

Fast $GL(n)$ -Invariant Framework for Tensors Regularization

Yaniv Gur · Ofer Pasternak · Nir Sochen

Received: 27 January 2008 / Accepted: 17 November 2008 / Published online: 5 December 2008
© Springer Science+Business Media, LLC 2008

Abstract We propose a novel framework for regularization of symmetric positive-definite (SPD) tensors (e.g., diffusion tensors). This framework is based on a local differential geometric approach. The manifold of symmetric positive-definite (SPD) matrices, P_n , is parameterized via the Iwasawa coordinate system. In this framework distances on P_n are measured in terms of a natural $GL(n)$ -invariant metric. Via the mathematical concept of fiber bundles, we describe the tensor-valued image as a section where the metric over the section is induced by the metric over P_n . Then, a functional over the sections accompanied by a suitable data fitting term is defined. The variation of this functional with respect to the Iwasawa coordinates leads to a set of $\frac{1}{2}n(n+1)$ coupled equations of motion. By means of the gradient descent method, these equations of motion define a Beltrami flow over P_n . It turns out that the local coordinate approach via the Iwasawa coordinate system results in very simple numerics that leads to fast convergence of the algorithm. Regularization results as well as results of fibers tractography for DTI are presented.

Keywords Beltrami framework · Differential geometry · PDEs · Tensor-valued image · Diffusion tensor imaging

Y. Gur (✉) · N. Sochen
Department of Applied Mathematics, Tel Aviv University,
Ramat Aviv, Tel Aviv 69978, Israel
e-mail: yanivg@post.tau.ac.il

N. Sochen
e-mail: sochen@post.tau.ac.il

O. Pasternak
School of Computer Science, Tel Aviv University, Ramat Aviv,
Tel Aviv 69978, Israel
e-mail: oferpas@post.tau.ac.il

1 Introduction

Tensors are nowadays an important tool for image analysis as well as a source for physical information. Tensor-valued images are obtained from applications such as Structure Tensor (ST) analysis (Weickert 1998) and Diffusion Tensor Imaging (DTI) (Pierpaoli et al. 1996).

In this paper we focus on symmetric positive-definite tensors. We suggest a general solution for the regularization problem of tensor-valued images of this type, and we demonstrate the proposed framework on DTI datasets. This framework is based on Riemannian geometry which is the mathematical basis for various works on manipulating of diffusion tensors (Batchelor et al. 2005; Fletcher and Joshi 2007; Moakher 2005).

Acquisition of high-resolution DT images generally leads to noisy images. In order to extract important features such as axon fiber bundles and to obtain smooth fiber tracts, the noise has to be removed. Therefore, regularization of Diffusion Tensor MRI data has been attracting much attention over the last few years (e.g., Ched'hotel et al. 2004; Coulon et al. 2001; Feddern et al. 2006; Moakher and Zerai 2007; Pasternak et al. 2008; Pennec et al. 2006; Tschumperlé and Deriche 2002; Wang et al. 2004 and the book Weickert and Hagen 2005). In DTI each image voxel is described in terms of a 3×3 symmetric positive-definite (SPD) matrix. Regularization of DTI data brings new challenges which have not been confronted before in the context of scalar image regularization. The main challenge is to perform a fast and efficient regularization process that preserves the properties of these matrices (i.e., symmetry and positive-definiteness) and respects the structure of the image.

Structure-preserving flows for SPD matrices may be defined via the polar decomposition (Tschumperlé and Deriche 2002). In this method the SPD matrix is decomposed

into its eigenvalues and eigenvectors. Then, regularization is achieved by working on the eigenvalues (diffusivities) and eigenvectors (directions) separately. However, the polar decomposition is not unique. There are two possible different decompositions for a SPD matrix. In these two decompositions the directions of the eigenvectors are opposite. Therefore, due to its non-uniqueness, artificial discontinuities may be created. As a result, one has to apply a heuristic realignment step in every iteration of the numerical solution. Also, determining the order of the eigenvalues is subject to a sorting bias, especially in oblate or prolate tensors (Basser and Pierpaoli 1996). A different approach is the Riemannian framework which was proposed by Pennec et al. (2006). In this framework a *global* affine-invariant Riemannian distance on P_n , the manifold of $n \times n$ SPD matrices, is defined. Then, by using this distance measure, a structure-preserving flow on P_n is constructed. However, the main deficiency of this framework is its running time. This framework makes intensive use of matrix operations such as inverse, log and exponent in every iteration, where these matrix operations are time consuming. In order to reduce running times, a different framework which is based on the Log-Euclidean distance was presented in Fillard et al. (2007). In that framework the flow was defined on the vector space of symmetric matrices. The log and exp matrix operations are then used to map the SPD matrices to symmetric matrices and back, respectively. Compared to the Riemannian framework, the use of the log and exp matrix operations was reduced significantly. However, the Log-Euclidean distance is no longer affine-invariant.

In this paper we present fast $GL(n)$ -invariant Riemannian framework for regularization of tensors in P_n . In order to overcome the deficiencies of the former methods a local coordinate approach is adopted. In this framework the P_n manifold is described via *local* coordinates. Then, P_n is turned into a Riemannian manifold by the definition of natural $GL(n)$ -invariant metric in terms of these coordinates. The local coordinate approach has two main advantages:

1. The numerics is fast and simple since it is free from matrix operations. In this method there are only operations between scalars to calculate (e.g., finite difference).
2. It is $GL(n)$ -invariant (i.e., invariant under the action of any real invertible matrix of order n).

A proper choice of coordinate system to parameterize P_n is important both analytically and numerically. In this paper we show that the *Iwasawa* decomposition of SPD matrices yields a natural coordinate system. The choice of this coordinate system simplifies the analytical as well as the numerical calculations.

The basis of our framework is the definition of a product space via the mathematical notion of fiber bundles. The product space is composed of two spaces: the image domain

(the base manifold), Ω , and the feature space of $n \times n$ symmetric positive-definite (SPD) matrices, P_n (the fiber). Using this concept a DT image may be described by the fiber bundle $\Omega \times P_3$ where Ω is in principal a two- or three-dimensional Euclidean space (DT slice or volume, respectively). A unique assignment of a feature space element (e.g., tensor) to each point of the base manifold (the image domain) is a *section* in the fiber bundle. By means of the Beltrami framework we define a functional over sections which is composed of a regularizer term and a data fitting term (fidelity term). The variation of this functional with respect to each one of the Iwasawa coordinates yields a set of $\frac{1}{2}n(n+1)$ coupled equations of motion (i.e., $n=3$ for DTI). These equations of motion lead to a Beltrami flow on P_n via the gradient-descent method. Regularization of DTI datasets is performed by extracting the Iwasawa coordinates from the data and then solving this set of equations for these coordinates.

The paper is organized as follows: In Sect. 2 we give a short introduction to fiber bundles and explain how a tensor-valued image is described via this concept. In Sect. 3 we discuss the basics of the Riemannian geometry of P_n . In this section we define a $GL(n)$ -invariant metric that turns P_n into a Riemannian symmetric space. The metric over the sections in the fiber bundle is then defined in terms of this metric. In Sect. 4 we introduce a functional over sections by means of the Beltrami framework. Using calculus of variations we derive the equations of motion with respect to the coordinates on the section. The parameterization of P_n via the Iwasawa coordinates is discussed in Sect. 5. We show that for P_3 in particular we get six coupled Beltrami equations for the six Iwasawa coordinates. Finally, in Sect. 6 we present regularization results on real volumetric DTI datasets. We show that using this framework we are able to denoise these datasets efficiently. Concluding remarks are drawn in Sect. 7.

2 Images as Fibered Space

Generally, an image may be described locally as the product of two spaces: The image domain (two- or three-dimensional usually) and the feature space which is composed of objects such as intensity, RGB values, orientation fields, etc. A product of two-spaces may be described mathematically via the concept of a fiber bundle. A fiber bundle is characterized by the total space E , the base manifold B , a structural group G (a group of homeomorphisms of the fiber F onto itself) together with a continuous surjective map $\pi : E \mapsto B$ which is called a projection. Since the dimension of E is higher than the dimension of B , the projection sends many points in E to one point in B . Locally, the total space E is described as product of the base manifold and the fiber: $E = B \times F$. Globally, this may not be the

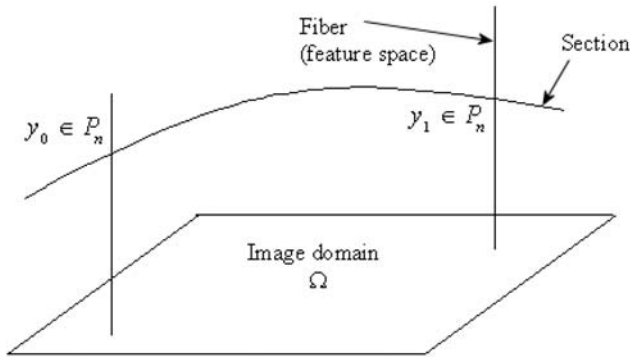


Fig. 1 Description of an image as a section of a fiber bundle. The base manifold B is the image domain Ω , the fiber is the feature space F (e.g., intensity values, color, DT). Here the image domain is Euclidean. Locally the bundle is described as a direct product $E = B \times F$ where E is the total space. In DTI the fiber is taken to be the space of 3×3 SPD matrices, P_3 . A DT image is composed of a particular selection of an SPD matrix to each point of the base manifold. Thus, it is a section in the fiber bundle. The section is referred to here as the image manifold

case since global features of the total space are not observed locally (e.g., the Möbius strip where the twist cannot be observed locally). When the identification of the total space with the direct product of the base space and the fiber is possible, this is known as a trivial bundle.

In computer vision the image domain is identified with the base manifold, and the feature space is identified with the fiber. The total space in this language is the *spatial-feature manifold*. For example, using this concept, a graylevel image is described locally by the product $\mathbb{R}^2 \times \mathbb{R}^+$. In this case the graylevels are the fiber. Since an image selects just one graylevel for each pixel, we need the mathematical notion of a *section*.

A section of the bundle is a mapping $f : B \rightarrow E$ such that $\pi(f(p)) = p$ for any $p \in B$. In our language, a given image is a section in a fiber bundle. For every pixel in the image domain we have a unique assignment of a feature space object. In this paper we refer the section as the *image manifold*.

If we go back to DTI then we have a tensor-valued image where for each pixel there is unique assignment of a three-dimensional SPD matrix. Thus, a DT image is a section in the fiber bundle with the map $f : \mathbb{R}^m \rightarrow \mathbb{R}^m \times P_3$, where \mathbb{R}^m is in practice the two- or three-dimensional Euclidean image domain (DTI slice or volume, respectively) and P_3 , which is the fiber, denotes the space of 3×3 SPD matrices. Thus, the spatial-feature manifold is described by the product $\mathbb{R}^m \times P_3$. The structural group in this case is simply $GL(3, \mathbb{R})$ which is the group of real invertible matrices of order 3. The action of $GL(3, \mathbb{R})$ on P_3 is given by $Y[g] = g^T Y g$ for $Y \in P_3$ and $g \in GL(3, \mathbb{R})$. This description is extendable for any n with the fiber-bundle $\mathbb{R}^m \times P_n$. In this case the structural group is $GL(n, \mathbb{R})$.

3 The Riemannian Geometry of P_n

For regularization of volumetric DTI datasets we take the image domain to be a three-dimensional Euclidean space. In order to write the metric over the spatial-feature manifold, the metric over the feature space P_n has to be defined. In this section we briefly review the important facts about the Riemannian geometry of P_n . Detailed discussions on this topic can be found in Jorgenson and Lang (2005), Lang (1999).

The space of SPD matrices is a symmetric space of non-compact type (Terras 1988). Also, it was given in Lang (1999) as an example of a Riemannian manifold of nonpositive curvature. It may be identified with an open cone in \mathbb{R}^m where $m = n(n + 1)/2$, i.e., for any $V, W \in P_n$ and for any positive scalar $c > 0$ we have $V + W \in P_n$ and $cV, cW \in P_n$. Since P_n is a connected manifold, according to the Hopf-Rinow theorem any geodesic segment can be extended indefinitely. Consequently, any two points in this space may be joined by a minimizing geodesic where its length is the geodesic distance between the two points. Moreover, it has been shown in Jost (2001) that there is precisely one geodesic connecting any two points on P_n . The geodesic distance between any two points $A, B \in P_n$ is given by Lang (1999)

$$d(A, B) = \left(\sum_{i=1}^n \log^2 \lambda_i \right)^{1/2}, \tag{1}$$

where λ_i are the eigenvalues of the matrix $A^{-1}B$. This geodesic distance is the so-called Riemannian distance which was used in Pennec et al. (2006) as a distance measure between two SPD tensors. With respect to this distance function, P_n is a geodesically complete Riemannian space.

Without getting into details, the tangent space at every point $Y \in P_n$ may be identified with the vector space of $n \times n$ symmetric matrices, SYM_n . Thus, the Riemannian metric at the point Y is defined in terms of the scalar product on SYM_n as in Lang (1999)

$$ds_{P_n}^2 = \text{tr}((Y^{-1}dY)^2), \tag{2}$$

where $Y^{-1} = (y_{ij})_{1 \leq i, j \leq n}^{-1}$ and $dY = (dy_{ij})_{1 \leq i, j \leq n}$. This metric is by definition positive-definite (Lang 1999; Terras 1988). Also, by defining the action of $g \in GL(n)$ on $Y \in P_n$ as $Y[g] = g^T Y g$, it can be easily shown that this metric is invariant under the action of $GL(n)$: Let $W = Y[g]$ where the differential is given by $dW = dY[g]$, then, upon plugging everything in $ds_{P_n}^2$ it follows that

$$\begin{aligned} ds_{P_n}^2 &= \text{tr}((Y^{-1}dY)^2) = \text{tr}((gW^{-1}g^T g^{-T}dWg^{-1})^2) \\ &= \text{tr}((W^{-1}dW)^2). \end{aligned} \tag{3}$$

It can be easily verified that this metric is also invariant with respect to the inversion map $Y \mapsto Y^{-1}$. Thus, this map is

an involutive isometry on P_n with respect to this metric. Therefore, this metric turns P_n into a Riemannian symmetric space. Another type of Riemannian symmetric spaces, namely compact Lie groups has been discussed by two of the authors in Gur and Sochen (2007). A bi-invariant metric of the same form as (2) has been used to formulate a coordinate-free regularization framework for Lie groups.

After the metric over P_n is introduced, the metric over the spatial-feature manifold $M = \mathbb{R}^3 \times P_n$ is simply given by $ds_M^2 = \sum_{i=1}^3 dx_i^2 + \text{tr}((Y^{-1}dY)^2)$. Later we choose the coordinates to parameterize P_n .

4 Calculus of Variations in Fiber Bundles

Using functional analysis on sections in the fiber bundle we are able to derive a flow on the section. This will be done via the Beltrami framework (see Kimmel et al. 1997; Shafirir et al. 2005 and references therein). We now briefly review this framework’s main ideas.

4.1 The Beltrami Framework

Denote by (Σ, γ) the image manifold and its metric and by (M, h) the embedding spatial-feature manifold (the fiber bundle) and its metric, then, the section of interest is expressed locally by the map $X : \Sigma \rightarrow M$. A functional over the space of sections is given in local coordinates by the following expression

$$S(X) = \int_{\Sigma} d^n x \sqrt{\gamma} \gamma^{\mu\nu} \frac{\partial X^i}{\partial x^\mu} \frac{\partial X^j}{\partial x^\nu} h_{ij}(X). \tag{4}$$

The case where $n = 2$ is known in the literature as the Polyakov action (Polyakov 1981). In this paper we treat volumetric images (e.g., DTI volume) and hence $n = 3$. The integration is taken over the two- or three-dimensional image manifold where γ is the determinant of the image metric, $(\gamma^{\mu\nu})$ denotes the inverse of the image manifold’s metric tensor and (h_{ij}) is the embedding space metric tensor (the spatial-feature fiber-bundle manifold). The coordinates in the spatial-feature space are denoted by X^i . The values of μ and ν range from 1 to $\dim \Sigma$ and the values of the i and j indices range from 1 to $\dim M$.

Using calculus of variations with respect to the embedding coordinates X^i , we obtain the Euler-Lagrange equations for this action:

$$\begin{aligned} & \frac{1}{\sqrt{\gamma}} \partial_\mu \left(\sqrt{\gamma} \gamma^{\mu\nu} \partial_\nu X^i \right) + \Gamma_{jk}^i \gamma^{\mu\nu} \partial_\mu X^j \partial_\nu X^k \\ & = \Delta_\gamma + \Gamma_{ij}^k \langle \nabla X^j, \nabla X^k \rangle_\gamma = 0, \end{aligned} \tag{5}$$

where Δ_γ is the Laplace-Beltrami operator that is the generalization of the Laplacian to manifolds with metric γ . The

solutions to these Euler-Lagrange equations are known as *harmonic maps*. By the gradient descent method we obtain a set of PDEs (i.e., Beltrami equations) with respect to the embedding coordinates. However, since the coordinates of the image domain are fixed, the interesting equations are for the coordinates of the fiber (P_n in our case). Hence,

$$X_t^i = \frac{1}{\sqrt{\gamma}} \partial_\mu \left(\sqrt{\gamma} \gamma^{\mu\nu} \partial_\nu X^i \right) + \Gamma_{jk}^i \gamma^{\mu\nu} \partial_\mu X^j \partial_\nu X^k, \tag{6}$$

where $i = 1, \dots, (\dim(M) - \dim(\Sigma))$ and where Γ_{jk}^i are the Christoffel symbols. When the embedding space is Euclidean all the Christoffel symbols vanish. The Christoffel symbols are calculated with respect to (h_{ij}) as follows:

$$\Gamma_{jk}^i = \frac{1}{2} h^{il} (\partial_j h_{lk} + \partial_k h_{jl} - \partial_l h_{jk}). \tag{7}$$

The Beltrami equations are solved together at each iteration using standard finite-differences schemes. The components of the induced metric and the Christoffel symbols are evaluated at each iteration. An important result of this framework is that no constraint on the positive definiteness of the matrices is needed. Indeed the ellipticity of the Laplace-Beltrami operator that generates the flow implies the validity of the extremum principle. This in turns mean that *the positive definiteness of the initial condition is a necessary and sufficient condition for the flow to stay on P_n .*

The Induced Metric

The metric over the image manifold (or section) is induced from the fiber bundle’s metric. In this way the flow depends on the geometry of the data and not only on the geometry of the image domain. Moreover, the induced metric is a dynamical variable which changes along the flow.

We assume an isometric embedding, e.g., infinitesimal distances on the spatial-feature manifold, M , are equal infinitesimal distances on the image manifold, Σ . This assumption yields the pull-back metric induced by $X : \gamma = X^*h$. Its components are given by

$$\gamma_{\mu\nu}(x) = \frac{\partial X^i}{\partial x^\mu} \frac{\partial X^j}{\partial x^\nu} h_{ij}(X). \tag{8}$$

Thus, the induced metric is actually calculated via the chain rule. Consequently, for a DTI volume where $M = \mathbb{R}^3 \times P_3$ we may write $dY = \sum_{i=1}^3 \frac{\partial Y}{\partial x_i} dx_i$ where $Y \in P_3$. Then, we plug this expression into the metric on M to get the components of the induced metric:

$$\gamma_{\mu\nu} = \delta_{\mu\nu} + \beta \text{tr}(Y^{-1} \partial_\mu Y Y^{-1} \partial_\nu Y), \tag{9}$$

where $\mu, \nu = 1, \dots, 3$ are the indices of the local coordinates on the image manifold. The β parameter is introduced

here in order to determine the ratio between the feature space and the image domain distances. It also determines the nature of the flow. In the limit $\beta \rightarrow 0$ the Laplace-Beltrami operator reduces to the usual Laplacian and hence we obtain

$$(\gamma_{\mu\nu}) = \begin{pmatrix} 1 + \beta \text{tr}((Y^{-1}Y_{x_1})^2) & \beta \text{tr}((Y^{-1}Y_{x_1})(Y^{-1}Y_{x_2})) & \beta \text{tr}((Y^{-1}Y_{x_1})(Y^{-1}Y_{x_3})) \\ \beta \text{tr}((Y^{-1}Y_{x_1})(Y^{-1}Y_{x_2})) & 1 + \beta \text{tr}((Y^{-1}Y_{x_2})^2) & \beta \text{tr}((Y^{-1}Y_{x_2})(Y^{-1}Y_{x_3})) \\ \beta \text{tr}((Y^{-1}Y_{x_1})(Y^{-1}Y_{x_3})) & \beta \text{tr}((Y^{-1}Y_{x_2})(Y^{-1}Y_{x_3})) & 1 + \beta \text{tr}((Y^{-1}Y_{x_3})^2) \end{pmatrix}. \tag{10}$$

5 Iwasawa Coordinate-Based Formalism

In Sect. 4.1 we have defined the induced metric. Now we would like to express this metric in terms of the coordinates on P_3 , explicitly.

The P_3 space is identified with an open cone in R^6 . The 6 different entries of the symmetric matrix may be identified with the Cartesian coordinates for R^6 . Therefore, it is straightforward to parameterize P_3 using these coordinates. However, taking the trace of $(Y^{-1}dY)^2$ where $Y = (y_{ij})_{1 \leq i, j \leq 3}$ and $dY = (dy_{ij})_{1 \leq i, j \leq 3}$, one gets long and cumbersome expressions with complicated terms involving the y_i 's in the denominator. Hence, the metric tensor is cumbersome and there are 78 Christoffel symbols associated with it. As a result, the numerical implementation is problematic.

An important issue in analysis on manifolds is the right choice of coordinate system for the problem. We may try to choose a different coordinate system in order to simplify the analytical calculations as well as the numerics.

Another candidate is the coordinate system associated with the polar decomposition. However, as we have mentioned in the introduction this coordinate system is problematic because of the non-uniqueness of the polar decomposition. Luckily, there is another set of coordinates called Iwasawa coordinates which corresponds to the Iwasawa decomposition of symmetric matrices. There are partial, as well as, full Iwasawa decompositions (Jorgenson and Lang 2005; Terras 1988). The full Iwasawa decomposition is unique and obtained by applying repeatedly the partial Iwasawa decomposition to the matrices $Y \in P_n$. Surprisingly, the Iwasawa coordinates turn out to be the natural parameterization on P_n from analytical as well as from numerical considerations. In this paper we use the term ‘‘Iwasawa decomposition’’ to describe the full Iwasawa decomposition.

5.1 The Iwasawa Decomposition

Iwasawa has proved that every connected semisimple Lie group G admits a unique representation as a product $G = KAN$ of an orthogonal subgroup K , an Abelian subgroup A and a nilpotent subgroup N (Rosenfeld 1997). In particular,

a linear flow. Small β values lead to a multi-channel total variation (TV) norm (Blomgren and Chan 1998) while large β values lead to potential surfaces (Sochen et al. 1998).

In matrix form the induced metric is given by

for every invertible real matrix of order n , $G = GL(n, \mathbb{R})$, K is the orthogonal group $K = O(n)$, A is a positive diagonal matrix and N is a strictly upper-triangular matrix with ones on its diagonal.

However, the space of SPD matrices, P_n , is not a Lie-group but it is identified with the quotient space $GL(n, \mathbb{R})/O(n)$. Therefore, any $Y \in P_n$ may be identified with an Iwasawa decomposition of the form $Y = AN$. In this case the operation between A and N is not an ordinary matrix multiplication. The homeomorphism $GL(n, \mathbb{R})/O(n) \rightarrow P_n$ is given by the operation $g^T g, \forall g \in GL(n, \mathbb{R})$. Hence, any $Y \in P_n$ may be decomposed uniquely as follows: $Y = N^T AN$ (Jorgenson and Lang 2005; Terras 1988).

Calculating the metric over P_n with respect to the Iwasawa coordinates is done by a substitution of the Iwasawa decomposition into $ds_{P_n}^2$. This yields

$$ds_{P_n}^2 = \text{tr}((Y^{-1}dY)^2) = \text{tr}((A^{-1}dA)^2) + 2\text{tr}(A^{-1}A[dNN^{-1}]), \tag{11}$$

where the second term plays an important role here. It is the source for the coupling between the Iwasawa coordinates.

Let us now derive the metric tensor for the spatial-feature manifold $\mathbb{R}^3 \times P_3$.

5.2 Diffusion Tensors

The diffusion process of water molecules in the brain may be modeled by a diffusion tensor. The diffusion tensor is a 3×3 SPD matrix. Let us write the Iwasawa decomposition for an SPD matrix:

$$Y = A[N] = N^T AN = \begin{pmatrix} 1 & 0 & 0 \\ w_4 & 1 & 0 \\ w_5 & w_6 & 1 \end{pmatrix} \begin{pmatrix} w_1 & 0 & 0 \\ 0 & w_2 & 0 \\ 0 & 0 & w_3 \end{pmatrix} \begin{pmatrix} 1 & w_4 & w_5 \\ 0 & 1 & w_6 \\ 0 & 0 & 1 \end{pmatrix}, \tag{12}$$

where in this case we identify w_1, \dots, w_6 with the Iwasawa coordinates.

The metric tensor for the Iwasawa decomposition is obtained by using (11)

$$(h_{ij}) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{w_1^2} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{w_2^2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{1}{w_3^2} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{2w_1(w_3+w_2w_6^2)}{w_2w_3} & -\frac{2w_1w_6}{w_3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\frac{2w_1w_6}{w_3} & \frac{2w_1}{w_3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \frac{2w_2}{w_3} \end{pmatrix}. \tag{13}$$

The positive-definiteness of Y implies the constraint $w_1, w_2, w_3 > 0$. Hence, the components of the metric tensor cannot be singular.

There are only 26 non-trivial Christoffel symbols associated with this metric tensor. With respect to this metric tensor, the induced metric is calculated easily using (8) and the following set of six coupled Beltrami equations for the Iwasawa coordinates is constructed:

$$\begin{aligned} X_t^i &= \Delta_\gamma X^i + \Gamma_{jk}^i \langle \nabla X^j, \nabla X^k \rangle_\gamma = 0, \\ X_{t=0}^i &= w_i(t=0), \quad i = 1, \dots, 6. \end{aligned} \tag{14}$$

5.3 DTI Data Fitting Term

The coupled equations in (6) correspond to the *smoothing* term. However, to keep the regularized data best fit the original one we add a physical data fitting term to the functional in (4). This results in the following cost functional:

$$\begin{aligned} L(X) &= \frac{\alpha}{2} \int_\Sigma d^3x \sum_{k=1}^N \left(e^{-bq_k^T Dq_k} - \hat{E}(q_k) \right)^2 \\ &\quad + \frac{1}{2} \int_\Sigma d^3x \sqrt{\gamma} \gamma^{\mu\nu} \partial_\mu X^i \partial_\nu X^j h_{ij}(X), \end{aligned} \tag{15}$$

where $i = 1, \dots, 6$. The term $e^{-bq_k^T Dq_k}$ is the signal attenuation model (Basser et al. 1994) where D is the diffusion tensor, q_k is the k 'th applied magnetic field gradient direction and b is a parameter of the measurement. The measured (noisy) signal is denoted by $\hat{E}(q_k)$.

The variation of this functional with respect to the Iwasawa coordinates yields the following equations of motion

$$\begin{aligned} -\alpha b \frac{1}{\sqrt{\gamma}} \sum_{k=1}^N \left(e^{-bq_k^T Dq_k} - \hat{E}(q_k) \right) e^{-bq_k^T Dq_k} \left(q_k^T \frac{\partial D}{\partial X^i} q_k \right) \\ + \frac{1}{\sqrt{\gamma}} \partial_\mu \left(\sqrt{\gamma} \gamma^{\mu\nu} \partial_\nu X^i \right) + \Gamma_{jk}^i \gamma^{\mu\nu} \partial_\mu X^j \partial_\nu X^k = 0. \end{aligned} \tag{16}$$

Then, the gradient descent equations are given by

$$\begin{aligned} X_t^i &= -\alpha b \frac{1}{\sqrt{\gamma}} \sum_{k=1}^N \left(e^{-bq_k^T Dq_k} - \hat{E}(q_k) \right) \\ &\quad \times e^{-bq_k^T Dq_k} \left(q_k^T \frac{\partial D}{\partial X^i} q_k \right) \\ &\quad + \frac{1}{\sqrt{\gamma}} \partial_\mu \left(\sqrt{\gamma} \gamma^{\mu\nu} \partial_\nu X^i \right) + \Gamma_{jk}^i \gamma^{\mu\nu} \partial_\mu X^j \partial_\nu X^k. \end{aligned} \tag{17}$$

The derivatives of D with respect to the Iwasawa coordinates are calculated as follows:

$$\begin{aligned} D &= \begin{pmatrix} 1 & 0 & 0 \\ w_4 & 1 & 0 \\ w_5 & w_6 & 1 \end{pmatrix} \begin{pmatrix} w_1 & 0 & 0 \\ 0 & w_2 & 0 \\ 0 & 0 & w_3 \end{pmatrix} \begin{pmatrix} 1 & w_4 & w_5 \\ 0 & 1 & w_6 \\ 0 & 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} w_1 & w_1w_4 & w_1w_5 \\ w_1w_4 & w_2 + w_1w_4^2 & w_1w_4w_5 + w_2w_6 \\ w_1w_5 & w_1w_4w_5 + w_2w_6 & w_3 + w_1w_5^2 + w_2w_6^2 \end{pmatrix}. \end{aligned} \tag{18}$$

Thus, for example

$$\frac{\partial D}{\partial x_1} = \begin{pmatrix} 1 & x_4 & x_5 \\ x_4 & x_4^2 & x_4x_5 \\ x_5 & x_4x_5 & x_5^2 \end{pmatrix}. \tag{19}$$

6 Experimental Results

The regularization framework was tested on in vivo, human noisy data that was acquired on a GE Signa 3T. A DW-EPI sequence was used with the following parameters: $\Delta/\delta = 31/25$ ms, a 22 cm FOV, matrix size of 128×128 and 1.7 mm slice thickness with 72 slices covering the entire brain. The TE was 88 ms and the gradient strength was 4 G/cm, resulting in a b value of 1000 s/mm² measured in 33 non-collinear gradient directions. In addition, a single non weighted b_0 image was acquired. The sequence was

repeated 4 times. In order to avoid intrinsic pulsative brain motion artifacts, the sequence was gated to the cardiac cycle with effective TR of 30 R–R intervals. The MRI protocol was approved by the local IRB committee, and informed consent was obtained from a healthy volunteer.

6.1 Data Analysis

The collected volumes were first corrected for patient motion by rigid-body spatial transformation (using SPM2) accompanied with gradient-orientation compensation. A noisy tensor field was then generated by applying DTI on a single repetition of the DW-EPI sequence. Then, it was regularized by the proposed framework (50 iterations with fidelity weight $\alpha = 1$ and $\beta = 1$). Finally, the regularized tensors dataset was compared with the initial noisy dataset. Also, it was compared with the mean of the four repetitions dataset. In order to use the noisy tensor field as an initial guess for

the Beltrami framework minimization, the attenuation signal of any negative eigenvalued voxel was replaced with the mean signal of its positive neighbors (if there were any, otherwise it was omitted). Visualization of the tensor fields was generated by spectral decomposition of the three different tensor fields for the rendering of fractional anisotropy (FA) maps (Basser and Pierpaoli 1996), RGB color schemes (Pajevic and Pierpaoli 1999), and three-dimensional tractography objects. The latter were acquired by brute-force streamline tractography (Mori et al. 2005) overlaid on SPGR anatomical images, co-registered with a b_0 diffusion image. A FA threshold of 0.2 was used for fiber initialization and termination.

6.2 Comparison between the Three Cases

The high voxel resolution yielded low-grade tensor images for the single repetition dataset (Fig. 2, left column): The noise is easily visible in all parts of the image, including




Fig. 2 FA images. Presented are mid-Axial (*top row*), mid-Sagittal (*middle row*) and mid-Coronal (*bottom row*) slices. The Beltrami regularized maps (*middle column*) show significantly lower noise level

compared to the original noisy images (*left column*). The regularized images resemble the high SNR images obtained by 4-repetitions (*right column*)

in deep white matter structures such as the genu of the corpus callosum, and especially in lateral brain areas where it is difficult to recognize small fiber bundles. This noise effect is less noticeable in the 4-repetitions dataset (Fig. 2, right column). In this case, finer fibers at the intersection with gray matter become visible. Therefore, we may expect that the outcome of the regularization process are images that are more similar to the 4-repetitions images. This is indeed the case (Fig. 2, middle column): The noise level of the regularized images have been improved significantly comparing to the original, single repetition images. Anatomically expected homogeneous fiber areas appear smoother with much less gray level value changes, and the images become similar to the 4-repetitions images. While the FA maps demonstrate that the eigenvalues have been regularized, we can further see in the color-coded schemes (Fig. 3) that the Beltrami regularization has also regularized the principal orientations of the tensors. The colors better represent different fiber bundles, and demonstrate again that the regularized images are

not only smoother but also preserve the edges between these bundles.

Very similar structures appear in the 4-repetitions color schemes, reassuring the assumption of piece-wise smooth organization of fibers, and supporting the regularized results. The difference between the regularized and the 4-repetitions images is the noise level where the regularized image seems to have a better one. However, the image edges of the 4-repetitions images are sharper.

The anisotropic regularization of the eigenvalues and the orientations can be presented in one image via ellipsoid representation. In Fig. 4 this is demonstrated for the corpus callosum fiber.

In order to quantify the regularization results we measure the similarity between the images. This is done via a *global* distance measure between tensors. There are some candidates for this measure. We have chosen the natural one which is the Riemannian distance in (1). Using this distance measure, we can compare the distance of the regularized image from the 4-repetitions image with the distance of the one




Fig. 3 Colorcode images. The same slices as in Fig. 2 coded by the tensor's principal orientation. Orientations were regularized as well, while important edges between different fiber bundles were preserved

repetition image from the 4-repetitions image. In Table 1 the mean distance and the standard deviation for the tensors at different ROIs are presented (see Fig. 5 for the ROI map). These differences were found statistically significant for all ROIs, with the exception of the corpus callosum ROI. In order to explain this exception, we should notice that the mean distance for this ROI had much larger values than for other ROIs. This distance, however, cannot be seen in the visual quality of this ROI compared with the other ROIs. This effect may be explained by the fact that the corpus callosum is a very dense fiber that is best fitted with extremely anisotropic tensors, where the smallest eigenvalues tend to zero. The Riemannian distance is less stable for this kind of

tensors, pulling the distance measures towards infinity, and therefore the variance is increased.

Further visualization and validation of the regularized results is obtained from tractography images. Identical ROIs were chosen for the three datasets to acquire streamline tractography of the cingulum. The ROIs were placed on two coronal slices, posterior and anterior to the mid-coronal slice, on anatomical landmarks identified with the cingulum. The streamlines generated had to pass through both ROIs. The cingulum is a fiber bundle passing between parts of the limbic system and is wrapped around the corpus callosum. It connects the anterior part of the cingulate gyrus (aCG), which is in proximity to the genu of the corpus callosum (gcc), passing superior to the body of the corpus callosum (bcc) and ends at the posterior part of the cingulate gyrus (pCG) next to the splenium of the corpus callosum (scc). From the posterior part of the cingulate gyrus it projects to the amygdala and hippocampus complex.

In Fig. 6 we can see the trajectories obtained for the three datasets. The three images were able to delineate the complete tract from the anterior to the posterior cingulate gyrus. Interestingly, the noisy dataset had also delineated the projection to the amygdala and hippocampus, whereas the 4-repetitions delineation was terminated at the posterior cingulate gyrus. As for the regularized dataset, it had also obtained the complete expected trajectory. Further comparing the three images reveals that the regularized dataset provided a smooth and robust shape of fibers, where the noisy dataset presents many spurious fibers that seem as noise. Observing




Fig. 4 First row from left to right: The original DTI axial slice, regularized slice using the full Iwasawa coordinates. The ellipsoids are color coded according to the directions of the principal components in the RGB plane. Second row: Enlargement of the marked area before the regularization (left) and after the regularization (right)

Fig. 5 ROIs map. The 5 ROIs selected for distance measurements, labeled by their ROI id number




Table 1 Similarity measures for the 4-repetitions image, the noisy image and the regularized image. All ROIs show that the distance of the regularized image from the 4-repetitions image is significantly lower

than the distance of the noisy image from the 4-repetitions image with exception of ROI no. 2

ROI id	ROI name	N	Noisy image	Regularized image	Significance
1	Internal capsule	37	0.3229 ± 0.2843	0.1557 ± 0.1286	yes
2	Corpus callosum	109	2.4306 ± 3.8796	2.2169 ± 3.7279	no
3	Lateral fibers	560	0.2004 ± 0.3161	0.1629 ± 0.3014	yes
4	Frontal fibers	1139	0.4591 ± 1.1183	0.3855 ± 1.0952	yes
5	Entire slice	5256	0.3892 ± 1.2077	0.3309 ± 1.0780	yes

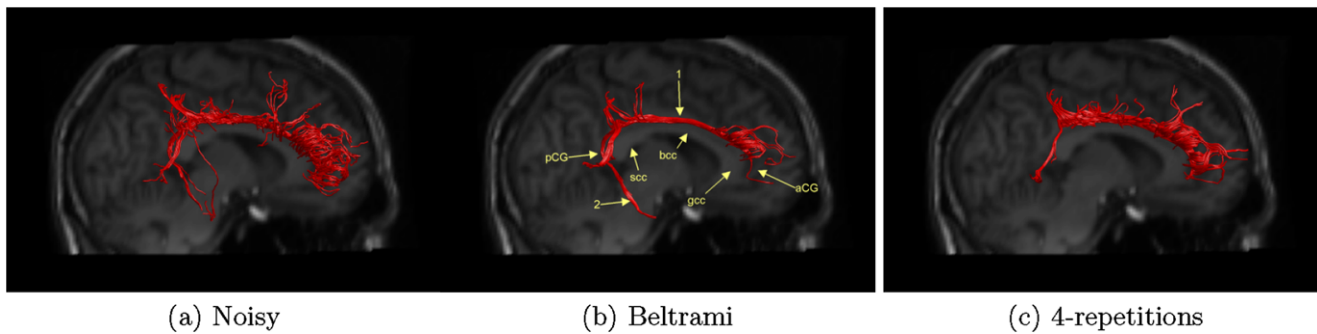


Fig. 6 The cingulum. In the three images (the noisy, regularized and the 4-repetitions) the cingulum between the aCG to the pCG has been delineated. In the regularized image the projection to the amygdala and

hippocampus was also obtained (marked by 2). The regularized tract has much less spurious fibers and is smoother than the noisy tract




Fig. 7 Cingulum zoomed. Area 1 in Fig. 6b is zoomed. This reveals that the regularized fiber orientations change smoothly along neighboring voxels, unlike sharp edges and discontinuities in the noisy dataset

a zoomed perspective of the same images (Fig. 7), reveals much finer differences between the images. Here it is easy to notice that the regularized dataset provided much more continuous and smooth trajectories than in the noisy dataset. In the noisy dataset the effect of noise is seen as sharp turns along the fiber as well as discontinuities along the tract. Comparing the regularized result with the 4-repetitions result shows that the regularized trajectories seems smoother than the 4-repetitions ones.

Since in most of its trajectory the cingulum is very close to the corpus callosum, small orientational noise is expected to cause deviations that might cross to the neighboring bundle and follow it to the other hemisphere. The two-ROIs method used here is supposed to eliminate most of these fibers. However, it is limited to the area that is between these ROIs. In areas posterior and anterior to the chosen ROIs there is still the possibility of noise to causing deviation on the corpus callosum. Observing the trajectory from a different viewing angle (Fig. 8) demonstrates clearly some of the false positives fibers that cross to the other hemisphere. This visualization shows us that indeed in the noisy dataset there are many fibers that deviate to the corpus callosum tract, where both in the regularized and the 4-repetitions datasets the number of such fibers is dramatically decreased.

As a conclusion from the tractography of the cingulum fiber bundle, the regularization process provides smoother, more robust, and less false positive fibers than in the original noisy dataset. Also, it preserves the anatomical shape of the fiber and increases its separation from other fiber bundles.

7 Summary

In this paper we have presented a novel geometric framework for regularization of data in P_n , the space of SPD matrices of order n . The basis of this framework is the description of P_n as a Riemannian manifold in terms of *local* coordinates on the manifold. We have shown that a suitable coordinate system to parameterize P_n are the *Iwasawa* coordinates. This coordinate system has advantages over other coordinate systems used by others (in the context of DTI data processing). Describing the P_n manifold via the Iwasawa coordinates simplifies the analytical as well as the numerical calculations. As a result, the proposed algorithm converges very fast and the results are satisfactory. We have shown that, using calculus of variations on P_n via a functional over sections, a set of coupled Beltrami equations for the local coordinates on P_n is derived. Regularization of DTI data is




Fig. 8 Deviation to the corpus callosum. Since the cingulum is very close to the corpus callosum, small deviations caused by noise may switch the trajectory to the corpus callosum trajectory. Any fiber that crosses to the opposite hemisphere is false positive. The regularized

image reduced dramatically the number of false positive fibers obtained in the original data. Similar results were obtained in the 4-repetitions data

performed by extracting the Iwasawa coordinates from the SPD matrices and then solving numerically a set of six coupled Beltrami equations with respect to these coordinates. We have demonstrated the efficiency of this framework on real DTI volume.

We believe that the local coordinate approach we presented here incorporated with the Iwasawa coordinates may be extended to yield fast and accurate algorithms for other applications of DTI processing (e.g., segmentation, registration, etc.). These research directions are under current study.

References

- Basser, P. J., & Pierpaoli, C. (1996). Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, *111*, 209–219.
- Basser, P. J., Mattiello, J., & LeBihan, D. (1994). MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, *66*, 259–267.
- Batchelor, P. G., Moakher, M., Atkinson, D., Calamante, F., & Connelly, A. (2005). A rigorous framework for diffusion tensor calculus. *Magnetic Resonance in Medicine*, *53*, 221–225.
- Blomgren, P., & Chan, T. F. (1998). Total variation methods for restoration of vector valued images. *IEEE Transactions on Image Processing*, *7*, 304–309.
- Chefd'hotel, C., Tschumperlé, D., Deriche, R., & Faugeras, O. (2004). Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision*, *20*, 147–162.
- Coulon, O., Alexander, D. C., & Arridge, S. R. (2001). A regularization scheme for diffusion tensor magnetic resonance images. In *IPMI '01: Proceedings of the 17th international conference on information processing in medical imaging* (pp. 92–105). London, UK: Berlin: Springer.
- Feddern, C., Weickert, J., Burgeth, B., & Welk, M. (2006). Curvature-driven PDE methods for matrix-valued images. *International Journal of Computer Vision*, *69*, 91–103.
- Fillard, P., Arsigny, V., Pennec, X., & Ayache, N. (2007). Clinical DT-MRI estimation, smoothing and fiber tracking with Log-Euclidean metrics. *IEEE Transactions on Medical Imaging*, *26*(11), 1472–1482.
- Fletcher, P. T., & Joshi, S. (2007). Riemannian geometry for the statistical analysis of diffusion tensor data. *Signal Processing*, *87*, 250–262.
- Gur, Y., & Sochen, N. (2007). Coordinate-free diffusion over compact Lie groups. In *Lecture notes in computer science: Vol. 4485. Proceedings of the 1st international conference on scale-space and variational methods* (pp. 580–591). Ischia Island, Italy, May 2007. Berlin: Springer.
- Jorgenson, J., & Lang, S. (2005). *Pos_n(R) and Eisenstein series. Lecture notes in mathematics*. Berlin: Springer.
- Jost, J. (2001). *Riemannian geometry and geometric analysis*. New York: Springer.
- Kimmel, R., Sochen, N., & Malladi, R. (1997). From high energy physics to low level vision. In *SCALE-SPACE '97: Proceedings of the first international conference on scale-space theory in computer vision* (pp. 236–247). London, UK: Berlin: Springer.
- Lang, S. (1999). *Fundamentals of differential geometry*. New York: Springer.
- Moakher, M. (2005). A differential geometric approach to the geometric mean of symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications*, *26*(3), 735–747.
- Moakher, M., & Zerai, M. (2007). Riemannian curvature-driven flows for tensor-valued data. In *Lecture notes in computer science: Vol. 4485. In proceedings of the 1st international conference on scale-space and variational methods* (pp. 592–602). Ischia Island, Italy: Berlin: Springer.
- Mori, S., Wakana, S., van Zijl, P. C. M., & Nagae-Poetscher, L. M. (2005). *MRI atlas of human white matter*. Amsterdam: Elsevier.
- Pajevic, S., & Pierpaoli, C. (1999). Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, *42*, 526–540.
- Pasternak, O., Sochen, N., Intrator, N., & Assaf, Y. (2008). Variational multiple-tensors fitting of fiber-ambiguous DW-MRI voxels. *MRI*, *26*(8), 1133–1144.
- Pennec, X., Fillard, P., & Ayache, N. (2006). A Riemannian framework for tensor computing. *International Journal of Computer Vision*, *66*, 41–66.
- Pierpaoli, C., Jezzard, P., Basser, P. J., Barnett, A., & Chiro, G. D. (1996). Diffusion tensor MRI of the human brain. *Radiology*, *201*, 637–648.
- Polyakov, A. M. (1981). Quantum geometry of bosonic strings. *Phys. Lett. B*, *103*, 207–210.
- Rosenfeld, B. (1997). *The geometry of Lie groups*. Dordrecht: Kluwer Academic.
- Shafirir, D., Sochen, N., & Deriche, R. (2005). Regularization of mappings between implicit manifolds of arbitrary dimension and codimension. In *Proceedings of the 3rd IEEE workshop on variational, geometric and level-set methods (VLSM) in computer vision*.

- Sochen, N., Kimmel, R., & Malladi, R. (1998). A general framework for low level vision. *IEEE Transactions in Image Processing*, 7, 310–318. Special Issue on Geometry Driven Diffusion.
- Terras, A. (1988). *Harmonic analysis on symmetric spaces and applications* (Vol. 2). New York: Springer.
- Tschumperlé, D., & Deriche, R. (2002). Orthonormal vector sets regularization with PDE's and applications. *International Journal of Computer Vision*, 50(3), 237–252.
- Wang, Z., Vemuri, B. C., Chen, Y., & Mareci, T. H. (2004). A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging*, 23, 930–939.
- Weickert, J. (1998). *Anisotropic diffusion in image processing*. Stuttgart: Teubner.
- Weickert, J., & Hagen, H. (Eds.) (2005). *Visualization and processing of tensor fields*. Berlin: Springer.