarithms. In Section 2, we give an overview of the theory of Log-Euclidean metrics, detailed in [1]. In particular, we briefly compare these new metrics to affine-invariant metrics. In Section 3, we highlight the differences between the three frameworks with experimental results on synthetic and real DT-MRI data in interpolation and regularization. Computations are very simple and experimentally much faster in the Log-Euclidean than in the affine-invariant framework.

2 Presentation of the Log-Euclidean Framework

Complete proofs for all the results presented in this Section are given in [1].

Existence and Uniqueness of the Logarithm A tensor S has a unique *symmetric matrix logarithm* $L = \log(S)$. It verifies $S = \exp(L)$ where \exp is the matrix exponential. Conversely, each symmetric matrix is associated to a tensor by the exponential. L is obtained from S by changing its eigenvalues into their natural logarithms, which can be done easily in an orthonormal basis in which S (and L) is diagonal.

A Vector Space Structure on 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 addition “+” and the scalar multiplication “.” with the matrix exponential. This defines on tensors the *logarithmic multiplication* \odot and the *logarithmic scalar multiplication* \otimes , given by:

$$\begin{cases} S_1 \odot S_2 \stackrel{def}{=} \exp(\log(S_1) + \log(S_2)) \\ \lambda \otimes S \stackrel{def}{=} \exp(\lambda \cdot \log(S)) = S^\lambda. \end{cases} \quad (1)$$

The logarithmic multiplication is *commutative* and coincides with matrix multiplication whenever the two tensors S_1 and 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 inherited from square matrices.

Log-Euclidean Metrics When one considers only the multiplication \odot on the tensor space, one has a *Lie group* structure [11], i.e. a space which is both a smooth manifold and a group in which algebraic operations are *smooth mappings*. Among Riemannian metrics in Lie groups, the most convenient in practice, when they exist, are *bi-invariant metrics*, i.e. distances that are invariant by multiplication and inversion. For the tensor Lie group, bi-invariant metrics exist and are particularly simple. 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:

$$\text{dist}(S_1, S_2) = \|\log(S_1) - \log(S_2)\|. \quad (2)$$

Boundary Problems and Symmetry Contrary to the classical Euclidean framework on tensors, one can clearly see from Eq. (2) that matrices with null or negative eigenvalues are at an infinite distance from tensors and will not appear in practical computations. Moreover, distances are not changed by inversion.

Invariance by Similarity Log-Euclidean metrics are not affine-invariant. However, some of them are invariant by similarity (orthogonal transformation and scaling). This means that if tensors are covariance matrices, computations on tensors using these metrics will be invariant with respect to a change of coordinates obtained by a similarity. The similarity-invariant Log-Euclidean metric used throughout this article is given by:

$$\text{dist}(S_1, S_2) = \left(\text{Trace}(\{\log(S_1) - \log(S_2)\}^2)\right)^{\frac{1}{2}}. \quad (3)$$

Euclidean Calculus in the Logarithmic Domain The tensor vector space 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 [1]. In particular, the Log-Euclidean mean of N tensors with arbitrary positive weights

$(w_i)_{i=1}^N$ such that $\sum_{i=1}^N w_i = 1$ is a direct generalization of the geometric mean of positive numbers and is given explicitly by:

$$\mathbb{E}_{LE}(S_1, \dots, S_N) = \exp\left(\sum_{i=1}^N w_i \log(S_i)\right). \quad (4)$$

This is remarkable: *in this framework, the processing of tensors is simply Euclidean in the logarithmic domain. Final results obtained on logarithms are mapped back to the tensor domain with the exponential.* Hence, statistical tools or PDEs are readily generalized to tensors in this framework.

Comparison with Affine-Invariant Metrics As shown experimentally in Section 3, Log-Euclidean computations provide results similar to their affine-invariant equivalent, presented in [12]. The reason is the two families of metrics provide two generalizations of the geometric mean of positive numbers on tensors. Contrary to the Log-Euclidean mean, there is in general no closed form for the affine-invariant mean but rather a *barycentric equation*. Nevertheless, the determinants of the two means are both equal to the scalar geometric mean of the determinants of the averaged tensors [1]. This explains their likeness and the absence of swelling effect in both cases. This resemblance between the two means propagates to general computations which involve averaging, such as interpolation, extrapolation and regularization. The two means are even identical in a number of cases, in particular when averaged tensors commute. Yet they are *not* equal in general: Log-Euclidean means are slightly more anisotropic.

3 Experimental Results

3.1 Bilinear and Trilinear Interpolation

Often, voxels in clinical DT images are quite anisotropic. But algorithms tracking white matter fascicles are more efficient with isotropic voxels [3]. An adequate interpolation method is therefore important for such algorithms.

Fig. 1 shows the results obtained for the bilinear interpolation of four tensors with three methods: Euclidean (linear interpolation of coefficients), affine-invariant and Log-Euclidean bilinear interpolations. There is a pronounced swelling effect in the Euclidean case, which is not physically acceptable. On the contrary, both Riemannian interpolations provide the same geometric interpolation of determinants. There is a slightly larger anisotropy in Log-Euclidean means, which is a general effect discussed in [1]. The computation of the affine-invariant mean is iterative (we use the Gauss-Newton method described in [12]), whereas the closed form given by Eq. (4) 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.

To compare the Euclidean and Riemannian bilinear interpolations on real data, we have reconstructed by bilinear interpolation a down-sampled DTI slice.

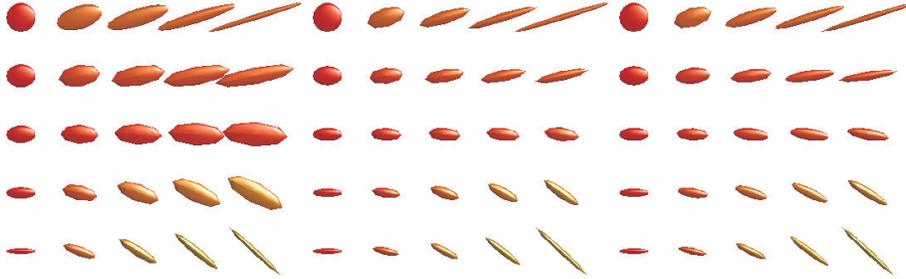


Fig. 1. Bilinear interpolation of 4 tensors at the corners of a grid. Left: Euclidean interpolation. **Middle:** affine-invariant interpolation. **Right:** Log-Euclidean interpolation. Note the characteristic swelling effect observed in the Euclidean case, which is not present in both Riemannian frameworks. Note also that Log-Euclidean means are slightly more anisotropic than their affine-invariant counterparts. The coloring of ellipsoids is based on the direction of dominant eigenvectors.

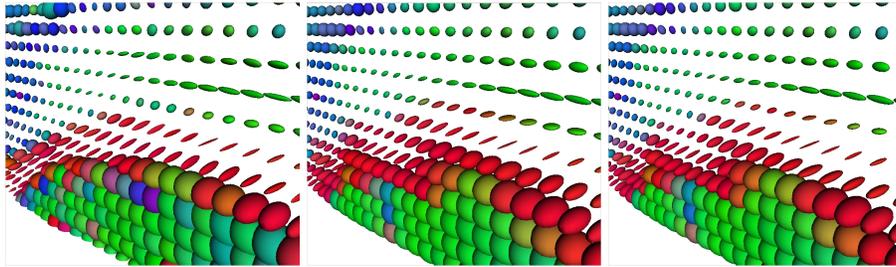


Fig. 2. Bilinear interpolation in a real DTI slice. Left: Original DTI slice, before down-sampling. **Middle:** Euclidean interpolation. **Right:** Log-Euclidean interpolation. Half the columns and lines of the original DTI slice were removed before reconstruction with a bilinear interpolation. The slice is taken in the mid-sagittal plane and displayed in perspective. 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.

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. 2 show that the tensors corresponding to the corpus callosum are better reconstructed in the Log-Euclidean case. Affine-invariant results are very close to Log-Euclidean results and not shown here.

3.2 Regularization of Tensor Fields

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 [12], Riemannian metrics provide a general framework to provide such a regularization. We focus here on a typical Riemannian crite-

tion for the regularization of tensor fields. An anisotropic regularization can be obtained by the minimization of a Φ -functional [4] $Reg(S)$ given by:

$$Reg(S) = \int_{\text{Image}} \Phi(\|\nabla S\|_{S(x)}(x)) dx. \quad (5)$$

$\Phi(s)$ is a function penalizing large values of the norm of the *spatial gradient* ∇S of the image. Here, we use $\Phi(s) = \sqrt{1 + s^2/\kappa^2}$. Contrary to the Euclidean case, the norm of ∇S *depends explicitly on the current point* $S(x)$ and is given by:

$$\|\nabla S\|_{S(x)}^2 = \sum_{i=1}^d \left\| \frac{\partial S}{\partial x_i}(x) \right\|_{S(x)}^2. \quad (6)$$

In general and in particular in the affine-invariant case, this dependence on the current point leads to complex resolution methods. Practically, this implies in the affine-invariant case an intensive use of matrix inverses, square roots, exponentials and logarithms [12]. But in the Log-Euclidean framework, the general Riemannian formulation is extremely simplified. The reason is that the dependence on the current tensor disappears on logarithms [1]:

$$\|\nabla S\|_{S(x)} = \|\nabla \log(S)\|. \quad (7)$$

Consequently, *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. Interestingly, mathematical issues such as existence and uniqueness of PDEs on tensors in the Log-Euclidean framework are simply particular cases of the classical theory of PDEs on vector-valued images.

In the following experiments, 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 and the geodesic marching scheme described in [12] 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 [4]. Homogeneous Neumann boundary conditions are used, $\kappa = 0.05$, $dt = 0.1$ and 100 iterations are performed.

As a first example, we restore a noisy synthetic image of tensors. Results are shown in Fig. 3: the negative impact of the Euclidean swelling effect is clearly visible. On the contrary, both Riemannian frameworks yield proper results, the only (small) difference being slightly more anisotropy for Log-Euclidean results.

Let us now turn to a real DTI volume of the brain with $128 \times 128 \times 30$ voxels with spatial dimensions of $1.875 \times 1.875 \times 4 \text{ mm}^3$. The b-value is 1000 s.mm^{-2} . As shown in Fig. 4, both Riemannian results are qualitatively very satisfactory: the smoothing is done without blurring the edges. They are also very similar to each other, with only slightly more anisotropy in the Log-Euclidean case. As before, the Euclidean results are marred by a pronounced swelling effect. Computations are much faster in the Log-Euclidean case: 30 minutes instead of 122 minutes for affine-invariant results on a Pentium Xeon 2.8 GHz with 1 Go of RAM.

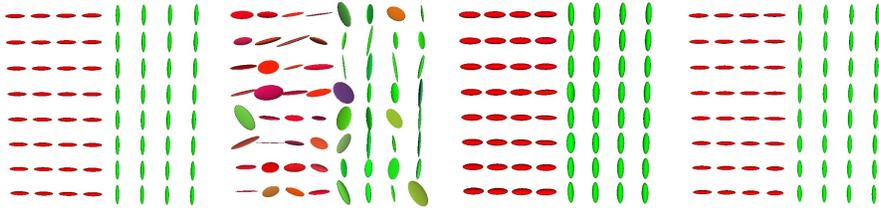


Fig. 3. Regularization of a synthetic DTI slice. **Left:** original synthetic data. **Middle Left:** noisy data. **Middle Right:** Euclidean regularization. **Right:** Log-Euclidean regularization. The original data is properly reconstructed in the Log-Euclidean case, as opposed to the Euclidean case where the result is marred by the swelling effect.

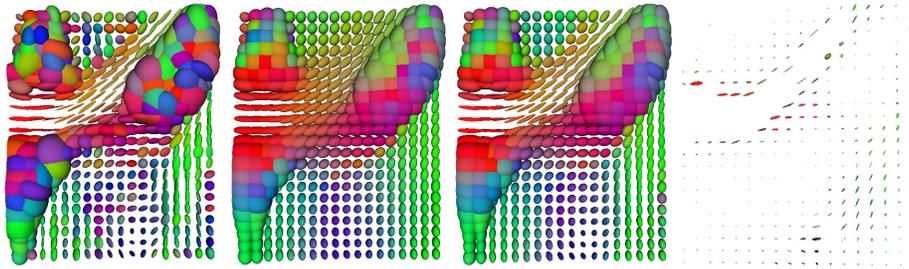


Fig. 4. Regularization of a real DTI volume. **Left:** close-up on the top right ventricle and nearby. **Middle Left:** Euclidean regularization. **Middle Right:** Log-Euclidean regularization. **Right:** highly magnified view ($\times 100$) of the absolute value (the absolute value of eigenvalues is taken) of the difference between Log-Euclidean and affine-invariant results. Note that there is *no* tensor swelling in the Riemannian cases, contrary to the Euclidean case. Log-Euclidean and affine-invariant results are very similar, the only difference being slightly more anisotropy in Log-Euclidean results.

4 Discussion and Perspectives

In this work, we have presented a particularly simple and efficient Riemannian framework for tensor calculus, called Log-Euclidean. As in the affine-invariant case, the defects of the Euclidean framework are corrected with Log-Euclidean metrics, but without any unnecessary technicality. Indeed, Riemannian computations on tensors are converted into Euclidean computations on vectors in this novel framework. In practice, classical statistical tools and PDEs for vectors can be directly used on the matrix logarithms of tensors, which are simple vectors. Moreover, all usual operations on tensors can be efficiently carried out in this framework, like the joint estimation and smoothing of DTI from diffusion-weighted images, as shown in [6].

In future work, we will study in further details 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 this new framework to better model and reconstruct the anatomical variability of the human brain with tensors [7].

Acknowledgments and Pending Patent

The authors thank Denis Ducreux, MD, Kremlin-Bicêtre Hospital (France), for the DT-MRI data he kindly let us use in this work.

A patent is pending for the Log-Euclidean processing framework of tensors (Deposit Number 0503483, April 7th, 2005, property of INRIA, France).

References

1. V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Fast and simple computations on tensors with Log-Euclidean metrics. Research Report 5584, INRIA, May 2005.
2. P. Basser, J. Mattiello, and D. Le Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
3. P. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
4. C. Chefd’hotel, D. Tschumperlé, R. Deriche, and O. Faugeras. Regularizing flows for constrained matrix-valued images. *J. Math. Im. Vis.*, 20(1-2):147–162, 2004.
5. O. Coulon, D. Alexander, and S. Arridge. Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67, 2004.
6. P. Fillard, V. Arsigny, X. Pennec, and N. Ayache. Joint estimation and smoothing of clinical DT-MRI with a Log-Euclidean metric. Research Report RR-5607, INRIA, Sophia-Antipolis, France, June 2005.
7. P. Fillard, V. Arsigny, X. Pennec, P. Thompson, and N. Ayache. Extrapolation of sparse tensor fields: Application to the modeling of brain variability. In G. Christensen and M. Sonka, editors, *Proc. of IPMI’05*, LNCS, Glenwood springs, Colorado, USA, July 2005. Springer. To appear.
8. 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.
9. C. Lenglet, M. Rousson, R. Deriche, and O. Faugeras. Statistics on multivariate normal distributions: A geometric approach and its application to diffusion tensor MRI. Research Report 5242, INRIA, June 2004.
10. M. Moakher. A differential geometry approach to the geometric mean of symmetric positive-definite matrices. *SIAM Jour. on Mat. Anal. and Appl.*, 2004. To appear.
11. Xavier Pennec. Probabilities and statistics on Riemannian manifolds: a geometric approach. Research Report 5093, INRIA, January 2004.
12. Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A Riemannian framework for tensor computing. Research Report 5255, INRIA, July 2004.
13. C. Poupon, C. A. Clark, V. Frouin, J. Regis, I. Bloch, D. Le Bihan, and J.-F. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *Neuroimage*, 12(2):184–95, August 2000.
14. Z. Wang, B. Vemuri, Y. Chen, and T. H. Mareci. A constrained variational principle for simultaneous smoothing and estimation of the diffusion tensors from complex DWI data. *IEEE TMI*, 23(8):930–939, 2004.
15. C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization of diffusion tensor MRI. *Media*, 6:93–108, 2002.