

Constrained Flows of Matrix-Valued Functions: Application to Diffusion Tensor Regularization

C. Chefd'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras

INRIA Sophia-Antipolis, France
`{cchefd,dtschump,der,faugeras}@sophia.inria.fr`

Abstract. Nonlinear partial differential equations (PDE) are now widely used to regularize images. They allow to eliminate noise and artifacts while preserving large global features, such as object contours. In this context, we propose a geometric framework to design PDE flows acting on constrained datasets. We focus our interest on flows of matrix-valued functions undergoing orthogonal and spectral constraints. The corresponding evolution PDE's are found by minimization of cost functionals, and depend on the natural metrics of the underlying constrained manifolds (viewed as Lie groups or homogeneous spaces). Suitable numerical schemes that fit the constraints are also presented. We illustrate this theoretical framework through a recent and challenging problem in medical imaging: the regularization of diffusion tensor volumes (DT-MRI).

Introduction

In the last decade, variational methods and nonlinear PDE's have been widely used to tackle computer vision problems, such as image restoration [19, 24, 26, 35], segmentation [22], stereo-based 3D reconstruction [12], image inpainting [4, 6], or image matching and optical flow estimation [1, 2, 11] (among other examples). Solutions to these problems, whether they are curves, surfaces, or images, are obtained by continuously deforming an initial estimation through a flow defined by a PDE. The corresponding evolution equations derive either from simple local heuristics or from cost functional minimizations. In the context of image and data restoration, which is considered in this paper, the idea is to achieve a selective smoothing that removes noise while preserving large global features, such as object contours (discontinuities of the signal). For this purpose, one generally uses anisotropic nonlinear diffusion PDE's, which often derive from variational principles (such as the ϕ -function formulation [7, 19, 24, 25] which has proven its efficiency for gray-valued image restoration).

Generalizing these algorithms to multi-valued datasets is now attracting a growing interest. In effect, the increase of computer performances has made possible the implementation of PDE's acting on large fields of vectors or matrices. Recent works point toward this direction, with the definition of methods for color image restoration [26, 28, 31, 35], or direction field regularization [5, 17, 23, 29, 33]. A flow on multi-valued data is generally not a straightforward generalization

of its scalar counterpart, and new theoretical developments are involved. For instance, in the case of direction field restoration, vectors are constrained to have a unit norm. This yields a significant modification of the corresponding diffusion PDE's by introducing a coupling term between vector components.

The aim of this paper is to provide some formal and numerical tools to deal efficiently with this type of problems, with a particular emphasis on flows of matrix-valued functions undergoing orthogonal and spectral constraints. Our final objective concerns the restoration of diffusion tensor images of the human brain (DT-MRI) [21], modeled as fields of symmetric positive definite matrices (i.e. constrained to have positive eigenvalues). We start from the fact that most evolution PDE's can be viewed as ordinary differential equations (ODE) on a suitable function space. When nonlinear constraints are involved, the corresponding flows belong to an infinite dimensional submanifold¹ of the initial search space. Usually, this geometric viewpoint is hidden behind the use of Lagrange multipliers in constrained variational principles. It also appears implicitly in the nonlinear heat equation, borrowed from harmonic map theory, which has been recently proposed to regularize direction fields [5, 29]. The interest of a fully geometric interpretation of the suitable function space is twofold. First, it provides a simple and unified framework to build constrained flows for matrix-valued functions, whether or not they derive from variational principles. Second, it naturally yields, through the use of exponential maps, suitable numerical schemes that are also constraint preserving.

This approach is presented as follows. Section 1 introduces a geometric setting to define constrained flows and their proper numerical implementations. Section 2 discusses the application of these methods to the definition of three different regularization techniques for DT-MRI datasets. Finally, section 3 presents results of numerical experiments on a real DT-MRI volume of the brain, leading to the construction of regular tissue fiber maps in the white matter.

1 Geometric Interpretation of Constraint Preserving Flows

1.1 Flows on Manifold

Let us consider the smooth trajectory $t \mapsto p(t)$ of a point moving on an arbitrary manifold \mathcal{M} . Any such trajectory can be defined, at least locally, as the integral curve of a smooth vector field V , and satisfies the evolution equation

$$\frac{\partial p(t)}{\partial t} = V(p(t)) \quad (\text{where } p(0) \in \mathcal{M}). \quad (1)$$

Recall that since V is a vector field, we have $\forall p \in \mathcal{M}, V(p) \in T_p\mathcal{M}$, where $T_p\mathcal{M}$ denotes *the tangent space of \mathcal{M} at p* . An example of such an integral curve, induced by a vector field on a sphere, is illustrated in Fig. 1a.

¹ Throughout this paper, we borrow some tools from differential geometry and Lie group theory. We refer to [18, 20] for the corresponding background.

Despite its simplicity, eq. (1) contains the key elements to build all the constrained flows presented below. Let us assume that \mathcal{M} is a submanifold corresponding to some nonlinear constraints on its embedding space \mathcal{P} . Then, we can say that a flow is *constraint preserving* if it is tangent to \mathcal{M} . Its form follows readily from the expression of the tangent space at any point p of the constrained set. This argument is valid for both finite and infinite dimensional manifolds. For our application, we will consider manifolds of the form $\mathcal{M} = \mathcal{F}(\Omega, \mathcal{N})$, where $\mathcal{F}(\Omega, \mathcal{N})$ is a set of functions $X : \Omega \rightarrow \mathcal{N}$, defined on an open subset Ω of the Euclidean space, and which take values in a constrained matrix manifold \mathcal{N} . This infinite dimensional manifold of mappings could be defined on a suitable Banach or Hilbert space [16, 18, 20]. For a given point $X \in \mathcal{F}(\Omega, \mathcal{N})$, its tangent space is identified to the set of functions:

$$T_X \mathcal{F}(\Omega, \mathcal{N}) = \{x \mapsto V(x), \text{ such that } V(x) \in T_{X(x)} \mathcal{N}\},$$

where $T_{X(x)} \mathcal{N}$ is the tangent space to \mathcal{N} at $X(x)$.

In the following, \mathcal{M} is viewed as a submanifold of the linear space $\mathcal{P} = \mathcal{F}(\Omega, \mathbb{R}^{n \times n})$, where $\mathbb{R}^{n \times n}$ denotes the set of $n \times n$ real matrices². Note that in the definition of the tangent space, we assumed that $\mathcal{F}(\Omega, \mathcal{N})$ inherits its constrained geometry directly from its codomain \mathcal{N} . We do not consider additional structural constraints, such as invertibility, which would yield more complex families of mappings (see for instance the groups of diffeomorphisms used in the image matching context [30]). In this setting, an evolution equation on $\mathcal{F}(\Omega, \mathcal{N})$ must satisfy:

$$\forall t > 0, \forall x \in \Omega, \quad \begin{cases} \frac{\partial X(x)}{\partial t} = V(x), & \text{with } V(x) \in T_{X(x)} \mathcal{N}, \\ X_0 = X_{(t=0)} \in \mathcal{F}(\Omega, \mathcal{N}). \end{cases} \quad (2)$$

In order to specialize this equation to the practical cases studied in this paper, we can also notice that a wide range of orthogonal and spectral constraints on matrices define *Lie groups* and *homogeneous spaces*. They are submanifolds of $\mathbb{R}^{n \times n}$ which present some nice algebraic properties. In particular, the tangent space at any given point of a matrix Lie group is easily defined in terms of its tangent space at identity (its Lie algebra) [18]. This is particularly important since equation (2), together with the expression of the tangent space at any point on \mathcal{N} , will yield directly the constrained flows we are looking for. Regardless of the type of constraints discussed in the section 2, we will describe each point on \mathcal{N} in terms of its extrinsic coordinates in $\mathbb{R}^{n \times n}$. The main reason is that a suitable parametrization of a constrained set is generally difficult, since several charts are often necessary to cover the underlying manifold. The constraint set may even have several connected components (the group of orthogonal matrices, for

² We naturally extend to $\mathcal{F}(\Omega, \mathbb{R}^{n \times n})$ all the operators existing on $\mathbb{R}^{n \times n}$. Let $X, Y, Z \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n})$, the expression $Z = XY$ corresponds to the equality $Z(x) = X(x)Y(x)$, $\forall x \in \Omega$ (product of matrices).

instance, has two connected components: the sets of rotations and anti-rotations (we call “anti-rotation” a reflection combined with a rotation)).

1.2 Gradient Flows and Cost Functionals

One way of building a constraint preserving flow (2) for a specific purpose is to minimize a cost functional f . The corresponding evolution equation is usually defined as a gradient flow that follows the direction of steepest descent toward the nearest critical point. To build such a flow, \mathcal{M} is endowed with a *Riemannian metric* $\langle \cdot, \cdot \rangle_p$ which defines for all $p \in \mathcal{M}$ a scalar product between any two vectors of the tangent space $T_p\mathcal{M}$. Recall that the gradient is then defined as the unique vector field ∇f such that

$$\forall v \in T_p\mathcal{M}, \quad df_p(v) = \langle \nabla f(p), v \rangle_p,$$

where $df_p(v)$ denotes the *first variation*³ of f at p in the direction v . The corresponding minimizing flow is given by

$$\frac{\partial p(t)}{\partial t} = -\nabla f(p(t)) \quad (p(0) \in \mathcal{M}).$$

When \mathcal{M} is a linear function space equipped with the L^2 -metric defined by $\langle v, w \rangle = \int_{\Omega} v(x) \cdot w(x) dx$, the equation $\nabla f(p) = 0$ corresponds to the usual Euler-Lagrange equation of variational calculus.

In our framework, we choose to endow $\mathcal{M} = \mathcal{F}(\Omega, \mathcal{N})$ with the metric $\langle U, V \rangle_X = \int_{\Omega} \langle U(x), V(x) \rangle_{X(x)} dx$, where $\langle \cdot, \cdot \rangle_p^{\mathcal{N}}$ denotes a suitable Riemannian metric on \mathcal{N} , usually an *invariant* or a *normal metric* in the case of Lie groups and homogeneous spaces [15]. In practice, when the L^2 -gradient is known for an arbitrary cost functional f on $\mathcal{F}(\Omega, \mathbb{R}^{n \times n})$, we will be able to reuse its expression to define the gradient of the restriction of f to $\mathcal{F}(\Omega, \mathcal{N})$. Up to a change of metric, the corresponding transformation is an orthogonal projection on the suitable tangent space.

1.3 Exponential Maps and Numerical Schemes

Standard numerical methods for differential equations are usually not suitable for constrained flows. Note that with a simple Euler method, the numerical integration of equation (1) yields the explicit scheme

$$p_{(t+dt)} = p_{(t)} + V(p_{(t)}) dt \quad \text{where} \quad p_0 = p_{(t=0)} \in \mathcal{M}, \quad (3)$$

where $dt > 0$ corresponds to the integration step. Since we are using extrinsic coordinates, there is a risk of stepping out from \mathcal{M} for each (non infinitesimal) displacement. In fact, by writing (3) we implicitly assumed that the integral

³ For an arbitrary curve p_{ϵ} on \mathcal{M} , such that $p_0 = p$ and $\frac{\partial p_{\epsilon}}{\partial \epsilon} \Big|_{\epsilon=0} = v$, one computes the first variation of f from the definition $df_p(v) = \frac{\partial f(p_{\epsilon})}{\partial \epsilon} \Big|_{\epsilon=0}$.




Fig. 1. (a) Integral curve of a vector field on a manifold \mathcal{M} , (b) Geodesic steps on a Riemannian manifold.

curve lies in a linear space. In our situation, we must adjust the integration method to accommodate the true constrained geometry of \mathcal{M} . The first idea is to project after each step the point p_{t+dt} on the underlying manifold. We propose here to avoid this post-processing using a geometric integration step relying on *exponential maps*, following directly a geodesic path between each iteration. This technique is borrowed from existing numerical methods on Riemannian manifolds [15, 27]. It consists of replacing (3) by

$$p_{t+dt} = \exp_{p_t}(dt V(p_t)) \quad p_0 \in \mathcal{M}, \quad (4)$$

where \exp_p denotes the exponential map at $p \in \mathcal{M}$. Recall that for a given point p on a (finite dimensional) manifold \mathcal{M} , and a vector $v \in Tp\mathcal{M}$, there exists a unique geodesic $t \mapsto \gamma_v(t)$ such that

$$\gamma_v(0) = p \quad \text{and} \quad \gamma'_v(0) = v.$$

The exponential map $\exp_p : T_p\mathcal{M} \rightarrow \mathcal{M}$ is then defined by $\exp_p(v) = \gamma_v(1)$ ($\forall v$ such that $\gamma_v(1)$ exists). For dt sufficiently small, equation (4) corresponds to a geodesic step of length

$$\|dtV(p_t)\|_{p_t} = \langle dtV(p_t), dtV(p_t) \rangle_{p_t}^{1/2},$$

in the direction of $V(p_t)$ (see Fig. 1b).

In general, the equation of a geodesic satisfies a second order differential equation. In this paper, we use the fact that for simple matrix manifolds, closed form solutions are often available (and thus the exponential map). Note that on linear manifolds endowed with the Euclidean metric, geodesics are straight lines. For instance, the tangent space at X in $\mathbb{R}^{n \times n}$ is defined by $T_X\mathbb{R}^{n \times n} \simeq \mathbb{R}^{n \times n}$ and we have $\forall X, V \in \mathbb{R}^{n \times n}$, $\exp_X V = X + V$. Then, equation (4) reduces to the original Euler integration method.

In our framework, these ideas extend naturally to the definition of a numerical integration scheme on $\mathcal{F}(\Omega, \mathcal{N})$. Given $\exp_p : T_p\mathcal{N} \rightarrow \mathcal{N}$ the exponential map on

\mathcal{N} , we introduce on $\mathcal{F}(\Omega, \mathcal{N})$ a step-forward operator $\mathbf{exp}_X V : T_X \mathcal{F}(\Omega, \mathcal{N}) \mapsto \mathcal{F}(\Omega, \mathcal{N})$ defined by

$$\forall x \in \Omega, x \mapsto (\mathbf{exp}_X V)(x) = \exp_{X(x)} V(x).$$

One can easily check that $\mathbf{exp}_X V$ yields a consistent numerical integrator (in time) for the continuous equation (2), such that

$$X_{(t+dt)}(x) = \left(\mathbf{exp}_{X_{(t)}} dt V_{(t)} \right)(x) = \exp_{X_{(t)}(x)} dt V_{(t)}(x), \quad (5)$$

with $V_{(t)}(x) \in T_{X_{(t)}(x)} \mathcal{N}$ and $X_0 \in \mathcal{F}(\Omega, \mathcal{N})$.

Regarding the discretization in space of $V_{(t)}$ (on Ω), finite difference schemes developed for unconstrained equations still apply. Actually, constraints on the codomain \mathcal{N} of $\mathcal{F}(\Omega, \mathcal{N})$ only change the time discretization scheme. If we had to deal with functions $\Omega \rightarrow \mathbb{R}^{n \times n}$ whose spatial domain Ω was an arbitrary manifold (for instance an implicit surface defined by a level set function), we could combine this technique with the framework introduced in [3].

2 Application of Constraint Preserving Flows to Diffusion Tensor Regularization

2.1 Context and Notations

We propose now to use the previous tools to address the problem of diffusion tensor regularization. Let $P(n)$ and $O(n)$ denote respectively the sets of symmetric positive definite, and orthogonal real matrices of size $n \times n$. Recall that

$$\begin{aligned} \forall X \in P(n), \quad X_{i,j} &= X_{j,i} \quad \text{and} \quad \forall a \in \mathbb{R}^n \setminus \{0\}, \quad a^T X a > 0, \quad (6) \\ \forall X \in O(n), \quad X X^T &= X^T X = I, \quad (I \text{ is the identity matrix}). \quad (7) \end{aligned}$$

We are mainly interested in regularizing fields $X : \Omega \rightarrow P(3)$ of diffusion tensors coming from DT-MRI imaging. This recent and non-invasive 3D medical imaging modality consists in measuring the water molecule motion in the tissue fibers, using magnetic resonance techniques. Each voxel $X(x) \in P(3)$ of the acquired image X is defined by a symmetric and positive definite 3×3 matrix that defines the local water molecule motion (i.e. the fiber orientations) [9, 10, 21, 14, 34]. These fiber structures are not explicitly given by the matrices $X(x)$, but can be retrieved by spectral decomposition $X(x) = U(x)D(x)U(x)^T$, where $U : \Omega \rightarrow O(3)$ is the field of 3×3 *orthogonal matrices* composed of the unit eigenvectors of X , and D is the field of the 3×3 matrices $D(x) = \text{diag}(\lambda_1, \lambda_2, \lambda_3)$ corresponding to its *positive eigenvalues*. Actually, the spectral decomposition splits the diffusion tensor field X into an *orientation feature* U (main directions of the water motion) and a *diffusivity feature* D (velocity of the water motion). A natural representation of X is then produced using ellipsoids whose axes and radii are respectively given by the eigenvectors in U and the corresponding




Fig. 2. View of a 3D diffusion tensor field $X : \Omega \rightarrow \mathcal{P}(3)$.

positive eigenvalues $\lambda_1, \lambda_2, \lambda_3$ (see Fig. 2).

In the following, the mathematical reasoning developed in section 1 is applied through a progressive analysis of the diffusion tensor regularization problem. We refer to the research report version of this paper [8] for a detailed derivation of the results presented below. Note that this framework allows us to generalize previous regularization methods acting directly on the matrix representation of DT-MRI data [9, 10, 32].

2.2 Symmetric Positive Definite (SPD) Constraint

By analogy with scalar image regularization methods, the first idea to restore a DT-MRI field is to find a matrix-valued flow minimizing a regularization functional f ,

$$\min_{X \in \mathcal{F}(\Omega, \mathcal{P}(3))} f(X) = \int_{\Omega} \phi(\|\nabla_{\Omega} X\|) + \frac{\alpha}{2} (X - X_0)^2 d\Omega, \quad (8)$$

(where $\|\nabla_{\Omega} X\| = \sqrt{\sum_{i,j} \|\nabla_{\Omega} X_{i,j}\|^2}$ and $\nabla_{\Omega} = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z})^T$ is the *spatial gradient* in Ω), while preserving the symmetric positive definite constraint (6). The fixed parameter $\alpha > 0$ prevents the expected solution from being too different from the given noisy field X_0 , while $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is an *increasing function* which controls the regularization behavior (this ϕ -function formulation was first introduced in the context of scalar image regularization with edge preservation [7, 19, 24]). Note that the functional (8) is invariant by matrix transposition and thus takes into account implicitly the symmetry property of the evolving matrices.

Constraint preserving flow: The simplest way to introduce flows on $\mathcal{F}(\Omega, \mathcal{P}(n))$ is to consider the factorization of any symmetric positive definite matrix X into a product $R^T R$ where $R \in \text{GL}(n, \mathbb{R})$, the Lie group of invertible $n \times n$ real matrices. Its Lie algebra $\mathfrak{gl}(n)$ is the set of real valued matrices, and we can identify any tangent vector W at point $R \in \text{GL}(n, \mathbb{R})$ with a vector V of $\mathfrak{gl}(n)$ through $W = RV$. This leads to flows on $\mathcal{F}(\Omega, \text{GL}(n))$ which are given by

$$\frac{\partial R}{\partial t} = R V \quad \text{where } V \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n}) \quad \text{and} \quad R_{(t=0)} = R_0 \in \mathcal{F}(\Omega, \text{GL}(n)).$$

One simply combines this expression with the factorization of X to obtain flows that act directly on $\mathcal{F}(\Omega, \mathbb{P}(n))$.

Gradient flow: Now, if we let G denotes the L^2 -gradient of f on $\mathcal{F}(\Omega, \mathbb{R}^{3 \times 3})$. We show in [8] that the natural gradient flow corresponding to the restriction of f to $\mathcal{F}(\Omega, \mathbb{P}(3))$ (endowed with its normal metric) satisfies

$$\frac{\partial X}{\partial t} = -((G + G^T) X^2 + X^2 (G + G^T)). \quad (9)$$

In our case, the matrix $G = (G_{i,j})$ is defined component by component with

$$G_{i,j} = \alpha(X_{i,j} - X_{0,i,j}) - \operatorname{div} \left(\frac{\phi'(\|\nabla_\Omega X\|)}{\|\nabla_\Omega X\|} \nabla_\Omega X_{i,j} \right). \quad (10)$$

Note that the PDE (9) allows to evolve directly the matrix coefficients while preserving the symmetric positive definiteness.

Numerical scheme: For equation (9), since the underlying manifold is an open region of a linear space, a simple Euler step could be used, provided the initial function X_0 is far enough from the constraint set boundaries. Using the exponential map theory presented before, we can introduce a more suitable scheme:

$$X_{(t+dt)} = \exp(-X_{(t)}(G + G^T) dt)^T X_{(t)} \exp(-X_{(t)}(G + G^T) dt), \quad (11)$$

where $\exp(A) = \sum_{i=0}^{\infty} \frac{A^i}{i!}$ denotes the *matrix exponential*, numerically implemented using a Padé approximation, as described in [13]. One can easily check by induction that this scheme satisfies the constraint. Let $K_{(t)} = \exp(-X_{(t)}(G_{(t)} + G_{(t)}^T) dt)$. If we assume $X_{(t)} \in \mathcal{F}(\Omega, \mathbb{P}(n))$, we have $\forall x \in \Omega, \forall v \in \mathbb{R}^n \setminus \{0\}, v^T X_{(t)}(x) v > 0$. Thus, $\forall x \in \Omega, \forall w \in \mathbb{R}^n \setminus \{0\}$,

$$\begin{aligned} w^T X_{(t+dt)}(x) w &= w^T K_{(t)}^T(x) X_{(t)}(x) K_{(t)}(x) w \\ &= (K_{(t)}(x) w)^T X_{(t)}(x) (K_{(t)}(x) w) > 0. \end{aligned}$$

Experiments on synthetic data: An application of the symmetric positive definite regularizing flow (9) on a noisy, synthetic 3D tensor field is illustrated in Fig. 3. The tensor field is quite well restored, but suffers from an *eigenvalue swelling effect*, despite the anisotropic diffusion obtained with the ϕ -function formulation. Actually, when the noise does not affect the tensor orientations as much as the diffusivities (which is the case here), a global matrix regularization process quickly smoothes the diffusivities in regard of the orientations. Therefore, to avoid this eigenvalue smoothing, we have to split the regularization of the tensor orientations and diffusivities. Note that the same kind of behavior appears in color image restoration, where several methods propose to separate the chromaticity and the brightness information from a color vector, and restore them with different equations [5, 17, 29, 31].




Fig. 3. Symmetric positive definite flow on a synthetic tensor field (with $\phi(s) = 2\sqrt{1 + s^2} - 2$).

2.3 Orthogonal Flow

Eigenvalue over-smoothing can be avoided by regularizing the orientation part U and the diffusivity part D of a diffusion tensor field X , with different coupled equations. This problem has been already considered in [10, 32]. Note that the regularization of the diffusivity field D can be handled with classic vector regularization PDE's [26, 28, 31, 35], which must however satisfy the maximum principle in order to preserve the eigenvalue positiveness. We can then focus on the regularization of the orientation feature, i.e. the regularization flow that minimizes the following ϕ -functional:

$$\min_{U \in \mathcal{F}(\Omega, \mathcal{O}(3))} f(U) = \int_{\Omega} \phi(\|\nabla_{\Omega} U\|) + \frac{\alpha}{2} (U - U_0)^2 d\Omega.$$

Constraint preserving flow: By definition, $\mathcal{O}(n) = \{X \in \text{GL}(n, \mathbb{R}) / X^T X = I\}$. Its tangent space at X is $T_X \mathcal{O}(n) = \{XA, A \in \mathfrak{so}(n)\}$, where $\mathfrak{so}(n)$ denotes its Lie algebra, the vector space of *skew-symmetric matrices*. Thus, any flow on $\mathcal{F}(\Omega, \mathcal{O}(n))$ satisfies an equation of the form

$$\frac{\partial X}{\partial t} = X A, \quad \text{where} \quad \begin{cases} A \in \mathcal{F}(\Omega, \mathfrak{so}(n)) \\ X_0 = X_{(t=0)} \in \mathcal{F}(\Omega, \mathcal{O}(n)). \end{cases} \quad (12)$$

Gradient flow: The gradient flow for the restriction of f to $\mathcal{F}(\Omega, \mathcal{O}(n))$ with respect to its bi-invariant metric⁴ is given by (see [8])

$$\frac{\partial X}{\partial t} = -\nabla f = -X \{X, G\} = X G^T X - G,$$

where $\{A, B\} = A^T B - B^T A$ denotes the generalized Lie brackets, and $G = (G_{i,j})$ is the usual gradient:

$$G_{i,j} = \alpha(U_{i,j} - U_{0,i,j}) - \text{div} \left(\frac{\phi'(\|\nabla_{\Omega} U\|)}{\|\nabla_{\Omega} U\|} \nabla_{\Omega} U_{i,j} \right).$$

⁴ $\langle XA, XB \rangle_X = - \int_{\Omega} \text{Tr}(A(x)B(x)) dx, \forall A, B \in \mathcal{F}(\Omega, \mathfrak{so}(n)).$

Sketch of Proof: Since $\nabla f \in T_X \mathcal{F}(\Omega, \mathcal{O}(n))$, we have $\nabla f = XA$, with $A \in \mathcal{F}(\Omega, \mathfrak{so}(n))$. The skew-symmetric matrix function A is then identified from the different expressions of the first variation (details in [8]).

Numerical scheme: On $\mathcal{O}(n)$, we have $\forall V \in T_X \mathcal{O}(n)$, $\exp_X V = X e^{X^{-1}V} = X e^{X^T V}$, which yields directly the following implementation of eq. (12):

$$X_{(t+dt)} = X_{(t)} \exp(-dt \{X_{(t)}, G\}). \quad (13)$$

Remark: Note that in this case, for 3×3 skew-symmetric matrices, the matrix exponential can be expressed with *Rodrigues' formula*:

$$\exp(A) = I + \frac{\sin \|a\|}{\|a\|} A + \frac{1 - \cos \|a\|}{\|a\|^2} A^2 \quad \text{where} \quad A = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \quad (\in \mathfrak{so}(3)),$$

and $a = [a_1 \ a_2 \ a_3]^T$. This equation can be used to improve the computational efficiency of the numerical implementation of eq. (13).

Experiments on synthetic data: The full regularization process of a diffusion tensor field X , using an orthogonal flow reads as follows:

1. Retrieve the orientation field U and the diffusivity field D of a tensor field X , i.e. $\forall x \in \Omega$, $X(x) = U(x)D(x)U(x)^T$.
2. Process the orientation field U with our orthogonal-constrained matrix flow (12). Note that one has to take care of the non-uniqueness of the orientation field U , as mentioned in [10, 9, 32]. Flipping one eigenvector $u^{[i]}$ (a column of U) does not change the tensor, but may affect the anisotropic regularization process. To overcome this problem, a local eigenvector alignment process is made before applying the PDE on each tensor of the field X . The neighboring eigenvector directions are aligned with the current one, by constraining the dot product between them to be positive:

$$\forall y \in \mathcal{V}(x), \quad u^{[i]}(y) := \text{sign}(u^{[i]}(x) \cdot u^{[i]}(y)) \ u^{[i]}(y),$$

where $\mathcal{V}(x)$ is a neighborhood of x . This local operation allows to act only on the orientation feature of the tensors. The importance of this alignment step is shown on Fig. 4c.

3. Finally, process the diffusivity field D with a classic vector scheme, as in [35]:

$$\frac{\partial \lambda_i}{\partial t} = \alpha (\lambda_{0_i} - \lambda_i) + \text{div}(\Sigma \nabla_{\Omega} \lambda_i),$$

where Σ is a matrix that drive the diffusion process and may contain some *a-priori* physiological information.

As we can see in Fig. 4, the tensor field seems well restored. The final result is closer to the original image, and does not suffer from eigenvalue swelling. Note that the restoration of tensor orientations was initially proposed in [9, 10, 32] using similar continuous flows but different numerical integration methods.

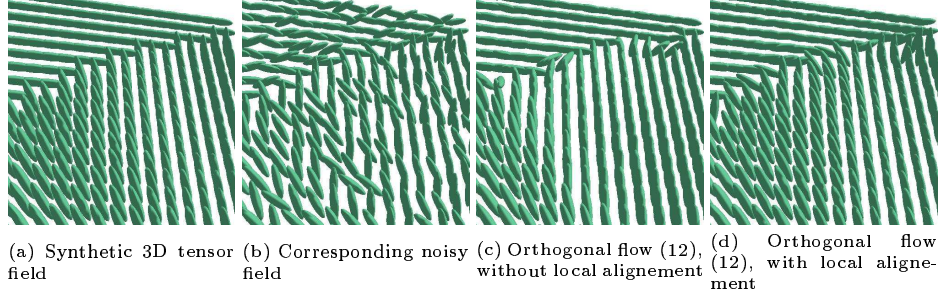


Fig. 4. Orthogonal flow on a synthetic 3D tensor field.

2.4 Isospectral flow

To avoid the eigenvector alignment step needed by the previous method (12), we propose here a simple way of regularizing directly tensor orientations. It consists in applying on the initial tensor field an *isospectral flow*, i.e. a regularizing flow that preserves the tensor diffusivities, while regularizing the tensor orientations.

Constraint preserving flow: We first consider the homogeneous space $M(Q) = \{H^T Q H, H \in O(n)\}$ of all the matrices orthogonally equivalent to a symmetric matrix Q . If all the values of X_0 had the same eigenvalues (the same as the matrix Q), we could model our configuration space as $\mathcal{F}(\Omega, M(Q))$. However, if we allow the eigenvalues of X_0 to vary spatially, we need to introduce a slightly more generic manifold $\mathcal{F}_{X(0)}^{\text{iso}}(\Omega) = \{H^T X(0) H, H \in \mathcal{F}(\Omega, O(n))\}$. From the expression of its tangent space $T_X \mathcal{F}_{X(0)}^{\text{iso}}(\Omega) = \{[X, A], A \in \mathcal{F}(\Omega, \mathfrak{so}(n))\}$ (the Lie brackets $[\cdot, \cdot]$ corresponds to the commutator $[A, B] = AB - BA$), the expected *isospectral* flows must satisfy

$$\frac{\partial X(t)}{\partial t} = [X, A(t)], \quad A(t) \in \mathcal{F}(\Omega, \mathfrak{so}(n)).$$

Gradient flow: Given the L^2 -gradient G on $\mathcal{F}(\Omega, \mathbb{R}^{n \times n})$ (defined by (10)), we show in [8] that a natural minimizing flow for f on $\mathcal{F}_{X(0)}^{\text{iso}}(\Omega)$ is defined by

$$\frac{\partial X}{\partial t} = [X, [X, -(G + G^T)]] . \quad (14)$$

One can also obtain a suitable integration scheme for eq. (14), derived from the expression of the exponential map on the orthogonal group, and such that

$$\forall x \in \Omega, \quad X_{(t+dt)}(x) = A_{(t)}(x)^T X_{(t)}(x) A_{(t)}(x), \quad \text{with} \\ A_{(t)}(x) = e^{-dt[G^T(x) + G(x), X_{(t)}(x)]}. \quad (15)$$

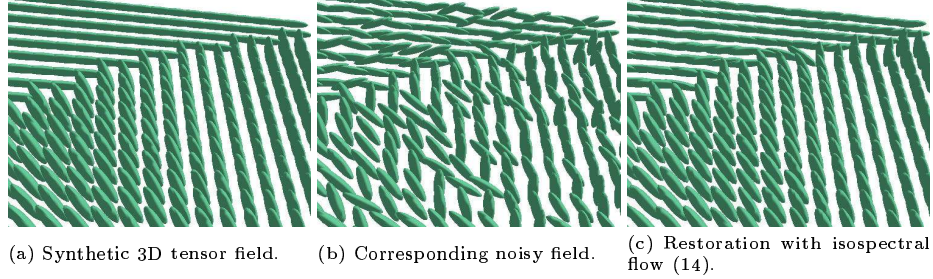


Fig. 5. Isospectral flow on a synthetic 3×3 tensor field.

Experimental results: Equation (14) was applied to the usual noisy, synthetic tensor field. This experiment is illustrated in Fig. 5.

As expected, the results are similar to the ones obtained with the decomposition and the orthogonal constraints. However, the computation cost is significantly reduced (no local alignment steps are needed anymore for regularizing tensor orientations). The PDE (14) applies directly on the original matrix coefficients $X_{i,j}$.

3 Results on Real DT-MRI Datasets

We have tested the proposed diffusion tensor regularization PDE's (9), (12), (14), on a real $128 \times 128 \times 56$ DT-MRI dataset⁵ of the brain (Fig. 6a,b). On the left side of each figure, we represent the tensors as ellipsoids, while the right side shows the corresponding fibers, tracked by following the main direction of each tensor. These fibers show the path followed by the water molecules. To make visualization easier, only one slice of the processed volume is represented. The computations have been carried out with a Pentium 1Ghz processor, 1Go RAM, and the corresponding computation times are given for the full volume restoration process. We noticed that the geodesic step approach allows time steps to be relatively high, and very few iterations are needed in practice. The ϕ -function used for these experiments is the “hypersurface” function, $\phi(s) = 2\sqrt{1 + s^2} - 2$ proposed in [7]. The results we obtained clearly illustrates the behavior of the different methods:

- The symmetric positive definite flow (9) tends to swell eigenvalues.
- The orthogonal constrained flow (12) works well with a local alignment step (Fig. 6e), but fails otherwise (Fig. 6d).
- The isospectral flow (14) has a quite similar behavior, but is more computationally efficient.

⁵ The authors would like to thank J.-F. Mangin and J.-B. Poline (SHFJ-CEA) for providing us with the DT-MRI data (this work was partially supported by ARC MC2). We also thank R. Fournier for his visualization tool “TensView”.

Even if the physiological validation of these results remains to be done, our methods seem to correct most artifacts (due to the image acquisition process) and retrieve the main global structures of the fiber network. It opens new perspectives to construct an accurate fiber map model of the brain.

Conclusion

In this paper, we proposed a geometric approach to construct general flows acting on constrained matrix-valued images. We also introduced numerical schemes that naturally fit into this framework thanks to the use of exponential maps. The efficiency and versatility of this setting has been illustrated with the problem of diffusion tensor regularization. Three different regularizing flows were introduced (symmetric positive definite, orthogonal, isospectral). Some of them generalize and extend previous works in this field. Numerical experiments show promising results even if, at this time, no physiological a-priori knowledge has been taken into account (for instance in the choice of a suitable ϕ -function). Future developments will include applications to other computer vision problems.

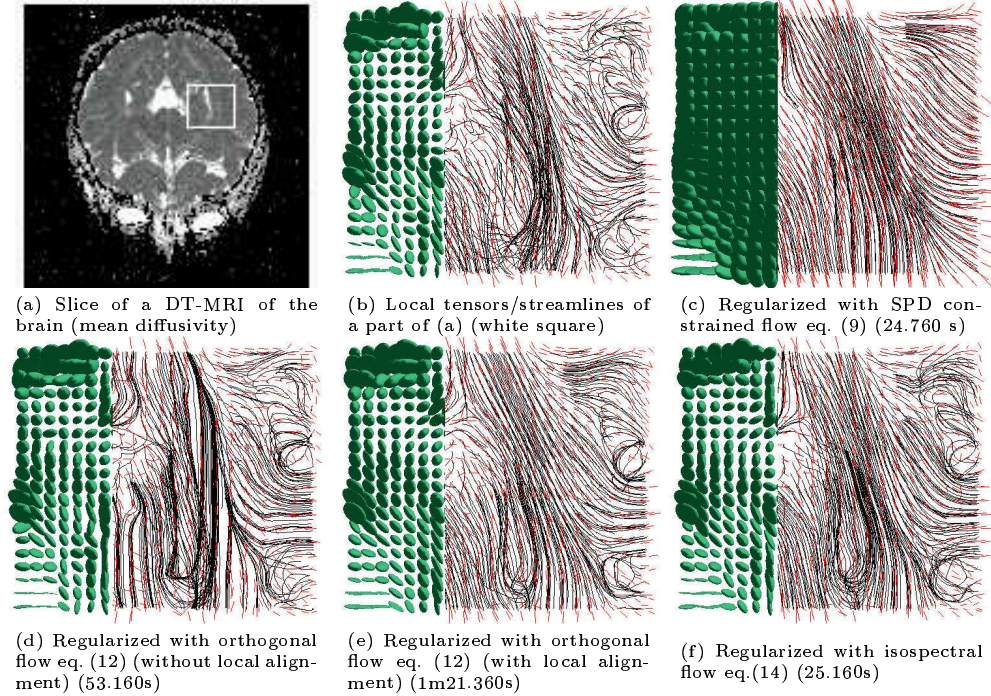


Fig. 6. Real diffusion tensor dataset (DT-MRI) regularization, using three different constrained methods (last numbers represent the computation time).

References

1. L. Alvarez, R. Deriche, J. Weickert, and J. Sánchez. Dense disparity map estimation respecting image discontinuities: A PDE and scale-space based approach. *International Journal of Visual Communication and Image Representation*, 2000.
2. Y. Amit. A nonlinear variational problem for image matching. *SIAM Journal on Scientific Computing*, 15(1), January 1994.
3. M. Bertalmio, L.T. Cheng, S. Osher, and G. Sapiro. Variational problems and partial differential equations on implicit surfaces: The framework and examples in image processing and pattern formation. *UCLA Research Report*, June 2000.
4. M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In Kurt Akeley, editor, *Proceedings of the SIGGRAPH*, pages 417–424. ACM Press, ACM SIGGRAPH, Addison Wesley Longman, 2000.
5. T. Chan and J. Shen. Variational restoration of non-flat image features : Models and algorithms. *Research Report, Computational and applied mathematics, department of mathematics Los Angeles*, June 1999.
6. T. Chan and J. Shen. Mathematical models for local deterministic inpaintings. Technical Report 00-11, Department of Mathematics, UCLA, Los Angeles, March 2000.
7. P. Charbonnier, G. Aubert, M. Blanc-Féraud, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of the International Conference on Image Processing*, volume II, pages 168–172, 1994.
8. C. Chefd'hotel, D. Tschumperlé, O. Faugeras and R. Deriche. Geometric Integration of Constraint Preserving Flows and Applications to Image Processing. *INRIA Research Report*, to appear, 2002.
9. O. Coulon, D.C. Alexander, and S.R. Arridge. A geometrical approach to 3d diffusion tensor magnetic resonance image regularisation. Technical Report, Department of Computer Science, University College London., 2001.
10. O. Coulon, D.C. Alexander, and S.R. Arridge. A regularization scheme for diffusion tensor magnetic resonance images. In *XVIIth International Conference on Information Processing in Medical Imaging*, 2001.
11. R. Deriche, P. Kornprobst, and G. Aubert. Optical flow estimation while preserving its discontinuities: A variational approach. In *Proceedings of the 2nd Asian Conference on Computer Vision*, volume 2, pages 71–80, Singapore, December 1995.
12. O. Faugeras and R. Keriven. Variational principles, surface evolution, PDE's, level set methods and the stereo problem. *IEEE Transactions on Image Processing*, 7(3):336–344, March 1998.
13. G. Golub and C. Van Loan. *Matrix computations*. The John Hopkins University Press, Baltimore, Maryland, second edition, 1989.
14. G. Granlund and H. Knutsson. *Signal Processing for Computer Vision*. Kluwer Academic Publishers, 1995.
15. U. Helmke and J. Moore. *Optimization and Dynamical Systems*. Springer, 1994.
16. V. Kac, editor. *Infinite Dimensional Lie Groups with Applications*. Mathematical Sciences Research Institute Publications 4. Springer, 1985.
17. R. Kimmel and N. Sochen. Orientation diffusion or how to comb a porcupine. Technical Report 2000-02, CIS, 2000. Accepted to special issue on PDEs in Image Processing, Computer Vision, and Computer Graphics, Journal of Visual Communication and Image Representation, 2000.

18. W. Klingenberg. *Riemannian Geometry*. de Gruyter Studies in Mathematics 1. Walter de Gruyter, 1982.
19. P. Kornprobst, R. Deriche, and G. Aubert. Nonlinear operators in image restoration. In *Proceedings of the International Conference on Computer Vision and Pattern Recognition*, pages 325–331, Puerto Rico, June 1997. IEEE Computer Society, IEEE.
20. S. Lang. *Differential Manifolds*. Springer, 1985.
21. D. Le Bihan. Methods and applications of diffusion MRI. In I.R. Young, editor, *Magnetic Resonance Imaging and Spectroscopy in Medicine and Biology*. John Wiley and Sons, 2000.
22. D. Mumford and J. Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42:577–684, 1989.
23. P. Perona. Orientation diffusions. *IEEE Transactions on Image Processing*, 7(3):457–467, March 1998.
24. P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990.
25. L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992.
26. G. Sapiro. *Geometric Partial Differential Equations and Image Analysis*. Cambridge University Press, 2001.
27. S. Smith. *Optimization Techniques on Riemannian Manifolds*. Hamiltonian and Gradient Flows, Algorithms and Control. American Mathematical Society, 1994, Fields Institute for Research in Mathematical Sciences, A. Bloch editor, 1994.
28. N. Sochen, R. Kimmel, and R. Malladi. A geometrical framework for low level vision. *IEEE Transaction on Image Processing, Special Issue on PDE based Image Processing*, 7(3):310–318, 1998.
29. B. Tang, G. Sapiro, and V. Caselles. Diffusion of general data on non-flat manifolds via harmonic maps theory : The direction diffusion case. *The International Journal of Computer Vision*, 36(2):149–161, February 2000.
30. A. Trounev. Diffeomorphisms groups and pattern matching in image analysis. *International Journal of Computer Vision*, 28(3):213–21, 1998.
31. D. Tschumperlé and R. Deriche. Constrained and unconstrained PDE's for vector image restoration. In Ivar Austvoll, editor, *Proceedings of the 10th Scandinavian Conference on Image Analysis*, pages 153–160, Bergen, Norway, June 2001.
32. D. Tschumperlé and R. Deriche. Diffusion tensor regularization with constraints preservation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Kauai Marriott, Hawaii, December 2001.
33. D. Tschumperlé and R. Deriche. Regularization of orthonormal vector sets using coupled PDE's. In *Proceedings of the 1st IEEE Workshop on Variational and Level Set Methods (VLSM'01)*, July 2001.
34. B. Vemuri, Y. Chen, M. Rao, T. McGraw, T. Mareci, and Z. Wang. Fiber tract mapping from diffusion tensor MRI. In *Proceedings of the 1st IEEE Workshop on Variational and Level Set Methods in Computer Vision (VLSM'01)*, July 2001.
35. J. Weickert. *Anisotropic Diffusion in Image Processing*. Teubner-Verlag, Stuttgart, 1998.
36. C.-F. Westin, S. Maier, B. Khidhir, P. Everett, F. Jolesz, and R. Kikinis. Image processing for diffusion tensor magnetic resonance imaging. In *Proceedings of the Second International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'99)*, Springer-Verlag, 1999.