A Riemannian Framework for Tensor Computing

Xavier Pennec, Pierre Fillard, Nicholas Ayache

► To cite this version:

Xavier Pennec, Pierre Fillard, Nicholas Ayache. A Riemannian Framework for Tensor Computing. International Journal of Computer Vision, 2006, 66 (1), pp.41–66. 10.1007/s11263-005-3222-z . inria-00614990

HAL Id: inria-00614990

<https://hal.inria.fr/inria-00614990>

Submitted on 17 Aug 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Riemannian Framework for Tensor Computing

Xavier Pennec, Pierre Fillard, Nicholas Ayache

EPIDAURE / ASCLEPIOS Project-team, INRIA Sophia-Antipolis
2004 Route des Lucioles BP 93, F-06902 Sophia Antipolis Cedex, FRANCE

February 8 2005

Abstract

Tensors are nowadays a common source of geometric information. In this paper, we propose to endow the tensor space with an affine-invariant Riemannian metric. We demonstrate that it leads to strong theoretical properties: the cone of positive definite symmetric matrices is replaced by a regular and complete manifold without boundaries (null eigenvalues are at the infinity), the geodesic between two tensors and the mean of a set of tensors are uniquely defined, etc. We have previously shown that the Riemannian metric provides a powerful framework for generalizing statistics to manifolds. In this paper, we show that it is also possible to generalize to tensor fields many important geometric data processing algorithms such as interpolation, filtering, diffusion and restoration of missing data. For instance, most interpolation and Gaussian filtering schemes can be tackled efficiently through a weighted mean computation. Linear and anisotropic diffusion schemes can be adapted to our Riemannian framework, through partial differential evolution equations, provided that the metric of the tensor space is taken into account. For that purpose, we provide intrinsic numerical schemes to compute the gradient and Laplace-Beltrami operators. Finally, to enforce the fidelity to the data (either sparsely distributed tensors or complete tensors fields) we propose least-squares criteria based on our invariant Riemannian distance which are particularly simple and efficient to solve.

Keywords: Tensors, Diffusion Tensor MRI, Regularization, Interpolation, Extrapolation, PDE, Riemannian Manifold, Affine-invariant Metric.

1 Introduction

Positive definite symmetric matrices (so-called tensors in this article) are often encountered in image processing, for instance as covariance matrices for characterizing statistics on deformations, or as an encoding of the principal diffusion directions in Diffusion Tensor Imaging (DTI). The measurements of these tensors is often noisy in real applications and we would

like to perform estimation, smoothing and interpolation of fields of this type of features. The main problem is that the tensor space is a manifold that is not a vector space with the usual additive structure. As symmetric positive definite matrices constitute a convex half-cone in the vector space of matrices, many usual operations (like the mean) are stable in this space. However, problems arise when estimating tensors from data (in standard DTI, the estimated symmetric matrix could have negative eigenvalues), or when smoothing fields of tensors: the numerical schemes used to solve the Partial Differential Equation (PDE) may sometimes lead to negative eigenvalues if the time step is not small enough. Even when a SVD is performed to smooth independently the rotation (eigenvectors basis trihedron) and eigenvalues, there is a continuity problem around equal eigenvalues.

In previous works [Pennec, 1996, Pennec and Ayache, 1998], we used invariance requirements to develop some basic probability tools on transformation groups and homogeneous manifolds. This statistical framework was then reorganized and extended in [Pennec, 1999, Pennec, 2004] for general Riemannian manifolds, invariance properties leading in some case to a natural choice for the metric. In this paper, we show how this theory can be applied to tensors, leading to a new intrinsic computing framework for these geometric features with many important theoretical properties as well as practical computing properties.

In the remaining of this section, we quickly investigate some connected works on tensors. Then, we summarize in Section 2 the main ideas of the statistical framework we developed on Riemannian manifolds. The aim is to exemplify the fact that choosing a Riemannian metric “automatically” determines a powerful framework to work on the manifold through the introduction of a few tools from differential geometry. In order to use this Riemannian framework on our tensor manifold, we propose in Section 3 an affine-invariant Riemannian metric on tensors. We demonstrate that it leads to very strong theoretical properties, as well as some important practical algorithms such as an intrinsic geodesic gradient descent. Section 4 focuses on the application of this framework to an important geometric data processing problem: interpolation of tensor values. We show that this problem can be tackled efficiently through a weighted mean optimization. However, if weights are easy to define for regularly sampled tensors (e.g. for linear or tri-linear interpolation), the problem proved to be more difficult for irregularly sampled values.

With Section 5, we turn to tensors field computing, and more particularly filtering. If the Gaussian filtering may still be defined through weighted means, the partial differential equation (PDE) approach is slightly more complex. In particular, the metric of the tensor space has to be taken into account when computing the magnitude of the spatial gradient of the tensor field. Thanks to our Riemannian framework, we propose efficient numerical schemes to compute the gradient, its amplitude, and the Laplace-Beltrami operator used in linear diffusion. We also propose an adjustment of this manifold Laplacian that realizes an anisotropic filtering. Finally, Section 6 focuses on simple statistical approaches to regularize and restore missing values in tensor fields. Here, the use of the Riemannian distance inherited from the chosen metric is fundamental to define least-squares data attachment criteria for dense and sparsely distributed tensor fields that lead to simple implementation schemes in our intrinsic computing framework.

1.1 Related work

Quite an impressive literature has now been issued on the estimation and regularization of tensor fields, especially in the context of Diffusion Tensor Imaging (DTI) [Basser et al., 1994, Le Bihan et al., 2001, Westin et al., 2002]. Most of the works dealing with the geometric nature of the tensors has been performed for the discontinuity-preserving regularization of the tensor fields using Partial Differential Equations (PDEs). For instance, [Coulon et al., 2004] anisotropically restores the principal direction of the tensor, and uses this regularized directions map as an input for the anisotropic regularization of the eigenvalues. A quite similar idea is adopted in [Tschumperlé, 2002], where a spectral decomposition $W(x) = U(x) D(x) U(x)^T$ of the tensor field is performed at each points to independently regularize the eigenvalues and eigenvectors (orientations). This approach requires an additional reorientation step of the rotation matrices due to the non-uniqueness of the decomposition (each eigenvector is defined up its sign and there may be joint permutations of the eigenvectors and eigenvalues) in order to avoid the creation of artificial discontinuities. Another problem arises when two or more eigenvalues become equal: a whole subspace of unit eigenvectors is possible, and even a re-orientation becomes difficult. An intrinsic integration scheme for PDEs that uses the exponential map has been added in [Chefd'hotel et al., 2002], and allows to perform PDEs evolution on the considered manifold without re-projections. In essence, this is an infinitesimal version of the intrinsic gradient descent technique on manifolds we introduced in [Pennec, 1996, Pennec, 1999] for the computation of the mean.

The affine-invariant Riemannian metric we detail in Section 3.3 may be traced back to the work of [Nomizu, 1954] on affine invariant connections on homogeneous spaces. It is implicitly hidden under very general theorems on symmetric spaces in many differential geometry textbooks [Kobayashi and Nomizu, 1969, Helgason, 1978, Gamkrelidze, 1991] and sometimes considered as a well known result as in [Bhatia, 2003]. In statistics, it has been introduced as the Fisher information metric [Skovgaard, 1984] to model the geometry of the multivariate normal family. The idea of the invariant metric came to the mind of the first author during the IPMI conference in 2001 [Coulon et al., 2001, Batchelor et al., 2001], as an application to diffusion tensor imaging (DTI) of the statistical methodology on Riemannian manifolds previously developed (and summarized in the next Section). However, this idea was not exploited until the end of 2003, when the visit of P. Thompson (UCLA, USA) raised the need to interpolate tensors that represent the variability from specific locations on sulci to the whole volume. The expertise of the second author on DTI [Fillard et al., 2003] provided an ideal alternative application field. During the writing of this paper, we discovered that the invariant metric has been independently proposed by [Förstner and Moonen, 1999] to deal with covariance matrices, and very recently by [Fletcher and Joshi, 2004] for the analysis of principal modes of sets of diffusion tensors. By looking for a suitable metric on the space of Gaussian distributions for the segmentation of diffusion tensor images, [Lenglet et al., 2004a, Lenglet et al., 2004b] also end-up with the same metric. It is interesting to see that completely different approaches, relying on an affine-invariant requirement on the one hand, and relying on an information measure to evaluate the distance between distributions on the other hand, lead to the same metric on the tensor space. However, to our knowledge, this Riemannian metric has not been promoted as a complete computing

framework, as we propose in this paper.

2 Statistics on geometric features

We summarize in this Section the theory of statistics on Riemannian manifolds developed in [Pennec, 1999, Pennec, 2004]. The aim is to exemplify the fact that choosing a Riemannian metric “automatically” determines a powerful framework to work on the manifold through the use of a few tools from differential geometry.

In the geometric framework, one can specify the structure of a manifold $\mathcal{M}$ by a *Riemannian metric*. This is a continuous collection of scalar products on the tangent space at each point of the manifold. Thus, if we consider a curve on the manifold, we can compute at each point its instantaneous speed vector and its norm, the instantaneous speed. To compute the length of the curve, we can proceed as usual by integrating this value along the curve. The distance between two points of a connected Riemannian manifold is the minimum length among the curves joining these points. The curves realizing this minimum for any two points of the manifold are called geodesics. The calculus of variations shows that geodesics are the solutions of a system of second order differential equations depending on the Riemannian metric. In the following, we assume that the manifold is *geodesically complete*, i.e. that the definition domain of all geodesics can be extended to $\mathbb{R}$. This means that the manifold has no boundary nor any singular point that we can reach in a finite time. As an important consequence, the Hopf-Rinow-De Rham theorem states that there always exists at least one minimizing geodesic between any two points of the manifold (i.e. whose length is the distance between the two points).

Figure 1: **Left:** The tangent planes at points x and y of the sphere S_2 are different: the vectors v and w of $T_x\mathcal{M}$ cannot be compared to the vectors t and u of $T_y\mathcal{M}$. Thus, it is natural to define the scalar product on each tangent plane. **Right:** The geodesics starting at x are straight lines in the exponential map and the distance along them is conserved.

2.1 Exponential chart

Let x be a point of the manifold that we consider as a local reference and $\overrightarrow{xy}$ a vector of the tangent space $T_x\mathcal{M}$ at that point. From the theory of second order differential equations, we

know that there exists one and only one geodesic starting from that point with this tangent vector. This allows to develop the manifold in the tangent space along the geodesics (think of rolling a sphere along its tangent plane at a given point). The geodesics going through the reference point are transformed into straight lines and the distance along these geodesics is conserved (at least in a neighborhood of x).

The function that maps to each vector $\vec{xy} \in T_x\mathcal{M}$ the point y of the manifold that is reached after a unit time by the geodesic starting at x with this tangent vector is called the *exponential map*. This map is defined in the whole tangent space $T_x\mathcal{M}$ (since the manifold is geodesically complete) but it is generally one-to-one only locally around 0 in the tangent space (i.e. around x in the manifold). In the sequel, we denote by $\vec{xy} = \log_x(y)$ the inverse of the exponential map: this is the smallest vector such that $y = \exp_x(\vec{xy})$. If we look for the maximal definition domain, we find out that it is a star-shaped domain delimited by a continuous curve C_x called the *tangential cut-locus*. The image of C_x by the exponential map is the cut locus $\mathcal{C}_x$ of point x . This is the closure of the set of points where several minimizing geodesics starting from x meet. On the sphere $\mathcal{S}_2(1)$ for instance, the cut locus of a point x is its antipodal point and the tangential cut locus is the circle of radius π .

The exponential map within this domain realizes a chart called *the exponential chart*. It covers all the manifold except the cut locus of the reference point x , which has a null measure. In this chart, geodesics starting from x are straight lines, and the distance from the reference point are conserved. This chart is somehow the “most linear” chart of the manifold with respect to the reference point x .

2.2 Practical implementation

In fact, most of the usual operations using additions and subtractions may be reinterpreted in a Riemannian framework using the notion of *bipoint*, an antecedent of vector introduced during the 19th Century. Indeed, one defines vectors as equivalent classes of bipoints (oriented couples of points) in a Euclidean space. This is possible because we have a canonical way (the translation) to compare what happens at two different points. In a Riemannian manifold, we can still compare things locally (by parallel transportation), but not any more globally. This means that each “vector” has to remember at which point of the manifold it is attached, which comes back to a bipoint.

However, one can also see a vector $\vec{xy}$ (attached at point x) as a vector of the tangent space at that point. Such a vector may be identified to a point on the manifold using the geodesic starting at x with tangent vector $\vec{xy}$, i.e. using the exponential map: $y = \exp_x(\vec{xy})$. Conversely, the logarithmic map may be used to map almost any bipoint (x, y) into a vector $\vec{xy} = \log_x(y)$ of $T_x\mathcal{M}$. This reinterpretation of addition and subtraction using logarithmic and exponential maps is very powerful to generalize algorithms working on vector spaces to algorithms on Riemannian manifolds, as illustrated by Table 1. It is also very powerful in terms of implementation since we can practically express all the geometric operations in these terms: the implementation of $\log_x$ and $\exp_x$ is the basis of any programming on Riemannian manifolds, as we will see in the following.

2.3 Basic statistical tools

The Riemannian metric induces an infinitesimal volume element on each tangent space, and thus a measure $d\mathcal{M}$ on the manifold that can be used to measure random events on the manifold and to define the probability density function (if it exists) of these random elements. It is worth noticing that the induced measure $d\mathcal{M}$ represents the notion of *uniformity* according to the chosen Riemannian metric. This automatic derivation of the uniform measure from the metric gives a rather elegant solution to the Bertrand paradox for geometric probabilities [Poincaré, 1912, Kendall and Moran, 1963]. However, the problem is only shifted: which Riemannian metric do we have to choose? We address this question in Section 3 for real positive definite symmetric matrices (tensors): it turns out that requiring an invariance by the full linear group will lead to a very regular and convenient manifold structure.

Let us come back to the basic statistical tools. With the probability measure of a random element, we can integrate functions from the manifold to any vector space, thus defining the expected value of this function. However, we generally cannot integrate manifold-valued functions. Thus, one cannot define the mean or expected “value” of a random manifold element using a weighted sum or an integral as usual. One solution is to rely on a distance-based variational formulation: the Fréchet or Karcher expected features basically minimize globally (or locally) the variance. As the mean is now defined through a minimization procedure, its existence and uniqueness are not ensured any more (except for distributions with a sufficiently small compact support). In practice, one mean value almost always exists, and it is unique as soon as the distribution is sufficiently peaked. The properties of the mean are very similar to those of the modes (that can be defined as central Karcher values of order 0) in the vectorial case.

To compute the mean value, we designed in [Pennec, 1999, Pennec, 2004] an original Gauss-Newton gradient descent algorithm that essentially alternates the computation of the barycenter in the exponential chart centered at the current estimation of the mean value, and a re-centering step of the chart at the point of the manifold that corresponds to the computed barycenter (geodesic marching step). To define higher moments of the distribution, we used the exponential chart at the mean point: the random feature is thus represented as a random vector with null mean in a star-shaped domain. With this representation, there is no difficulty to define the covariance matrix and potentially higher order moments.

	Vector space	Riemannian manifold
Subtraction	$\overrightarrow{xy} = y - x$	$\overrightarrow{xy} = \log_x(y)$
Addition	$y = x + \overrightarrow{xy}$	$y = \exp_x(\overrightarrow{xy})$
Distance	$\text{dist}(x, y) = \ y - x\ $	$\text{dist}(x, y) = \ \overrightarrow{xy}\ _x$
Mean value (implicit)	$\sum_i \overrightarrow{xx_i} = 0$	$\sum_i \log_{\bar{x}}(x_i) = 0$
Gradient descent	$x_{t+\varepsilon} = x_t - \varepsilon \nabla C(x_t)$	$x_{t+\varepsilon} = \exp_{x_t}(-\varepsilon \nabla C(x_t))$
Linear (geodesic) interpolation	$x(t) = x_1 + t \overrightarrow{x_1x_2}$	$x(t) = \exp_{x_1}(t \overrightarrow{x_1x_2})$

Table 1: Re-interpretation of basic standard operations in a Riemannian manifold.

Based on this covariance matrix, we defined a Mahalanobis distance between a random and a deterministic feature that basically weights the distance between the deterministic feature and the mean feature using the inverse of the covariance matrix. Interestingly, the expected Mahalanobis distance of a random element with itself is independent of the distribution and is equal to the dimension of the manifold, as in the vectorial case.

As for the mean, we chose in [Pennec, 1996, Pennec, 1999, Pennec, 2004] a variational approach to generalize the Normal Law: we define it as the distribution that minimizes the information knowing the mean and the covariance. This amounts to consider a Gaussian distribution on the exponential chart centered at the mean point that is truncated at the cut locus (if there is one). However, the relation between the concentration matrix (the “metric” used in the exponential of the probability density function) and the covariance matrix is slightly more complex than the simple inversion of the vectorial case, as it has to be corrected for the curvature of the manifold. Last but not least, using the Mahalanobis distance of a normally distributed random feature, we can generalize the χ^2 law: we were able to show that it has the same density as in the vectorial case up to an order 3 in σ . This opens the way to the generalization of many other statistical tests, as we may expect similarly simple approximations for sufficiently centered distributions.

3 Working on the Tensor space

Let us now focus on the space $\mathcal{Sym}_n^+$ of positive definite symmetric matrices (tensors). The goal is to find a Riemannian metric with interesting enough properties. It turns out that it is possible to require an invariance by the full linear group (Section 3.3). This leads to a very regular manifold structure where tensors with null and infinite eigenvalues are both at an infinite distance of any positive definite symmetric matrix: the cone of positive definite symmetric matrices is replaced by a space which has an infinite development in each of its $n(n+1)/2$ directions. Moreover, there is one and only one geodesic joining any two tensors, and we can even define globally consistent orthonormal coordinate systems of tangent spaces. Thus, the structure we obtain is very close to a vector space, except that the space is curved.

3.1 Exponential, logarithm and square root of tensors

In the following, we will make an extensive use of a few functions on symmetric matrices. The exponential of any matrix can be defined using the series $\exp(A) = \sum_{k=0}^{+\infty} \frac{A^k}{k!}$. In the case of symmetric matrices, we have some important simplifications. Let $W = U D U^T$ be a diagonalization, where U is an orthonormal matrix, and $D = \text{DIAG}(d_i)$ is the diagonal matrix of the eigenvalues. We can write any power of W in the same basis: $W^k = U D^k U^T$. This means that we may factor out the rotation matrices in the series and map the exponential individually to each eigenvalue:

$$\exp(W) = \sum_{k=0}^{+\infty} \frac{W^k}{k!} = U \text{DIAG}(\exp(d_i)) U^T.$$

The series defining the exponential function converges for any (symmetric) matrix argument, but this is generally not the case for the series defining its inverse function: the

logarithm. However, any tensor can be diagonalized into $\Sigma = U \text{DIAG}(d_i) U^T$ with strictly positive eigenvalues d_i . Thus, the function

$$\log(\Sigma) = U (\text{DIAG}(\log(d_i))) U^T$$

is always well defined on tensors. Moreover, if all the eigenvalues are small enough ($|d_i - 1| < 1$), then the series defining the usual log converges and we have:

$$\log(\Sigma) = U \left(\text{DIAG} \left(\sum_{k=1}^{+\infty} \frac{(-1)^{k+1}}{k} (d_i - 1)^k \right) \right) U^T = \sum_{k=1}^{+\infty} \frac{(-1)^{k+1}}{k} (\Sigma - \text{Id})^k. \quad (1)$$

The logarithm we defined is obviously the inverse function of exp. Thus, the matrix exponential realizes a one-to-one mapping between the space of symmetric matrices to the space of tensors.

Classically, one defines the (left) square root of a matrix B as the set $\{B_L^{1/2}\} = \{A \in GL_n / A A^T = B\}$. One could also define the right square root: $\{B_R^{1/2}\} = \{A \in GL_n / A^T A = B\}$. For tensors, we define the square root as:

$$\Sigma^{1/2} = \{\Lambda \in \mathcal{Sym}_n^+ / \Lambda^2 = \Sigma\}.$$

The square root is always defined and moreover unique: let $\Sigma = U D^2 U^T$ be a diagonalization (with positives values for the d_i 's). Then $\Lambda = U D U^T$ is of course a square root of Σ , which proves the existence. For the uniqueness, let us consider two symmetric and positive square roots Λ_1 and Λ_2 of Σ . Then, $\Lambda_1^2 = \Sigma$ and $\Lambda_2^2 = \Sigma$ obviously commute and thus they can be diagonalized in the same basis: this means that the diagonal matrices D_1^2 and D_2^2 are equal. As the elements of D_1 and D_2 are positive, they are also equal and $\Lambda_1 = \Lambda_2$. Last but not least, we have the property that

$$\Sigma^{1/2} = \exp \left(\frac{1}{2} (\log \Sigma) \right).$$

3.2 An affine invariant distance

Let us consider the following action of the linear group GL_n on the tensor space $\mathcal{Sym}_n^+$:

$$A \star \Sigma = A \Sigma A^T \quad \forall A \in GL_n \quad \text{and} \quad \Sigma \in \mathcal{Sym}_n^+.$$

This group action corresponds for instance to the standard action of the affine group on the covariance matrix Σ_{xx} of a random variables x in $\mathbb{R}^n$: if $y = Ax + t$, then $\bar{y} = A\bar{x} + t$ and $\Sigma_{yy} = \mathbf{E}[(y - \bar{y})(y - \bar{y})^T] = A \Sigma_{xx} A^T$.

This action is naturally extended to tangent vectors in the same way: if $\Gamma(t) = \Sigma + t W + O(t^2)$ is a curve passing at Σ with tangent vector W , then the curve $A \star \Gamma(t) = A \Sigma A^T + t A W A^T + O(t^2)$ passes through $A \star \Sigma$ with tangent vector $A \star W$.

Following [Pennec and Ayache, 1998], any invariant distance on $\mathcal{Sym}_n^+$ verifies $\text{dist}(A \star \Sigma_1, A \star \Sigma_2) = \text{dist}(\Sigma_1, \Sigma_2)$. Choosing $A = \Sigma_1^{-1/2}$, we can reduce this to a pseudo-norm, or distance to the identity:

$$\text{dist}(\Sigma_1, \Sigma_2) = \text{dist} \left(\text{Id}, \Sigma_1^{-\frac{1}{2}} \Sigma_2 \Sigma_1^{-\frac{1}{2}} \right) = N \left(\Sigma_1^{-\frac{1}{2}} \Sigma_2 \Sigma_1^{-\frac{1}{2}} \right).$$

Moreover, as the invariance has to hold for any transformation, N should be invariant under the action of the isotropy group $\mathcal{H}(\text{Id}) = O_n = \{U \in GL_n / UU^T = \text{Id}\}$:

$$\forall U \in O_n, \quad N(U \Sigma U^T) = N(\Sigma).$$

Using the spectral decomposition $\Sigma = U D^2 U^T$, it is easy to see that $N(\Sigma)$ has to be a symmetric function of the eigenvalues. Moreover, the symmetry of the distance $\text{dist}(\Sigma, \text{Id}) = \text{dist}(\text{Id}, \Sigma)$ imposes that $N(\Sigma) = N(\Sigma^{(-1)})$. Thus, a good candidate is the sum of the squared logarithms of the eigenvalues:

$$N(\Sigma)^2 = \|\log(\Sigma)\|^2 = \sum_{i=1}^n (\log(\sigma_i))^2. \quad (2)$$

This “norm” verifies by construction the symmetry and positiveness. $N(\Sigma) = 0$ implies that $\sigma_i = 1$ (and conversely), so that the separation axiom is verified. However, we do not know any simple proof of the triangle inequality, which should read $N(\Sigma_1) + N(\Sigma_2) \geq N(\Sigma_1^{-1/2} \Sigma_2 \Sigma_1^{-1/2})$, even if we can verify it experimentally (see e.g. [Förstner and Moonen, 1999]).

3.3 An invariant Riemannian metric

Another way to determine the invariant distance is through the Riemannian metric. Let us take the most simple scalar product on the tangent space at the identity matrix: if W_1 and W_2 are tangent vectors (i.e. symmetric matrices, not necessarily definite nor positive), we define the scalar product to be the standard matrix scalar product $\langle W_1 | W_2 \rangle = \text{Tr}(W_1^T W_2)$. This scalar product is obviously invariant by the isotropy group O_n . Now, if W_1 and W_2 are two tangent vectors at Σ , we require their scalar product to be invariant by the action of any transformation: $\langle W_1 | W_2 \rangle_\Sigma = \langle A \star W_1 | A \star W_2 \rangle_{A \star \Sigma}$. This should be true in particular for $A = \Sigma^{-1/2}$, which allows us to define the scalar product at any Σ from the scalar product at the identity:

$$\langle W_1 | W_2 \rangle_\Sigma = \left\langle \Sigma^{-\frac{1}{2}} W_1 \Sigma^{-\frac{1}{2}} \mid \Sigma^{-\frac{1}{2}} W_2 \Sigma^{-\frac{1}{2}} \right\rangle_{\text{Id}} = \text{Tr} \left(\Sigma^{-\frac{1}{2}} W_1 \Sigma^{-1} W_2 \Sigma^{-\frac{1}{2}} \right).$$

One can easily verify that this definition is left unchanged if we use any other transformation $A = U \Sigma^{-1/2}$ (where U is a free orthonormal matrix) that transports Σ to the identity: $A \star \Sigma = A \Sigma A^T = U U^T = \text{Id}$.

To find the geodesic without going through the computation of Christoffel symbols, we may rely on a result from differential geometry [Gamkrelidze, 1991, Helgason, 1978, Kobayashi and Nomizu, 1963] which says that the geodesics for the invariant metrics on affine symmetric spaces are generated by the action of the one-parameter subgroups of the acting Lie group¹. Since the one-parameter subgroups of the linear group are given by the matrix exponential $\exp(t A)$, geodesics on our tensor manifold going through Σ with tangent vector W should have the following form:

$$\Gamma_{(\Sigma, W)}(t) = \exp(t A) \Sigma \exp(t A)^T \quad \text{with} \quad W = A \Sigma + \Sigma A^T. \quad (3)$$

¹To be mathematically correct, we should consider the quotient space $\text{Sym}_n^+ = GL_n^+ / SO_n$ instead of $\text{Sym}_n^+ = GL_n / O_n$ so that all spaces are simply connected.

For our purpose, we need to relate explicitly the geodesic to the tangent vector in order to define the exponential chart. Since Σ is a symmetric matrix, there is hopefully an explicit solution to the Sylvester equation $W = A\Sigma + \Sigma A^T$. We get $A = \frac{1}{2} (W\Sigma^{(-1)} + \Sigma^{1/2} Z \Sigma^{-1/2})$, where Z is a free skew-symmetric matrix. However, introducing this solution into the equation of geodesics (Eq. 3) does not lead to a very tractable expression. Let us look at an alternative solution.

Since our metric (and thus the geodesics) is invariant under the action of the group, we can focus on the geodesics going through the origin (the identity). In that case, a symmetric solution of the Sylvester equation is $A = \frac{1}{2}W$, which gives the following equation for the geodesic going through the identity with tangent vector W :

$$\Gamma_{(\text{Id}, W)}(t) = \exp\left(\frac{t}{2} W\right) \exp\left(\frac{t}{2} W\right)^T = \exp(t W).$$

We may observe that the tangent vector along this curve is the parallel transportation of the initial tangent vector. If $W = U \text{DIAG}(w_i) U^T$,

$$\frac{d\Gamma(t)}{dt} = U \text{DIAG}(w_i \exp(t w_i)) U^T = \Gamma(t)^{\frac{1}{2}} W \Gamma(t)^{\frac{1}{2}} = \Gamma(t)^{\frac{1}{2}} \star W.$$

By definition of our invariant metric, the norm of this vector is constant: $\|\Gamma(t)^{1/2} \star W\|_{\Gamma(t)^{1/2} \star \text{Id}}^2 = \|W\|_{\text{Id}}^2 = \|W\|_2^2$. This was expected since geodesics are parameterized by arc-length. Thus, the length of the curve between time 0 and 1 is

$$\mathcal{L} = \int_0^1 \left\| \frac{d\Gamma(t)}{dt} \right\|_{\Gamma(t)}^2 dt = \|W\|_{\text{Id}}^2.$$

Solving for $\Gamma_{(\text{Id}, W)}(1) = \Sigma$, we obtain the “norm” $N(\Sigma)$ of Eq.(2). Using the invariance of our metric, we easily obtain the geodesic starting from any other point of the manifold using our group action:

$$\Gamma_{(\Sigma, W)}(t) = \Sigma^{\frac{1}{2}} \star \Gamma_{(\text{Id}, \Sigma^{-1/2} \star W)}(t) = \Sigma^{\frac{1}{2}} \exp\left(t \Sigma^{-\frac{1}{2}} W \Sigma^{-\frac{1}{2}}\right) \Sigma^{\frac{1}{2}}.$$

Coming back to the distance $\text{dist}^2(\Sigma, \text{Id}) = \sum_i (\log \sigma_i)^2$, it is worth noticing that tensors with null eigenvalues are located as far from the identity as tensors with infinite eigenvalues: at the infinity. Thanks to the invariance by the linear group, this property holds for the distance to any (positive definite) tensor of the manifold. Thus, the original cone of positive definite symmetric matrices (a linear manifold with a flat metric but which is incomplete: there is a boundary at a finite distance) has been changed into a regular and complete (but curved) manifold with an infinite development in each of its $n(n+1)/2$ directions.

3.4 Exponential and logarithm maps

As a general property of Riemannian manifolds, geodesics realize a local diffeomorphism from the tangent space at a given point of the manifold to the manifold: $\Gamma_{(\Sigma, W)}(1) = \exp_{\Sigma}(W)$ associates to each tangent vector $W \in T_{\Sigma} \text{Sym}_n^+$ a point of the manifold. This mapping is

called the exponential map, because it corresponds to the usual exponential in some matrix groups. This is exactly our case for the exponential map around the identity:

$$\exp_{\text{Id}}(UDU^T) = \exp(UDU^T) = U \text{DIAG}(\exp(d_i)) U^T.$$

However, the Riemannian exponential map associated to our invariant metric has a more complex expression at other tensors:

$$\exp_{\Sigma}(W) = \Sigma^{\frac{1}{2}} \exp\left(\Sigma^{-\frac{1}{2}} W \Sigma^{-\frac{1}{2}}\right) \Sigma^{\frac{1}{2}}.$$

In our case, this diffeomorphism is global, and we can uniquely define the inverse mapping everywhere:

$$\log_{\Sigma}(\Lambda) = \Sigma^{\frac{1}{2}} \log\left(\Sigma^{-\frac{1}{2}} \Lambda \Sigma^{-\frac{1}{2}}\right) \Sigma^{\frac{1}{2}}.$$

Thus, $\exp_{\Sigma}$ gives us a collection of one-to-one and complete maps of the manifold, centered at any point Σ . As explained in Section 2.1, these charts can be viewed as the development of the manifold onto the tangent space along the geodesics. Moreover, as the manifold has a non-positive curvature [Skovgaard, 1984], there is no cut-locus and the statistical properties detailed in [Pennec, 2004] hold in their most general form. For instance, we have the existence and uniqueness of the mean of any distribution with a compact support [Kendall, 1990].

3.5 Induced and orthonormal coordinate systems

One has to be careful because the coordinate system of all these charts is not orthonormal. Indeed, the coordinate system of each chart is induced by the standard coordinate system (here the matrix coefficients), so that the vector $\overrightarrow{\Sigma\Lambda}$ corresponds to the standard derivative in the vector space of matrices: we have $\Lambda = \Sigma + \overrightarrow{\Sigma\Lambda} + O(\|\overrightarrow{\Sigma\Lambda}\|^2)$. Even if this basis is orthonormal at some points of the manifold (such as at the identity for our tensors), it has to be corrected for the Riemannian metric at other places due to the manifold curvature.

From the expression of the metric, one can observe that

$$\|\overrightarrow{\Sigma\Lambda}\|_{\Sigma}^2 = \|\log_{\Sigma}(\Lambda)\|_{\Sigma}^2 = \|\Sigma^{-\frac{1}{2}} \log_{\Sigma}(\Lambda) \Sigma^{-\frac{1}{2}}\|_{\text{Id}}^2 = \|\log(\Sigma^{-\frac{1}{2}} \star \Lambda)\|_2^2.$$

This shows that $\overrightarrow{\Sigma\Lambda}_{\perp} = \log(\Sigma^{-\frac{1}{2}} \star \Lambda) \in T_{\Sigma} \text{Sym}_n^+$ is the expression of the vector $\overrightarrow{\Sigma\Lambda}$ in an orthonormal basis. In our case, the transformation $\Sigma^{1/2} \in GL_n$ is moreover uniquely defined (as a positive square root) and is a smooth function of Σ over the complete tensor manifold. Thus, $\overrightarrow{\Sigma\Lambda}_{\perp}$ realizes an atlas of orthonormal exponential charts which is globally smooth with respect to the reference point² Σ . This group action approach was chosen in earlier works [Pennec, 1996, Pennec and Thirion, 1997, Pennec and Ayache, 1998] with what we called the placement function.

For some statistical operations, we need to use a minimal representation (e.g. 6 parameters for 3×3 tensors) in a (locally) orthonormal basis. This can be realized through the classical “Vec” operator that maps the element $a_{i,j}$ of a $n \times n$ matrix A to the $(i n + j)$ th

²On most homogeneous manifolds, this can only be realized locally. For instance, on the sphere, there is a singularity at the antipodal point of the chosen origin for any otherwise smooth placement function.

element $\text{Vec}(A)_{i\ n+j}$ of a $n \times n$ dimensional vector $\text{Vec}(A)$. Since we are working with symmetric matrices, we have only $n(n+1)/2$ independent coefficients, say the upper triangular part. However, the off-diagonal coefficients are counted twice in the L_2 norm at the identity: $\|W\|_2^2 = \sum_{i=1}^n w_{i,i}^2 + 2 \sum_{i < j \leq n} w_{i,j}^2$. Thus, to express our minimal representation in an orthonormal basis, we need to multiply the off diagonal terms by $\sqrt{2}$:

$$\text{Vec}_{\text{Id}}(W) = \left(w_{1,1}, \sqrt{2} w_{1,2}, w_{2,2}, \sqrt{2} w_{1,3}, \sqrt{2} w_{2,3}, w_{3,3}, \dots, \sqrt{2} w_{1,n}, \dots, \sqrt{2} w_{(n-1),n}, w_{n,n} \right)^T.$$

Now, for a vector $\vec{\Sigma\Lambda} \in T_{\Sigma} \text{Sym}_n^+$, we define its minimal representation in the orthonormal coordinate system as:

$$\text{Vec}_{\Sigma}(\vec{\Sigma\Lambda}) = \text{Vec}_{\text{Id}}(\vec{\Sigma\Lambda}_{\perp}) = \text{Vec}_{\text{Id}}\left(\Sigma^{-\frac{1}{2}} \vec{\Sigma\Lambda} \Sigma^{-\frac{1}{2}}\right) = \text{Vec}_{\text{Id}}\left(\log(\Sigma^{-\frac{1}{2}} \star \Lambda)\right).$$

The mapping Vec_{Σ} realizes an explicit isomorphism between $T_{\Sigma} \text{Sym}_n^+$ and $\mathbb{R}^{n(n+1)/2}$ with the canonical metric.

3.6 Gradient descent and PDEs: an intrinsic geodesic marching scheme

Let $f(\Sigma)$ be an objective function to minimize, Σ_t the current estimation of Σ , and $W_t = \partial_{\Sigma} f = [\partial f / \partial \sigma_{ij}]$ its matrix derivative at that point, which is of course symmetric. The principle of a first order gradient descent is to go toward the steepest descent, in the direction opposite to the gradient for a short time-step ε , and iterate the process. However, the standard operator $\Sigma_{t+1} = \Sigma_t - \varepsilon W_t$ is only valid for very short time-steps in the flat Euclidean matrix space, and we could easily go out of the cone of positive definite tensors. A much more interesting numerical operator is given by following the geodesic backward starting at Σ with tangent vector W_t during a time ε . This intrinsic gradient descent ensures that we cannot leave the manifold. It can easily be expressed using the exponential map:

$$\Sigma_{t+1} = \Gamma_{(\Sigma_t, W_t)}(-\varepsilon) = \exp_{\Sigma_t}(-\varepsilon W_t) = \Sigma_t^{\frac{1}{2}} \exp(-\varepsilon \Sigma_t^{-\frac{1}{2}} W_t \Sigma_t^{-\frac{1}{2}}) \Sigma_t^{\frac{1}{2}}.$$

This intrinsic scheme is trivially generalized to partial differential evolution equations (PDEs) on tensor fields such as $\partial_t \Sigma(x, t) = -W(x, t)$: we obtain $\Sigma(x, t+dt) = \exp_{\Sigma(x, t)}(-dt W(x, t))$.

3.7 Example with the mean value

Let $\Sigma_1 \dots \Sigma_N$ be a set of measures of the same Tensor. The Karcher or Fréchet mean is the set of tensors minimizing the sum of squared distances: $C(\Sigma) = \sum_{i=1}^N \text{dist}^2(\Sigma, \Sigma_i)$. In the case of tensors, the manifold has a non-positive curvature [Skovgaard, 1984], so that there is one and only one mean value $\bar{\Sigma}$ [Kendall, 1990]. Moreover, a necessary and sufficient condition for an optimum is a null gradient of the criterion. Thus, the intrinsic Newton gradient descent algorithm gives the following mean value at estimation step $t+1$:

$$\bar{\Sigma}_{t+1} = \exp_{\bar{\Sigma}_t} \left(\frac{1}{N} \sum_{i=1}^N \log_{\bar{\Sigma}_t}(\Sigma_i) \right) = \bar{\Sigma}_t^{\frac{1}{2}} \exp \left(\frac{1}{N} \sum_{i=1}^N \log \left(\bar{\Sigma}_t^{-\frac{1}{2}} \Sigma_i \bar{\Sigma}_t^{-\frac{1}{2}} \right) \right) \bar{\Sigma}_t^{\frac{1}{2}}. \quad (4)$$

Note that we cannot easily simplify this expression further as in general the data Σ_i and the mean value $\bar{\Sigma}_t$ cannot be diagonalized in a common basis. However, this gradient descent algorithm usually converges very fast (about 10 iterations, see Fig. 2).

3.8 Simple statistical operations on tensors

As described in [Pennec, 2004], we may generalize most of the usual statistical methods by using the exponential chart at the mean point. For instance, the empirical covariance matrix of a set of N tensors Σ_i of mean $\bar{\Sigma}$ is defined using the tensor product: $\frac{1}{N-1} \sum_{i=1}^n \bar{\Sigma} \bar{\Sigma}_i \otimes \bar{\Sigma} \bar{\Sigma}_i$. Using our *Vec* mapping, we may come back to more usual matrix notations and write its expression in a minimal representation with an orthonormal coordinate system:

$$\text{Cov} = \frac{1}{N-1} \sum_{i=1}^N \text{Vec}_{\bar{\Sigma}} \left(\bar{\Sigma} \bar{\Sigma}_i \right) \text{Vec}_{\bar{\Sigma}} \left(\bar{\Sigma} \bar{\Sigma}_i \right)^T.$$

One may also define the Mahalanobis distance

$$\mu_{(\bar{\Sigma}, \text{Cov})}^2(\Sigma) = \text{Vec}_{\bar{\Sigma}} \left(\bar{\Sigma} \Sigma \right)^T \text{Cov}^{(-1)} \text{Vec}_{\bar{\Sigma}} \left(\bar{\Sigma} \Sigma \right).$$

Looking for the probability density function that minimizes the information with a constrained mean and covariance, we obtain a generalization of the Gaussian distribution of the form:

$$N_{\bar{\Sigma}, \Gamma}(\Sigma) = k \exp \left(-\frac{1}{2} \mu_{\bar{\Sigma}, \Gamma}^2(\Sigma) \right).$$

The main difference with a Euclidean space is that we have a curvature to take into account: the invariant measure induced on the manifold by our metric is linked to the usual matrix measure by $d\mathcal{M}(\Sigma) = d\Sigma / \det(\Sigma)$. Likewise, the curvature slightly modifies the usual relation between the covariance matrix, the concentration matrix Γ and the normalization parameter k of the Gaussian distribution [Pennec, 2004]. These differences have an impact on the calculations using continuous probability density functions. However, from a practical point of view, we only deal with a discrete sample set of measurements, so that the measure-induced corrections are hidden. For instance, we can generate a random (generalized) Gaussian tensor using the following procedure: we sample $n(n+1)/2$ independent and normalized real Gaussian samples, multiply the corresponding vector by the square root of the desired covariance matrix (expressed in our *Vec* coordinate system), and come back to the tensor manifold using the inverse *Vec* mapping. Using this procedure, we can easily generate noisy measurements of known tensors (see e.g. Fig. 7).

To check the implementation of our charts and geodesic marching algorithms, we verified experimentally the central limit theorem. This theorem states that the empirical mean of N independently and identically distributed (IID) random variables with a variance γ^2 asymptotically follows a Gaussian law of variance γ^2/N , centered at the exact mean value. The principle of our experiments is now as follows. We randomly generated N random Gaussian tensors around a random tensor $\bar{\Sigma}$ with a variance of $\gamma^2 = 1$. We computed the mean $\hat{\Sigma}$ using the algorithm of Eq. 4. The convergence is clearly very fast (Fig. 2, left).

Now, if the error between the computed and the exact mean really follows a Gaussian law of variance γ^2/N , then the normalized Mahalanobis distance $\mu^2 = N \text{ dist}(\bar{\Sigma}, \hat{\Sigma})^2/\gamma^2$ should follow a χ_6^2 distribution. However, this simple experiment only gives us one measurement. Thus, to verify the distribution, we repeated this experiment with N varying from 10 to 1000. Figure 2 presents the histogram of the normalized Mahalanobis distances we obtain. The empirical distribution follows quite well the theoretical χ_6^2 distribution, as expected, with an empirical mean of 5.85 and a variance of 12.17 (expected values are 6 and 12). Moreover, a Kolmogorov-Smirnov test confirms that the distance between the empirical and theoretical cumulative pdf is not significant (p-value of 0.19).

Figure 2: Mean of random Gaussian tensors. **Left:** Typical evolution of the distance between successive iterations of the mean computation. The convergence is clearly very fast. **Right:** Histogram of the renormalized Mahalanobis distance $\mu^2 = N \text{ dist}(\bar{\Sigma}, \hat{\Sigma})^2/\gamma^2$ between the computed and the exact mean tensors. The curve is the pdf of the χ_6^2 distribution.

4 Tensor Interpolation

One of the important operations in geometric data processing is to interpolate values between known measurements. In 3D image processing, (tri-) linear interpolation is often used thanks to its very low computational load and comparatively much better results than nearest neighbor interpolation. Other popular methods include the cubic and, more generally, spline interpolations [Thévenaz et al., 2000, Meijering, 2002].

The standard way to define an interpolation on a regular lattice of dimension d is to consider that the interpolated function $f(x)$ is a linear combination of samples f_k at integer (lattice) coordinates $k \in \mathbb{Z}^d$: $f(x) = \sum_k w(x-k) f_k$. To realize an interpolation, the “sample weight” function w has to vanish at all integer coordinates except 0 where it has to be one. A typical example where the convolution kernel has an infinite support is the sinus cardinal interpolation. With the nearest-neighbor, linear (or tri-linear in 3D), and higher order spline interpolations, the kernel is piecewise polynomial, and limited to a few neighboring points in the lattice.

When it comes to an irregular sampling (i.e. a set of measurements f_k at positions x_k), interpolation may still be defined using a weighted mean: $f(x) = \sum_{k=1}^N w_k(x) f_k$. To ensure

that this is an interpolating function, one has to require that $w_i(x_j) = \delta_{ij}$ (where δ_{ij} is the Kronecker symbol). Moreover, the coordinates are usually normalized so that $\sum_{k=1}^N w_k(x) = 1$ for all position x within the domain of interest. Typical examples in triangulations or tetrahedrizations are barycentric and natural neighbor coordinates [Sibson, 1981].

4.1 Interpolation through weighted mean

To generalize interpolation methods defined using weighted means to our tensor manifold, let us assume that the sample weights $w_k(x)$ are defined as above in $\mathbb{R}^d$. Thanks to their normalization, the value $f(x)$ interpolated from vectors f_k verifies $\sum_{i=1}^N w_i(x) (f_i - f(x)) = 0$. Thus, similarly to the Fréchet mean, we can define the interpolated value $\Sigma(x)$ on our tensor manifold as the tensor that minimizes the weighted sum of squared distances to the measurements Σ_i : $C(\Sigma(x)) = \sum_{i=1}^N w_i(x) \text{dist}^2(\Sigma_i, \Sigma(x))$. Of course, we loose in general the existence and uniqueness properties. However, for positive weights, the existence and uniqueness theorems for the Karcher mean can be adapted. In practice, this means that we have a unique tensor that verifies $\sum_{i=1}^N w_i(x) \overrightarrow{\Sigma(x)\Sigma_i} = 0$. To reach this solution, it is easy to adapt the Gauss-Newton scheme proposed for the Karcher mean. The algorithm becomes:

$$\begin{aligned} \Sigma_{t+1}(x) &= \exp_{\Sigma_t(x)} \left(\sum_{i=1}^N w_i(x) \log_{\Sigma_t(x)}(\Sigma_i) \right) \\ &= \Sigma_t^{\frac{1}{2}}(x) \exp \left(\sum_{i=1}^N w_i(x) \log \left(\Sigma_t^{-\frac{1}{2}}(x) \Sigma_i \Sigma_t^{-\frac{1}{2}}(x) \right) \right) \Sigma_t^{\frac{1}{2}}(x). \end{aligned} \quad (5)$$

Once again, this expression cannot be easily simplified, but the convergence is very fast (usually less than 10 iterations as for the mean).

4.2 Example of the linear interpolation

The linear interpolation is simple as this is a walk along the geodesic joining the two tensors. For instance, the interpolation in the standard Euclidean matrix space would give $\Sigma'(t) = (1-t)\Sigma_1 + t\Sigma_2$. In our Riemannian space, we have the closed-form expression: $\Sigma(t) = \exp_{\Sigma_1}(t \log_{\Sigma_1}(\Sigma_2)) = \exp_{\Sigma_2}((1-t) \log_{\Sigma_2}(\Sigma_1))$ for $t \in [0; 1]$. We displayed in Fig. 3 the flat and the Riemannian interpolations between 2D tensors of eigenvalues (5,1) horizontally and (1,50) at 45 degrees, along with the evolution of the eigenvalues, their mean (i.e. trace of the matrix) and product (i.e. determinant of the matrix or volume of the ellipsoid).

With the standard matrix coefficient interpolation, the evolution of the trace is perfectly linear (which was expected since this is a linear function of the coefficients), and the principal eigenvalue regularly grows almost linearly, while the smallest eigenvalue slightly grows toward a local maxima before lowering. What is much more annoying is that the determinant (i.e. the volume) does not grow regularly in between the two tensors, but goes through a maximum. If we interpret our tensors as covariance matrices of Gaussian distributions, this means that the probability of a random point to be accepted as a realization of our distribution is larger in between than at the measurement points themselves! On the contrary,

Figure 3: **Top:** Linear interpolation between 2D tensors of eigenvalues (5,1) horizontally and (1,50) at 45 degrees. **Left:** interpolation in the standard matrix space (interpolation of the coefficients). **Right:** geodesic interpolation in our Riemannian space. **Bottom:** evolution of the eigenvalues, their mean (i.e. trace of the matrix) and product (i.e. determinant of the matrix or volume of the ellipsoid).

one can clearly see a regular evolution of the eigenvalues and of their product with the interpolation in our Riemannian space. Moreover, there is a much smoother rotation of the eigenvectors than with the standard interpolation.

4.3 Tri-linear interpolation

The bi- and tri-linear interpolation of tensors on a regular grid in 2D or 3D are almost as simple, except that we do not have any longer an explicit solution using geodesics since there are more than two reference points. After computing the (bi-) tri-linear weights with respect to the neighboring sites of the point we want to evaluate, we now have to go through the iterative optimization of the weighted mean (Eq. 5) to compute the interpolated tensor. We display an example in Figure 4. One can see that the volume of the tensors is much more important with the classical than with the Riemannian interpolation. We also get a much smoother interpolation of the principal directions with our method.

4.4 Interpolation of non regular measurements

When tensors are not measured on a regular grid but “randomly” localized in space, defining neighbors becomes an issue. One solution, proposed by [Sibson, 1981] and later used for

Figure 4: **Top left:** Bi-linear interpolation between the four 2D tensors at the corners in the standard matrix space (interpolation of the coefficients). **Top right:** Equivalent bi-linear interpolation in our Riemannian space. **Bottom left:** A slice of the tri-linear interpolation between 3D tensors in the standard matrix space (interpolation of the coefficients). **Bottom right:** Equivalent tri-linear interpolation in our Riemannian space.

surfaces by [Cazals and Boissonnat, 2001], is the natural neighbor interpolation. For any point x , its natural neighbors are the points of $\{x_i\}$ whose Voronoi cells are chopped off upon insertion of x into the Voronoi diagram. The weight w_i of each natural neighbor x_i is the proportion of the new cell that is taken away by x to x_i in the new Voronoi diagram. One important restriction of these interesting coordinates is that they are limited to the convex hull of the point set (otherwise the volume or surface of the cell is infinite).

Another idea is to rely on radial-basis functions to define the relative influence of each measurement point. For instance, a Gaussian influence would give a weight $w_i(x) = G_\sigma(x - x_i)$ to the measurement Σ_i located at x_i . Since weights need to be renormalized in our setup, this would lead to the following evolution equation:

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)} \left(\frac{\sum_{i=1}^N G_\sigma(x - x_i) \overrightarrow{\Sigma_t(x) \Sigma_i}}{\sum_{i=1}^N G_\sigma(x - x_i)} \right). \quad (6)$$

The initialization could be the (normalized) Gaussian mean in the matrix space. An example of the result of this evolution scheme is provided on top of Figure 10. However, this algorithm does not lead to an interpolation, but rather to an approximation, since the weights are not zero at other measurement points. Moreover, we have little control on the quality of this approximation. It is only at the limit where σ goes to zero that we end-up with a (non-continuous) closest point interpolation.

We will describe in Section 6.3 a last alternative that performs the interpolation and extrapolation of sparsely distributed tensor measurements using diffusion.

5 Filtering tensor fields

Let us now consider that we have a tensor field, for instance like in Diffusion Tensor Imaging (DTI) [Le Bihan et al., 2001], where the tensor is a first order approximation of the anisotropic diffusion of the water molecules at each point of the imaged tissues. In the brain, the diffusion is much favored in the direction of oriented structures (fibers of axons). One of the goal of DTI is to retrieve the main tracts along these fibers. However, the tensor field obtained from the images is noisy and needs to be regularized before being further analyzed. A naive but simple and often efficient regularization on signal or images is the convolution by a Gaussian. The generalization to tensor fields is quite straightforward using once again weighted means (Section 5.1 below). An alternative is to consider a regularization using diffusion. This will be the subject of Sections 5.3 and 5.4.

5.1 Gaussian Filtering

In the continuous setting, the convolution of a vector field $F_0(x)$ by a Gaussian is:

$$F(x) = \int_y G_\sigma(y - x) F_0(y) dy.$$

In the discrete setting, coefficients are renormalized since the neighborhood $\mathcal{V}$ is usually limited to points within one to three times the standard deviation:

$$F(x) = \frac{\sum_{u \in \mathcal{V}(x)} G_\sigma(u) F_0(x + u)}{\sum_{u \in \mathcal{V}(x)} G_\sigma(u)} = \arg \min_F \sum_{u \in \mathcal{V}(x)} G_\sigma(u) \|F_0(x + u) - F\|^2.$$

Like previously, this weighted mean can be solved on our manifold using our intrinsic gradient descent scheme. Starting from the measured tensor field $\Sigma_0(x)$, the evolution equation is

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)} \left(\frac{\sum_{u \in \mathcal{V}} G_\sigma(u) \overrightarrow{\Sigma_t(x) \Sigma_t(x + u)}}{\sum_{u \in \mathcal{V}} G_\sigma(u)} \right).$$

We illustrate in Fig. 5 the comparative Gaussian filtering of a slice of a DT MR image using the flat metric on the coefficient (since weights are positive, a weighted sum of positive definite matrices is still positive definite) and our invariant Riemannian metric. One can see a more important blurring of the corpus callosum fiber tracts using the flat metric.

However, the integration of this filtering scheme into a complete fiber tracking system would be necessary to fully evaluate the pros and cons of each metric.

Figure 5: Regularization of a DTI slice around the corpus callosum by isotropic Gaussian filtering. **Left:** raw estimation of the tensors. The color codes for the direction of the principal eigenvector (red: left/right, green: anterior/posterior, blue: top/bottom). **Middle:** Gaussian filtering of the coefficients (5x5 window, $\sigma = 2.0$). **Right:** equivalent filtering (same parameters) using the Riemannian metric.

5.2 Spatial gradient of Tensor fields

On a n -dimensional vector field $F(x) = (f_1(x_1, \dots, x_d), \dots, f_n(x_1, \dots, x_d))^T$ over $\mathbb{R}^d$, one may express the spatial gradient in an orthonormal basis as:

$$\nabla F^T = \left(\frac{\partial F}{\partial x} \right) = [\partial_1 F, \dots, \partial_d F] = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_1}{\partial x_d} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1} & \cdots & \frac{\partial f_n}{\partial x_d} \end{bmatrix}.$$

The linearity of the derivatives implies that we could use directional derivatives in more than the d orthogonal directions. This is especially well adapted to stabilize the discrete computations: the finite difference estimation of the directional derivative is $\partial_u F(x) = F(x+u) - F(x)$. By definition, the spatial gradient is related to the directional derivatives through $\nabla F^T u = \partial_u F(x)$. Thus, we may compute ∇F as the matrix that best approximates (in the least-square sense) the directional derivatives in the neighborhood $\mathcal{V}$ (e.g. 6, 18 or 26 connectivity in 3D):

$$\begin{aligned} \nabla F(x) &= \arg \min_G \sum_{u \in \mathcal{V}} \|G^T u - \partial_u F(x)\|^2 = \left(\sum_{u \in \mathcal{V}} u u^T \right)^{(-1)} \left(\sum_{u \in \mathcal{V}} u \partial_u F(x)^T \right) \\ &\simeq \left(\sum_{u \in \mathcal{V}} u u^T \right)^{(-1)} \left(\sum_{u \in \mathcal{V}} u (F(x+u) - F(x))^T \right). \end{aligned}$$

We experimentally found in other applications (e.g. to compute the Jacobian of a deformation field in non-rigid registration [Rey et al., 2002, p. 169]) that this gradient approximation scheme was more stable and much faster than computing all derivatives using convolutions, for instance by the derivative of the Gaussian.

To quantify the local amount of variability independently of the space direction, one usually takes the norm of the gradient: $\|\nabla F(x)\|^2 = \sum_{i=1}^d \|\partial_i F(x)\|^2$. Once again, this can be approximated using all directional derivatives in the neighborhood

$$\|\nabla F(x)\|^2 \simeq \frac{d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} \frac{\|F(x+u) - F(x)\|^2}{\|u\|^2}. \quad (7)$$

Notice that this approximation is consistent with the previous one only if the directions u are normalized to unity.

For a manifold valued field $\Sigma(x)$ define on $\mathbb{R}^d$, we can proceed similarly, except that the directional derivatives $\partial_i \Sigma(x)$ are now tangent vectors of $T_{\Sigma(x)}\mathcal{M}$. They can be approximated just like above using finite “differences” in our exponential chart:

$$\partial_u \Sigma(x) \simeq \overrightarrow{\Sigma(x) \Sigma(x+u)} = \Sigma(x)^{\frac{1}{2}} \log \left(\Sigma(x)^{-\frac{1}{2}} \Sigma(x+u) \Sigma(x)^{-\frac{1}{2}} \right) \Sigma(x)^{\frac{1}{2}}. \quad (8)$$

As observed in Section 3.5, we must be careful that this directional derivative is expressed in the standard matrix coordinate system (coefficients). Thus, the basis is not orthonormal: to quantify the local amount of variation, we have to take the metric at the point $\Sigma(x)$ into account, so that:

$$\|\nabla \Sigma(x)\|_{\Sigma(x)}^2 = \sum_{i=1}^d \|\partial_i \Sigma(x)\|_{\Sigma(x)}^2 \simeq \frac{d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} \frac{\left\| \log \left(\Sigma(x)^{-\frac{1}{2}} \Sigma(x+u) \Sigma(x)^{-\frac{1}{2}} \right) \right\|_2^2}{\|u\|^2}. \quad (9)$$

5.3 Filtering using PDEs

Regularizing a scalar, vector or tensor field F aims at reducing the amount of its spatial variations. The first order measure of such variations is the spatial gradient ∇F that we dealt with in the previous section. To obtain a regularity criterion over the domain Ω , we just have to integrate: $\text{Reg}(F) = \int_{\Omega} \|\nabla F(x)\|^2 dx$. Starting from an initial field $F_0(x)$, the goal is to find at each step a field $F_t(x)$ that minimizes the regularity criterion by gradient descent in the space of (sufficiently smooth and square integrable) functions.

To compute the first order variation, we write a Taylor expansion for an incremental step in the direction of the field H . Notice that $H(x)$ is a tangent vector at $F(x)$:

$$\text{Reg}(F + \varepsilon H) = \text{Reg}(F) + 2 \varepsilon \int_{\Omega} \langle \nabla F(x) \mid \nabla H(x) \rangle dx + O(\varepsilon^2).$$

We get the directional derivative: $\partial_H \text{Reg}(F) = 2 \int_{\Omega} \langle \nabla F(x) \mid \nabla H(x) \rangle dx$. To compute the steepest descent, we now have to find the gradient $\nabla \text{Reg}(F)$ such that for all variation H , we have $\partial_H \text{Reg}(F) = \int_{\Omega} \langle \nabla \text{Reg}(F)(x) \mid H(x) \rangle_{F(x)} dx$. Notice that $\nabla \text{Reg}(F)(x)$ and $H(x)$ are elements of the tangent space at $F(x)$, so that the scalar product should be taken at $F(x)$ for a tensor field.

Figure 6: Norm of the gradient of the tensor field. **Left:** computed on the coefficients with Eq. 7 (with the flat metric). **Middle:** we computed the directional derivatives with the exponential map (Eq. 8), but the norm is taken without correcting for the metric. As this should be very close to the flat gradient norm, we only display the difference image. The main differences are located on very sharp boundaries, where the curvature of our metric has the most important impact. However, the relative differences remains small (less than 10%), which shows the stability of both the gradient and the log/exp computation schemes. **Right:** Riemannian norm of the Riemannian gradient (Eq. 9). One can see much more detailed structures within the brain, which will now be preserved during an anisotropic regularization step.

The case of a scalar field Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a scalar field. Our regularization criterion is $Reg(f) = \int_{\Omega} \|\nabla f(x)\|^2 dx$. Let us introduce the divergence $\text{div}(\cdot) = \langle \nabla | \cdot \rangle$ and the Laplacian operator $\Delta f = \text{div}(\nabla f)$. The divergence is usually written $\nabla^T = (\partial_1, \dots, \partial_d)$, so that in an orthonormal coordinate system we have $\Delta f = \langle \nabla | \nabla f \rangle = \sum_{i=1}^d \partial_i^2 f$. Let now $G(x)$ be a vector field. Typically, we will use $G(x) = \nabla f(x)$. Using the standard differentiation rules, we have:

$$\text{div}(h G) = \langle \nabla | h G \rangle = h \text{div}(G) + \langle \nabla h | G \rangle.$$

Now, thanks to the Green's formula (see e.g. [Gallot et al., 1993]), we know that the flux going out of the boundaries of a (sufficiently smooth) region Ω is equal to the integral of the divergence inside this region. If we denote by n the normal pointing outward at a boundary point, we have:

$$\int_{\partial\Omega} \langle h G | n \rangle dn = \int_{\Omega} \text{div}(h G) = \int_{\Omega} h \text{div}(G) + \int_{\Omega} \langle \nabla h | G \rangle.$$

This result can also be interpreted as an integration by part in $\mathbb{R}^d$. Assuming homogeneous Neumann boundary conditions (gradient orthogonal to the normal on $\partial\Omega$: $\langle G | n \rangle = 0$), the flow across the boundary vanishes, and we are left with: $\int_{\Omega} \langle G | \nabla h \rangle = - \int_{\Omega} h \text{div}(G)$. Thus, coming back to our original problem, we have:

$$\partial_h Reg(f)(x) = 2 \int_{\Omega} \langle \nabla f(x) | \nabla h(x) \rangle dx = -2 \int_{\Omega} h(x) \Delta f(x) dx.$$

Since this last formula is no more than the scalar product on the space $L_2(\Omega, \mathbb{R})$ of square integrable functions, we end-up with the classical Euler-Lagrange equation: $\nabla Reg(f) =$

$-2\Delta f(x)$. The evolution equation used to filter the data is thus

$$f_{t+1}(x) = f_t(x) - \varepsilon \nabla \text{Reg}(f)(x) = f_t(x) + 2\varepsilon \Delta f_t(x).$$

The vector case Let us decompose our vector field $F(x)$ into its n scalar components $f_i(x)$. Likewise, we can decompose the $d \times n$ gradient ∇F into the gradient of the n scalar components $\nabla f_i(x)$ (columns). Thus, choosing an orthonormal coordinate system on the space $\mathbb{R}^n$, our regularization criterion is decoupled into n independent scalar regularization problems:

$$\text{Reg}(F)(x) = \sum_{i=1}^n \int_{\Omega} \|\nabla f_i(x)\|^2 dx = \sum_{i=1}^n \text{Reg}(f_i).$$

Thus, each component f_i has to be independently regularized with the Euler-Lagrange equation: $\nabla \text{Reg}(f_i) = -2\Delta f_i$. With the convention that the Laplacian is applied component-wise (so that we still have $\Delta F = \text{div}(\nabla F) = \nabla^T \nabla F = (\Delta f_1, \dots, \Delta f_n)^T$), we end-up with the vectorial equation:

$$\nabla \text{Reg}(F) = -2\Delta F \quad \text{for} \quad \text{Reg}(F) = \int_{\Omega} \|\nabla F(x)\|^2 dx.$$

The associated evolution equation is $F_{t+1}(x) = F_t(x) + 2\varepsilon \Delta F_t(x)$.

Tensor fields For a tensor field $\Sigma(x) \in \text{Sym}_n^+$ over $\mathbb{R}^d$, the procedure is more complex as we should use the covariant derivative (the connection) to differentiate vectors fields on our manifold. However, we may avoid the introduction of additional complex mathematical tools by coming back to the basic definitions. We summarize below the main ideas, while the full calculations are worked out in Appendix A.1. Let $(x_1, \dots, x_d)$ be an orthonormal coordinate system of $\mathbb{R}^d$. Our regularization criterion is:

$$\text{Reg}(\Sigma) = \int_{\Omega} \|\nabla \Sigma(x)\|_{\Sigma(x)}^2 dx = \sum_{i=1}^d \int_{\Omega} \|\partial_i \Sigma\|_{\Sigma}^2. \quad (10)$$

The idea is to write this criterion as the trace of sums and products of standard Euclidean matrices and to compute its directional derivative $\partial_W \text{Reg}$ for a perturbation field W . This expression contains of course derivatives $\partial_i W$ that we need to integrate. However, as everything is expressed in the standard Euclidean chart (matrix coefficients), and assuming the proper Neumann boundary conditions, we shall safely use the previous integration by part formula $\int_{\Omega} \text{Tr}((\partial_i W) \Lambda_i) = - \int_{\Omega} \text{Tr}(W (\partial_i \Lambda_i))$. Notice that we are using the matrix coefficients only as a chart and not as a metric. Eventually, we rewrite the obtained expression in terms of our Riemannian metric to obtain the formula defining the gradient of the criterion: $\partial_W \text{Reg} = \int_{\Omega} \langle W | \nabla \text{Reg} \rangle_{\Sigma}$. By identification, we get: $\nabla \text{Reg}(\Sigma) = -2 \Delta \Sigma$, where Δ is the Laplace-Beltrami operator on our manifold:

$$\Delta \Sigma = \sum_{i=1}^d \Delta_i \Sigma \quad \text{with} \quad \Delta_i \Sigma = \partial_i^2 \Sigma - (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma). \quad (11)$$

As we can see, the flat Euclidean second order directional derivatives $\partial_i^2 \Sigma$ are corrected by an additional term due to the curvature of our manifold. To conclude, the gradient descent on the regularization criterion with the intrinsic geodesic marching scheme of Section 3.6 leads to:

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)}(-\varepsilon \nabla \text{Reg}(\Sigma)(x)) = \exp_{\Sigma_t(x)}(2 \varepsilon \Delta \Sigma(x)).$$

For the numerical computation of the Laplacian, we may approximate the first and second order tensor derivative by their Euclidean derivatives. This gives a fourth order approximation of the Laplace-Beltrami operator (see Appendix A.2). However, this numerical scheme is extrinsic since it is based on (Euclidean) differences of tensors. We propose here an intrinsic scheme based on the exponential chart at the current point. We already know from Eq. 8 that $\overrightarrow{\Sigma(x)\Sigma(x+u)}$ is an approximation of the first order directional derivative $\partial_u \Sigma(x)$. We show in Appendix A.2 that $\overrightarrow{\Sigma(x)\Sigma(x+u)} + \overrightarrow{\Sigma(x)\Sigma(x-u)}$ is a forth order approximation of the Laplace Beltrami operator in the direction u :

$$\Delta_u \Sigma = \partial_u^2 \Sigma - 2 (\partial_u \Sigma) \Sigma^{(-1)} (\partial_u \Sigma) = \overrightarrow{\Sigma(x)\Sigma(x+u)} + \overrightarrow{\Sigma(x)\Sigma(x-u)} + O(\|u\|^4). \quad (12)$$

To compute the complete manifold Laplacian of Eq. 11, we just have to compute the above numerical approximations of the tensor field derivatives along d orthonormal basis vectors x_i . However, like for the computation of the gradient, we may improve the stability of the numerical scheme by averaging the derivatives in all possible directions in the neighborhood $\mathcal{V}$. Assuming a symmetric and isotropic neighborhood, we finally obtain:

$$\Delta \Sigma(x) = \frac{d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} \frac{\Delta_u^2 \Sigma(x)}{\|u\|^2} \simeq \frac{2d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} \frac{\overrightarrow{\Sigma(x)\Sigma(x+u)}}{\|u\|^2}. \quad (13)$$

5.4 Anisotropic filtering

In practice, we would like to filter within the homogeneous regions, but not across their boundaries. The basic idea is to penalize the smoothing in the directions where the derivative is important [Perona and Malik, 1990, Gerig et al., 1992]. If $c(\cdot)$ is a weighting function decreasing from $c(0) = 1$ to $c(+\infty) = 0$, this can be realized directly in the discrete implementation of the Laplacian (Eq. 13): the contribution $\Delta_u \Sigma$ of the spatial direction u to the Laplace-Beltrami operator is weighted by our decreasing function according to the norm $\|\partial_u \Sigma\|_\Sigma$ of the gradient in that direction. The important point here is that we should evaluate the norm of directional derivatives of the tensor field with our invariant metric. With our finite difference approximations, this leads to the following modified Laplacian:

$$\begin{aligned} \Delta_{aniso} \Sigma(x) &= \frac{d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} c\left(\frac{\|\partial_u \Sigma(x)\|_{\Sigma(x)}}{\|u\|}\right) \frac{\Delta_u^2 \Sigma(x)}{\|u\|^2} \\ &\simeq \frac{2d}{\text{Card}(\mathcal{V})} \sum_{u \in \mathcal{V}} c\left(\frac{\left\|\overrightarrow{\Sigma(x)\Sigma(x+u)}\right\|_{\Sigma(x)}}{\|u\|}\right) \frac{\overrightarrow{\Sigma(x)\Sigma(x+u)}}{\|u\|^2}. \end{aligned} \quad (14)$$

Figures 7 and 8 present example results of this very simple anisotropic filtering scheme on synthetic and real DTI images. We used the function $c(x) = \exp(-x^2/\kappa^2)$, where the threshold κ controls the amount of local regularization: for a gradient magnitude greater than 2 to 3 times κ , there is virtually no regularization, while the field is almost linearly smoothed for gradient magnitudes below a fraction (say 0.1) of κ . For both synthetic and real data, the histogram of the gradient norm is very clearly bimodal so that the threshold κ is easily determined.

In Fig. 7, we generated a tensor field with a discontinuity, and add independent Gaussian noises according to Section 3.8. The anisotropic smoothing perfectly preserves the discontinuity while completely smoothing each region. In this synthetic experiment, we retrieve tensor values that are very close to the initial tensor field. This could be expected since the two regions are perfectly homogeneous. After enough regularization steps, each region is a constant field equal to the mean of the 48 initially noisy tensors of the region. Thus, similarly to the Euclidean mean of identically and independently distributed measurements, we expect the standard deviation of the regularized tensors to be roughly $7 \simeq \sqrt{48}$ times smaller than the one of the noisy input tensors.

Figure 7: **Left:** 3D synthetic tensor field with a clear discontinuity. **Middle:** The field has been corrupted by a Gaussian noise (in the Riemannian sense). **Right:** result of the regularization after 30 iterations (time step $\varepsilon = 0.01$).

In Figure 8, we display the evolution of (a slice of) the tensor field, the norm of the gradient and the fractional anisotropy (FA) at different steps of the anisotropic filtering of a 3D DTI. The FA is based on the normalized variance of the eigenvalues. It shows the differences between an isotropic diffusion in the brain (where the diffusion tensor is represented by a sphere, FA=0) and a highly directional diffusion (cigar-shaped ellipsoid, FA=1). Consequently, the bright regions in the image are the potential areas where nervous fibers are located. One can see that the tensors are regularized in “homogeneous” regions (ventricles, temporal areas), while the main tracts are left unchanged. It is worth noticing that the fractional anisotropy is very well regularized even though this measure has almost nothing in common with our invariant tensor metric.

Figure 9 displays closeups around the ventricles to compare the different regularization methods developed so far. One can see that the Riemannian metric gives much less weight to large tensors, thus providing a regularization which is more robust to outliers. The anisotropic filtering further improves the results by preserving the discontinuities of the tensor scale (e.g. at the boundary of the ventricles), but also the discontinuities of the tensor orientation, which is exactly what is needed for fiber tracking in DTI.

Figure 8: Anisotropic filtering of a DTI slice (time step 0.01, $\kappa = 0.046$). From left to right: at the beginning, after 10 and after 50 iterations. **Top:** A 3D view of the tensors as ellipsoids. The color codes for the direction of the principal eigenvector. The results could be compared with the isotropic Gaussian filtering displayed in Figure 5. **Middle:** Riemannian norm of the gradient. **Bottom:** fractional anisotropy.

Figure 9: Closeup on the results of the different filtering methods around the splenium of the corpus callosum. The color codes for the direction of the principal eigenvector (red: left-right, green: posterior-anterior, blue: inferior-superior). **Top left:** Original image. **Top right:** Gaussian filtering using the flat metric (5x5 window, $\sigma = 2.0$). This metric gives too much weight to tensors with large eigenvalues, thus leading to clear outliers in the ventricles or in the middle of the splenium tract. **Bottom left:** Gaussian filtering using the Riemannian metric (5x5 window, $\sigma = 2.0$). Outliers disappeared, but the discontinuities are not well preserved, for instance in the ventricles at the level of the cortico-spinal tracts (upper-middle part of the images). **Bottom right:** Anisotropic filtering in the Riemannian framework (time step 0.01, 50 iterations). The ventricles boundary is very well conserved with an anisotropic filter and both isotropic (ventricles) and anisotropic (splenium) regions are regularized. Note that the U-shaped tracts at the boundary of the grey/white matter (lower left and right corners of each image) are preserved with an anisotropic filter and not with a Gaussian filter.

6 Regularization and restoration of tensor fields

The pure diffusion is efficient to reduce the noise in the data, but it also reduces the amount of information. Moreover, the amount of smoothing is controlled by the time of diffusion (time step ε times the number of iterations), which is not an easy parameter to tune. At an infinite diffusion time, the tensor field will be completely homogeneous (or homogeneous by part for some anisotropic diffusion schemes), with a value corresponding to the mean of the measurements over the region (with Neumann boundary conditions). Thus, the absolute minimum of our regularization criterion alone is of little interest.

To keep close to the measured tensor field $\Sigma_0(x)$ while still regularizing, a more theoretically grounded approach is to consider an optimization problem with a competition between a data attachment term and a possibly non-linear anisotropic regularization term:

$$C(\Sigma) = \text{Sim}(\Sigma, \Sigma_0) + \lambda \text{Reg}(\Sigma).$$

Like before, the intrinsic evolution equation leading to a local minimum is:

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)}(-\varepsilon (\nabla \text{Sim}(\Sigma, \Sigma_0) + \lambda \nabla \text{Reg}(\Sigma)(x))).$$

6.1 The regularization term

As we saw in the previous section, the simplest regularization criterion is the norm of the gradient of the field $\text{Reg}(F) = \int_{\Omega} \|\nabla F(x)\|^2 dx$. To preserve the discontinuities, the gradient of this criterion (the Laplacian) may be tailored to prevent the smoothing across them, as we have done in Section 5.4. However, there is no more convergence guarantee, since this anisotropic regularization “force” may not derive from a well-posed criterion (energy). Following the pioneer work of [Perona and Malik, 1990], there has been quite an extensive amount of work to propose well posed PDEs for the non-linear, anisotropic and non-stationary regularization of scalar and vector fields (see e.g. [Weickert, 1998, Sapiro, 2001] to cite only a few recent books). Some of these techniques were recently adapted to work on matrix valued fields [Weickert and Brox, 2002] (with the flat metric) or on the rotation manifolds [Tschumperlé and Deriche, 2002].

One of the main idea is to replace the usual simple regularization term $\text{Reg}(F) = \int_{\Omega} \|\nabla F(x)\|^2 dx$ by an increasing function Φ of the norm of the spatial gradient: $\text{Reg}(F) = \int_{\Omega} \Phi(\|\nabla F(x)\|) dx$. With some regularity conditions on the Φ -function [Aubert and Kornprobst, 2001], one can recompute the previous derivations with this Φ -function, and we end-up with:

$$\nabla \text{Reg}(F)(x) = -\text{div} \left(\frac{\Phi'(\|\nabla F\|)}{\|\nabla F\|} \nabla F \right) = - \sum_{i=1}^d \partial_i \left(\frac{\Phi'(\|\nabla F\|)}{\|\nabla F\|} \partial_i F \right).$$

This scheme was used in [Chefd’hotel et al., 2004] with the flat Euclidean metric on tensors, in conjunction with a geometric numerical integration scheme that preserves the rank of the matrix. Their conclusion was that the *rank/signature preserving flow tends to blend the orientation and diffusivity features (eigenvalue swelling effect)*. This rank-signature preserving flow is based on the matrix exponential but does not make any reference to a specific metric. Reformulated in our notations, their evolution equation is:

$\Sigma(x, t + dt) = \exp(dt A(x, t)) \Sigma(x, t) \exp(dt A(x, t))$, where $A(x, t)$ is implicitly related to the driving tangent vector field using: $\partial_t \Sigma(x, t) = -W(x, t) = A(x, t) \Sigma(x, t) + \Sigma(x, t) A(x, t)$. From the first expression of our geodesics (Eq. 3 in Section 3.3), we can see that this is a geodesic marching scheme for our metric. However, they use the flat Euclidean metric on coefficients to evaluate the norm of the spatial gradient of the tensor field. We claim that a consistent framework should use the invariant metric. This leads to very different driving gradient fields $W(x, t)$. Indeed large eigenvalues naturally dominate small ones with the flat metric on coefficients, which means that small eigenvalues are completely smoothed out. This produces the eigenvalue swelling effect they observed. On the contrary, the use of the invariant metric perfectly respects the discontinuities of small and large eigenvalues in our experiments of Sections 5 (even if the anisotropic diffusion PDE is not exactly the same).

We are currently investigating how to adapt the Φ -function formalism to our Riemannian tensor framework. The gradient of the modified criterion can be computed with the invariant metric like in Appendix A.1 to obtain a weighted manifold Laplacian with an additional anisotropic correction term. However, designing an efficient discrete computation scheme is more difficult. We may compute the directional derivatives using finite differences in the flat matrix space and use the intrinsic evolution scheme, but we believe that there are more efficient ways to do it using the exponential map. In the following, we keep the isotropic regularization based on the squared amplitude of the gradient.

6.2 A least-squares attachment term

Usually, one considers that the data (e.g. a scalar image or a displacement vector field $F_0(x)$) are corrupted by a uniform (isotropic) Gaussian noise independent at each spatial position. With a maximum likelihood approach, this amounts to considering a least-squares criterion $Sim(F) = \int_{\Omega} \|F(x) - F_0(x)\|^2 dx$. Like in the previous section, we compute the first order variation by writing the Taylor expansion

$$Sim(F + \varepsilon H) = Sim(F) + 2\varepsilon \int_{\Omega} \langle H(x) | F(x) - F_0(x) \rangle dx + O(\varepsilon^2).$$

This time, the directional derivative $\partial_H Sim(F)$ is directly expressed using a scalar product with H in the proper functional space, so that the steepest ascent direction is $\nabla Sim(F) = 2(F(x) - F_0(x))$.

On the tensor manifold, assuming a uniform (generalized) Gaussian noise independent at each position also leads to a least-squares criterion through a maximum likelihood approach. The only difference is that it uses our Riemannian distance:

$$Sim(\Sigma) = \int_{\Omega} \text{dist}^2(\Sigma(x), \Sigma_0(x)) dx = \int_{\Omega} \left\| \overrightarrow{\Sigma(x)\Sigma_0(x)} \right\|_{\Sigma(x)}^2 dx.$$

Thanks to the properties of the exponential map, one can show that the gradient of the squared distance is: $\nabla_{\Sigma} \text{dist}^2(\Sigma, \Sigma_0) = -2 \overrightarrow{\Sigma \Sigma_0}$ [Pennec, 2004]. One can verify that this is a tangent vector at Σ whereas $\overrightarrow{\Sigma_0 \Sigma}$ is not. Finally, we obtain a steepest ascent direction of our criterion which is very close to the vector case:

$$\nabla Sim(\Sigma)(x) = -2 \overrightarrow{\Sigma(x)\Sigma_0(x)}. \quad (15)$$

6.3 A least-squares attachment term for sparsely distributed tensors

Now, let us consider the case where we do not have a dense measure of our tensor field, but only N measures Σ_i at irregularly distributed sample points x_i . Assuming a uniform Gaussian noise independent at each position still leads to a least-squares criterion:

$$Sim(\Sigma) = \sum_{i=1}^N \text{dist}^2(\Sigma(x_i), \Sigma_i) = \int_{\Omega} \sum_{i=1}^N \text{dist}^2(\Sigma(x), \Sigma_i) \delta(x - x_i) dx.$$

In this criterion, the tensor field $\Sigma(x)$ is related to the data only at the measurement points x_i through the Dirac distributions $\delta(x - x_i)$. If the introduction of distributions may be dealt with for the theoretical differentiation of the criterion with respect to the continuous tensor field Σ , it is a real problem for the numerical implementation. In order to regularize the problem, we consider the Dirac distribution as the limit of the Gaussian function G_{σ} when σ goes to zero. Using that scheme, our criterion becomes the limit case $\sigma = 0$ of:

$$Sim_{\sigma}(\Sigma) = \int_{\Omega} \sum_{i=1}^N \text{dist}^2(\Sigma(x), \Sigma_i) G_{\sigma}(x - x_i) dx. \quad (16)$$

From a practical point of view, we need to use a value of σ which is of the order of the spatial resolution of the grid on which $\Sigma(x)$ is evaluated, so that all measures can at least influence the neighboring nodes.

Now that we came back to a smooth criterion, we may differentiate it exactly as we did for the dense measurement setup. The first order variation is:

$$Sim_{\sigma}(\Sigma + \varepsilon W) = Sim_{\sigma}(\Sigma) - 2\varepsilon \sum_{i=1}^N \int_{\Omega} \left\langle W(x) \mid G_{\sigma}(x - x_i) \overrightarrow{\Sigma(x)\Sigma_i} \right\rangle dx + O(\varepsilon^2),$$

so that we get:

$$\nabla Sim_{\sigma}(x) = -2 \sum_{i=1}^N G_{\sigma}(x - x_i) \overrightarrow{\Sigma(x)\Sigma_i}. \quad (17)$$

6.4 Interpolation through diffusion

With the sparse data attachment term (16) and the isotropic first order regularization term (10), we are looking for a tensor field that minimizes its spatial variations while interpolating (or more precisely approximating at the desired precision) the measurement values:

$$C(\Sigma) = \sum_{i=1}^N \int_{\Omega} G_{\sigma}(x - x_i) \text{dist}^2(\Sigma(x_i), \Sigma_i) + \lambda \int_{\Omega} \|\nabla \Sigma(x)\|_{\Sigma(x)}^2 dx.$$

According to the previous sections, the gradient of this criterion is

$$\nabla C(\Sigma)(x) = -2 \sum_{i=1}^N G_{\sigma}(x - x_i) \overrightarrow{\Sigma(x)\Sigma_i} - 2\lambda \Delta \Sigma(x).$$

Using our finite difference approximation scheme (Eq. 13), the intrinsic geodesic gradient descent scheme (Sec. 3.6) is finally:

$$\Sigma_{t+1}(x) = \exp_{\Sigma_t(x)} \left(\varepsilon \left\{ \sum_{i=1}^N G_\sigma(x - x_i) \overrightarrow{\Sigma(x)\Sigma_i} + \lambda' \sum_{u \in \mathcal{V}} \frac{\overrightarrow{\Sigma(x)\Sigma(x+u)}}{\|u\|^2} \right\} \right) \quad (18)$$

Last but not least, we need an initialization of the tensor field $\Sigma_0(x)$ to obtain a fully operational algorithm. This is easily done with any radial basis function approximation, for instance the renormalized Gaussian scheme that we investigated in Section 4.4. Figure 10 displays the result of this algorithm on the interpolation between 4 tensors. One can see that the soft closest point approximation is well regularized into a constant field equal to the mean of the four tensors if data attachment term is neglected. On the contrary, a very small value of λ is sufficient for regularizing the field between known tensors (as soon as σ is much smaller than the typical spatial distance between two measurements).

The choice of the initialization is a critical issue from a computational point of view. For instance, starting with a constant (or any harmonic) field is a bad idea: there is a null Laplacian everywhere, except at the immediate neighborhood of the sparse tensors, exactly where the data attachment term acts. Thus, we have a potentially destructive competition between the two terms of the criterion in very localized area. On the contrary, starting with a soft closest point approximation leads to a Laplacian which is non null on the boundaries of the Voronoi cells of the measurement points, i.e. the farthest possible place from the sparse measures. In that case, the Laplacian regularization will spread from these boundaries with no constraints until it reaches the counterbalancing forces of the data attachment term in the immediate vicinity of the sparse measurements. Thus, we may expect to reach the maximal efficiency in terms of convergence rate.

7 Conclusion

We propose in this paper an affine invariant metric that endows the space of positive definite symmetric matrices (tensors) with a very regular manifold structure. In particular, tensors with null and infinite eigenvalues are both at an infinite distance of any positive definite symmetric matrix: the cone of positive definite symmetric matrices is replaced by a space which has an infinite development in each of its $n(n+1)/2$ directions. Moreover, there is one and only one geodesic joining any two tensors, and we can even define globally consistent orthonormal coordinate systems of the tangent spaces. Thus, the structure we obtain is very close to a vector space, except that the space is curved. We exemplify some of the good metric properties for some simple statistical operations. For instance, the Karcher mean in Riemannian manifolds has to be defined through a distance-based variational formulation. With our invariant metric on tensors, the existence and uniqueness is insured, which is generally not the case.

A second contribution of the paper is the application of this framework to important geometric data processing problems such as interpolation, filtering, diffusion and restoration of tensor fields. We show that interpolation and Gaussian filtering can be tackled efficiently through a weighted mean computation. However, if weights are easy to define for regularly

Figure 10: Interpolation and extrapolation of tensor values from four measurements using diffusion. **Top left:** The four initial tensor measurements. **Top right:** Initialization of the tensor field using a soft closest point interpolation (mean of the four tensors with a renormalized spatial Gaussian influence). **Bottom left:** result of the diffusion without the data attachment term (1000 iterations, time-step $\varepsilon = 1$, $\lambda = +\infty$). **Bottom right:** result of the diffusion with an attachment term after (1000 iterations, time-step $\varepsilon = 1$, $\lambda = 0.01$, $\sigma = 1$ pixel of the reconstruction grid). The algorithm did in fact converge in about 100 iterations.

sampled tensors (e.g. for linear to tri-linear interpolation), the problem proved to be more difficult for irregularly sampled values. The solution we propose is to consider this type of interpolation as a statistical restoration problem where we want to retrieve a regular tensor field between (possibly noisy) measured tensor values at sparse points. This type of problem is usually solved using a PDE evolution equation. We show that the usual linear regularization (minimizing the magnitude of the gradient) and some anisotropic diffusion schemes can be adapted to our Riemannian framework, provided that the metric of the tensor space is taken into account. We also provide intrinsic numerical schemes for the computation of the gradient and Laplace-Beltrami operators. Finally, simple statistical considerations led us to propose least-squares data attachment criteria for dense and sparsely distributed tensor fields. The differentiation of these criteria is particularly efficient thanks to the use of the Riemannian distance inherited from the chosen metric.

From a theoretical point of view, this paper is a striking illustration of the general framework we are developing since [Pennec, 1996] to provide a rigorous computing environment for geometric objects. This framework is based on the choice of a Riemannian metric on one side, which leads to powerful differential geometry tools such as the exponential maps and geodesic marching techniques, and on the transformation of linear combinations or integrals into minimization problems on the other side. The Karcher mean and the generalized Gaussian distribution are a typical example that we previously investigated [Pennec, 2004]. In the present paper, we provide new examples with interpolation, filtering and PDEs on Riemannian-valued fields.

Many research avenues are still left open, in particular the choice of the metric to use. In a more practical domain, we believe that investigating new intrinsic numerical schemes to compute the derivatives in the PDEs could lead to important gains in accuracy and efficiency. Last but not least, all the results presented in this paper still need to be confronted to other existing methods and validated in the context of medical DTI applications. We are currently investigating another very interesting application field in collaboration with P. Thompson and A. Toga at UCLA: the modeling and analysis of the variability of the brain anatomy.

References

- [Aubert and Kornprobst, 2001] Aubert, G. and Kornprobst, P. (2001). *Mathematical problems in image processing*, volume 147 of *Applied Mathematical Sciences*. Springer.
- [Basser et al., 1994] Basser, P., Mattiello, J., and Bihan, D. L. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267.
- [Batchelor et al., 2001] Batchelor, P., Hill, D., Calamante, F., and Atkinson, D. (2001). Study of the connectivity in the brain using the full diffusion tensor from MRI. In Insana, M. and Leahy, R., editors, *Proc. of the 17th Int. Conf. on Information Processing in Medical Imaging (IPMI 2001)*, volume 2082 of *LNCIS*, pages 121–133. Springer Verlag.
- [Bhatia, 2003] Bhatia, R. (2003). On the exponential metric increasing property. *Linear Algebra and its Applications*, 375:211–220.

- [Cazals and Boissonnat, 2001] Cazals, F. and Boissonnat, J.-D. (2001). Natural coordinates of points on a surface. *Comp. Geometry Theory and Applications*, 19:155–173.
- [Chefd’hotel et al., 2002] Chefd’hotel, C., Tschumperlé, D., Deriche, R., and Faugeras, O. (2002). Constrained flows of matrix-valued functions: Application to diffusion tensor regularization. In Heyden, A., Sparr, G., Nielsen, M., and Johansen, P., editors, *Proc. of ECCV 2002*, volume 2350 of *LNCS*, pages 251–265. Springer Verlag.
- [Chefd’hotel et al., 2004] Chefd’hotel, C., Tschumperlé, D., Deriche, R., and Faugeras, O. (2004). Regularizing flows for constrained matrix-valued images. *J. Math. Imaging and Vision*, 20(1-2):147–162.
- [Coulon et al., 2001] Coulon, O., Alexander, D., and Arridge, S. (2001). A regularization scheme for diffusion tensor magnetic resonance images. In Insana, M. and Leahy, R., editors, *Proc. of the 17th Int. Conf. on Information Processing in Medical Imaging (IPMI 2001)*, volume 2082 of *LNCS*, pages 92–105. Springer Verlag.
- [Coulon et al., 2004] Coulon, O., Alexander, D., and Arridge, S. (2004). Diffusion tensor magnetic resonance image regularization. *Medical Image Analysis*, 8(1):47–67.
- [Fillard et al., 2003] Fillard, P., Gilmore, J., Piven, J., Lin, W., and Gerig, G. (2003). Quantitative analysis of white matter fiber properties along geodesic paths. In Ellis, R. E. and Peters, T. M., editors, *Proc. of MICCAI’03, Part II*, volume 2879 of *LNCS*, pages 16–23, Montreal. Springer Verlag.
- [Fletcher and Joshi, 2004] Fletcher, P. T. and Joshi, S. C. (2004). Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In *Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis, ECCV 2004 Workshops CVAMIA and MMBIA, Prague, Czech Republic, May 15, 2004*, volume 3117 of *LNCS*, pages 87–98. Springer.
- [Förstner and Moonen, 1999] Förstner, W. and Moonen, B. (1999). A metric for covariance matrices. In Krumm, F. and Schwarze, V. S., editors, *Qua vadis geodesia...? Festschrift for Erik W. Grafarend on the occasion of his 60th birthday*, number 1999.6 in Tech. Report of the Dpt of Geodesy and Geoinformatics, pages 113–128. Stuttgart University.
- [Gallot et al., 1993] Gallot, S., Hulin, D., and Lafontaine, J. (1993). *Riemannian Geometry*. Springer Verlag, 2nd edition edition.
- [Gamkrelidze, 1991] Gamkrelidze, R., editor (1991). *Geometry I*, volume 28 of *Encyclopaedia of Mathematical Sciences*. Springer Verlag.
- [Gerig et al., 1992] Gerig, G., Kikinis, R., Kübler, O., and Jolesz, F. (1992). Nonlinear anisotropic filtering of MRI data. *IEEE Transactions on Medical Imaging*, 11(2):221–232.
- [Helgason, 1978] Helgason, S. (1978). *Differential Geometry, Lie groups, and Symmetric Spaces*. Academic Press.

- [Kendall and Moran, 1963] Kendall, M. and Moran, P. (1963). *Geometrical probability*. Number 10 in Griffin’s statistical monographs and courses. Charles Griffin & Co. Ltd.
- [Kendall, 1990] Kendall, W. (1990). Probability, convexity, and harmonic maps with small image I: uniqueness and fine existence. *Proc. London Math. Soc.*, 61(2):371–406.
- [Kobayashi and Nomizu, 1969] Kobayashi, S. and Nomizu, K. (1969). *Foundations of differential geometry*, volume II of *Interscience tracts in pure and applied mathematics*. John Wiley & Sons.
- [Le Bihan et al., 2001] Le Bihan, D., Manguin, J.-F., Poupon, C., Clark, C., Pappata, S., Molko, N., and Chabriat, H. (2001). Diffusion tensor imaging: Concepts and applications. *Journal Magnetic Resonance Imaging*, 13(4):534–546.
- [Lenglet et al., 2004a] Lenglet, C., Rousson, M., Deriche, R., and Faugeras, O. (2004a). Statistics on multivariate normal distributions: A geometric approach and its application to diffusion tensor MRI. Research Report 5242, INRIA.
- [Lenglet et al., 2004b] Lenglet, C., Rousson, M., Deriche, R., and Faugeras, O. (2004b). Toward segmentation of 3D probability density fields by surface evolution: Application to diffusion MRI. Research Report 5243, INRIA.
- [Meijering, 2002] Meijering, E. (2002). A chronology of interpolation: From ancient astronomy to modern signal and image processing. *Proceedings of the IEEE*, 90(3):319–342.
- [Nomizu, 1954] Nomizu, K. (1954). Invariant affine connections on homogeneous spaces. *American J. of Math.*, 76:33–65.
- [Pennec, 1996] Pennec, X. (1996). *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).
- [Pennec, 1999] Pennec, X. (1999). Probabilities and statistics on Riemannian manifolds: Basic tools for geometric measurements. In Cetin, A., Akarun, L., Ertuzun, A., Gurcan, M., and Yardimci, Y., editors, *Proc. of Nonlinear Signal and Image Processing (NSIP’99)*, volume 1, pages 194–198, June 20-23, Antalya, Turkey. IEEE-EURASIP.
- [Pennec, 2004] Pennec, X. (2004). Probabilities and statistics on Riemannian manifolds: A geometric approach. Research Report 5093, INRIA. An extended version will appear in the Int. Journal of Mathematical Imaging and Vision.
- [Pennec and Ayache, 1998] Pennec, X. and Ayache, N. (1998). Uniform distribution, distance and expectation problems for geometric features processing. *Journal of Mathematical Imaging and Vision*, 9(1):49–67. A preliminary version appeared as INRIA Research Report 2820, March 1996.
- [Pennec and Thirion, 1997] Pennec, X. and Thirion, J.-P. (1997). A framework for uncertainty and validation of 3D registration methods based on points and frames. *Int. Journal of Computer Vision*, 25(3):203–229.

- [Perona and Malik, 1990] Perona, P. and Malik, J. (1990). Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 12(7):629–639.
- [Poincaré, 1912] Poincaré, H. (1912). *Calcul des probabilités*. 2nd edition, Paris.
- [Rey et al., 2002] Rey, D., Subsol, G., Delingette, H., and Ayache, N. (2002). Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. *Medical Image Analysis*, 6(2):163–179.
- [Sapiro, 2001] Sapiro, G. (2001). *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press.
- [Sibson, 1981] Sibson, R. (1981). A brief description of natural neighbour interpolation. In Barnett, V., editor, *Interpreting Multivariate Data*, pages 21–36. John Wiley & Sons, Chichester.
- [Skovgaard, 1984] Skovgaard, L. (1984). A Riemannian geometry of the multivariate normal model. *Scand. J. Statistics*, 11:211–223.
- [Thévenaz et al., 2000] Thévenaz, P., Blu, T., and Unser, M. (2000). Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19(7):739–758.
- [Tschumperlé, 2002] Tschumperlé, D. (2002). *PDE-Based Regularization of Multivalued Images and Applications*. PhD thesis, University of Nice-Sophia Antipolis.
- [Tschumperlé and Deriche, 2002] Tschumperlé, D. and Deriche, R. (2002). Orthonormal vector sets regularization with PDE’s and applications. *International Journal on Computer Vision*, 50(3):237–252.
- [Weickert, 1998] Weickert, J. (1998). *Anisotropic Diffusion in Image Processing*. Teubner-Verlag.
- [Weickert and Brox, 2002] Weickert, J. and Brox, T. (2002). Diffusion and regularization of vector- and matrix-valued images. In Nashed, M. and Scherzer, O., editors, *Inverse Problems, Image Analysis, and Medical Imaging.*, volume 313 of *Contemporary Mathematics*, pages 251–268, Providence. AMS.
- [Westin et al., 2002] Westin, C., Maier, S., Mamata, H., Nabavi, A., Jolesz, F., and Kikinis, R. (2002). Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108.

A Tensor regularization: the Laplace-Beltrami operator

A.1 Gradient of the L^2 regularization of a tensor field

Let $\Sigma(x) \in \mathcal{Sym}_n^+$ be a tensor field over $\mathbb{R}^d$, and $(x_1, \dots, x_d)$ be an orthonormal coordinate system. To simplify the notations, we use in this section ∂_i for the spatial derivative $\partial/(\partial x_i)$ and we do not specify the (spatial) integration variable x . The L^2 regularization criterion is:

$$Reg(\Sigma) = \int_{\Omega} \|\nabla \Sigma(x)\|_{\Sigma(x)}^2 dx = \sum_{i=1}^d \int_{\Omega} \|\partial_i \Sigma\|_{\Sigma}^2 = \sum_{i=1}^d \int_{\Omega} \text{Tr}\left((\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)}\right).$$

Using the Taylor expansion $(\Sigma + \varepsilon W)^{(-1)} = \Sigma^{(-1)} - \varepsilon \Sigma^{(-1)} W \Sigma^{(-1)} + O(\varepsilon^2)$ in the Taylor expansion of our regularization criterion and identifying the first order term to $Reg(\Sigma + \varepsilon W) = Reg(\Sigma) + \varepsilon \partial_W Reg + O(\varepsilon^2)$, we get the directional derivative:

$$\partial_W Reg = 2 \sum_{i=1}^d \int_{\Omega} \text{Tr}\left((\partial_i W) \Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)} - (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)} W \Sigma^{(-1)}\right)$$

The main goal is to find out the field of tangent vectors $\nabla Reg(x) \in T_{\Sigma(x)} \mathcal{Sym}_n^+$ such that, by definition of the gradient, we have the equality: $\partial_W Reg = \int_{\Omega} \langle W | \nabla Reg \rangle_{\Sigma} dx$ for every field of tangent vectors $W(x) \in T_{\Sigma(x)} \mathcal{Sym}_n^+$. As the above expression of $\partial_W Reg$ is in the standard Euclidean chart (matrix coefficients), we shall safely use the computations of the previous sections. Notice that we are using the matrix coefficients only as a chart and not as a metric. Let $\Lambda_i = \Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)}$. We get:

$$\partial_W Reg = 2 \sum_{i=1}^d \int_{\Omega} \left(\text{Tr}((\partial_i W) \Lambda_i) - \langle W | (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \rangle_{\Sigma} \right)$$

Now, assuming the proper Neumann boundary conditions, we can apply the previous integration by part formula $\int_{\Omega} \text{Tr}((\partial_i W) \Lambda_i) = - \int_{\Omega} \text{Tr}(W (\partial_i \Lambda_i))$ to the first term:

$$\begin{aligned} \partial_W Reg &= -2 \sum_{i=1}^d \int_{\Omega} \left(\text{Tr}\left(W \Sigma^{(-1)} (\Sigma (\partial_i \Lambda_i) \Sigma) \Sigma^{(-1)}\right) + \langle W | (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \rangle_{\Sigma} \right) \\ &= -2 \sum_{i=1}^d \int_{\Omega} \langle W | \Sigma (\partial_i \Lambda_i) \Sigma + (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \rangle_{\Sigma} \end{aligned}$$

We have obtained the expression that defines the gradient of our regularization criterion:

$$\nabla Reg = -2 \sum_{i=1}^d \left(\Sigma (\partial_i \Lambda_i) \Sigma + (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma) \right)$$

To compute explicitly its value, let us observe first that $\partial_i \Sigma^{(-1)} = -\Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)}$ because $\partial_i (\Sigma^{(-1)} \Sigma) = 0$. Thus, thanks to the chain rule, we have:

$$\Sigma (\partial_i \Lambda_i) \Sigma = \Sigma \partial_i (\Sigma^{(-1)} (\partial_i \Sigma) \Sigma^{(-1)}) \Sigma = \partial_i^2 \Sigma - 2 (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma)$$

Eventually, we end up with $\nabla \text{Reg}(\Sigma) = -2\Delta\Sigma$, where Δ is the Laplace-Beltrami operator on our manifold:

$$\Delta\Sigma = \sum_{i=1}^d \Delta_i \Sigma \quad \text{with} \quad \Delta_i \Sigma = \partial_i^2 \Sigma - (\partial_i \Sigma) \Sigma^{(-1)} (\partial_i \Sigma)$$

A.2 Numerical implementation of the Laplace-Beltrami operator

From the Taylor expansion of a tensor field (considered as a matrix field) Σ at x , we have $\Sigma(x + \varepsilon u) = \Sigma(x) + \varepsilon \partial_u \Sigma(x) + \varepsilon^2 \partial_u^2 \Sigma(x)/2 + \varepsilon^3 \partial_u^3 \Sigma(x)/6 + O(\varepsilon^4)$. Thus, we may approximate the first and second order tensor derivatives by their Euclidean derivatives:

$$\begin{aligned} \partial_u \Sigma(x) &= \frac{1}{2} (\Sigma(x + u) - \Sigma(x - u)) + O(\|u\|^3) \\ \partial_u^2 \Sigma(x) &= (\Sigma(x + u) - \Sigma(x)) + (\Sigma(x - u) - \Sigma(x)) + O(\|u\|^4) \end{aligned}$$

This finally gives us a fourth order approximation of the Laplace-Beltrami operator in the spatial direction u :

$$\begin{aligned} \Delta_u \Sigma(x) &= \partial_u^2 \Sigma - 2 (\partial_u \Sigma) \Sigma^{(-1)} (\partial_u \Sigma) \\ &= \Sigma(x + u) + \Sigma(x - u) - 2\Sigma(x) \\ &\quad - \frac{1}{2} (\Sigma(x + u) - \Sigma(x - u)) \Sigma^{(-1)} (\Sigma(x + u) - \Sigma(x - u)) + O(\|u\|^4) \end{aligned}$$

However, this numerical scheme is extrinsic since it is based on (Euclidean) differences of tensors. We propose here an intrinsic scheme based on the exponential chart at the current point: we claim that $\overrightarrow{\Sigma(x)\Sigma(x+u)} + \overrightarrow{\Sigma(x)\Sigma(x-u)}$ is a forth order approximation of the Laplace Beltrami operator in the direction u . Indeed, we have

$$\begin{aligned} \overrightarrow{\Sigma(x)\Sigma(x+\varepsilon u)} &= \Sigma^{\frac{1}{2}}(x) \log \left(\Sigma^{-\frac{1}{2}}(x) \Sigma(x + \varepsilon u) \Sigma^{-\frac{1}{2}}(x) \right) \Sigma^{\frac{1}{2}}(x) \\ &= \Sigma^{\frac{1}{2}}(x) \log \left(\text{Id} + \varepsilon W + \frac{\varepsilon^2}{2} H + O(\varepsilon^3) \right) \Sigma^{\frac{1}{2}}(x) \end{aligned}$$

where we put $W = \Sigma^{-\frac{1}{2}} \partial_u \Sigma \Sigma^{-\frac{1}{2}}$ and $H = \Sigma^{-\frac{1}{2}} \partial_u^2 \Sigma \Sigma^{-\frac{1}{2}}$. From the Log series (Eq. 1), we get:

$$\overrightarrow{\Sigma(x)\Sigma(x+\varepsilon u)} = \Sigma^{\frac{1}{2}} \left[+\varepsilon W + \frac{\varepsilon^2}{2} H - \frac{1}{2} \left(\varepsilon^2 W^2 + \frac{\varepsilon^3}{2} (WH + HW) \right) + \frac{\varepsilon^3}{3} W^3 + O(\varepsilon^4) \right] \Sigma^{\frac{1}{2}}$$

The Taylor expansion of $\overrightarrow{\Sigma(x)\Sigma(x-\varepsilon u)}$ is obtained by replacing ε by $-\varepsilon$, so that we finally end up with $\overrightarrow{\Sigma(x)\Sigma(x+\varepsilon u)} + \overrightarrow{\Sigma(x)\Sigma(x-\varepsilon u)} = \Sigma^{\frac{1}{2}} [\varepsilon^2 H - \varepsilon^2 W^2 + O(\varepsilon^4)] \Sigma^{\frac{1}{2}}$, which proves that

$$\overrightarrow{\Sigma(x)\Sigma(x+u)} + \overrightarrow{\Sigma(x)\Sigma(x-u)} = \partial_u^2 \Sigma - 2 (\partial_u \Sigma) \Sigma^{(-1)} (\partial_u \Sigma) + O(\|u\|^4) = \Delta_u \Sigma + O(\|u\|^4)$$