

REGULARIZING FLOWS FOR CONSTRAINED MATRIX-VALUED IMAGES

C. CHEFD'HOTEL, D. TSCHUMPERLÉ, R. DERICHE, AND O. FAUGERAS

ABSTRACT. Nonlinear diffusion equations are now widely used to restore and enhance images. They allow to eliminate noise and artifacts while preserving large global features, such as object contours. In this context, we propose a differential-geometric framework to define PDEs acting on some manifold constrained datasets. We consider the case of images taking value into matrix manifolds defined by orthogonal and spectral constraints. We directly incorporate the geometry and natural metric of the underlying configuration space (viewed as a Lie group or a homogeneous space) in the design of the corresponding flows. Our numerical implementation relies on structure-preserving integrators that respect intrinsically the constraints geometry. The efficiency and versatility of this approach are illustrated through the anisotropic smoothing of diffusion tensor volumes in medical imaging. **Note: This is the draft of a paper published in *Journal of Mathematical Imaging and Vision* 20:147–162, 2004. Do not distribute.**

1. INTRODUCTION

Variational methods and nonlinear partial differential equations (PDEs) are now widely used to tackle computer vision problems, such as image restoration, segmentation, stereo-based 3D reconstruction, or optical flow estimation (see the textbooks [3, 27, 31, 44] and references therein for an overview). Solutions to these problems, whether they are curves, surfaces, images, or vector fields, are generally obtained by continuously deforming an initial estimate through a flow defined by a PDE. The corresponding evolution equations derive from simple local heuristics or from the minimization of cost functionals. In the context of image restoration, the idea is to achieve a selective smoothing that removes the noise while preserving large global features, such as object contours. For this purpose, one generally uses anisotropic diffusion PDEs (comprehensive reviews on nonlinear diffusion equations in image processing can be found in [3, 36, 44]). Generalizing these techniques from gray-valued images to multi-valued datasets has recently attracted a growing interest. This type of problem arises, for instance, when dealing with color images, direction fields, trajectories of camera orientations, DT-MRI volumes, or fields of statistical parameters in Doppler analysis¹. In these cases, the extension of standard methods is usually nontrivial due to the existence of additional point-wise constraints: Unit norm, orthogonality, positive definiteness (among others).

Several solutions have been proposed in recent works. The first approach was suggested by Perona in [29] for images taking value into the unit circle S^1 . In this

Date: August 9, 2004.

Key words and phrases. Image regularization, multi-valued data, constraints, Lie groups, homogeneous spaces, Riemannian metric, natural gradient, structure-preserving algorithms.

¹We thank F. Barbaresco who recently brought this problem to our attention (see [4]).

work, the problem of using a standard diffusion equation acting on a parametrization of S^1 was discussed. Then, more intrinsic geometric approaches were considered, by modeling the constrained dataset as a mapping $\mathbf{X} : \mathcal{D} \rightarrow \mathcal{N}$ from a domain manifold \mathcal{D} into a target manifold \mathcal{N} . In this setting, a geometric formulation of the Total Variation (TV) restoration model was given in [8], Beltrami flows were proposed in [23, 34, 4], and nonlinear heat equations borrowed from harmonic theory were presented in [25, 35, 43]. Most of these approaches were applied in the case where the target manifold was the unit hypersphere S^{n-1} . The problem of building PDEs acting on fields of orthogonal matrices was also discussed in [38, 39] using the formalism of Lagrange multipliers.

Representing the nonlinear geometry of \mathcal{D} and \mathcal{N} is one of the main issues. A very elegant way of dealing with a non-flat domain \mathcal{D} was suggested in [5] using an implicit formulation in terms of level-sets. Several solutions have been proposed to represent the values of \mathbf{X} , either using local coordinate charts on \mathcal{N} [29, 23], extrinsic coordinates in the embedding space [35, 38, 39], or another implicit representation as level-set of an auxiliary function [25]. Such a choice has a strong influence on the discretization of the problem. A complete parametrization requires to switch between coordinate charts during the evolution, while models using extrinsic coordinates may numerically violate the constraints (they require at least a re-projection).

In this work², our final objective is to extend and generalize some of these ideas to matrix-valued functions undergoing orthogonal or spectral constraints. The application of interest is the regularization of DT-MRI data in medical imaging. In a larger perspective, we propose to model PDEs acting on constrained multi-valued datasets as evolution equations on a suitable infinite-dimensional manifold of mappings. One should be aware that a rigorous definition of these concepts raises numerous mathematical technicalities that we merely start to address here. Our approach and derivations, developed by analogy with finite-dimensional problems, remain essentially formal in this paper, and will have to be further investigated and verified. However, the advantage of this geometric perspective is to yield simple tools to design and implement the constrained counterparts to widely used PDEs in image processing. Moreover, this setting offers a way to take into account spatially varying constraints (such as the isospectral constraint detailed in this work). In this case, datasets can not be modeled as mappings between manifolds.

We also strongly believe that the prior knowledge one may have on the constraints geometry should be integrated, not only in the design of the continuous equations, but also in the corresponding numerical schemes. For this purpose, we propose to generalize the idea of geodesic marching, suggested in [8], in the larger perspective of geometric integration methods [17].

This paper is organized as follows: We propose a geometric characterization of evolution equations on three constrained sets of matrix-valued mappings in Section 2. In Section 3, we introduce structure-preserving numerical integrators that respect the geometry of the underlying constrained sets. In Section 4, regularizing flows are built using a systematic approach to transform unconstrained gradient flows into their constrained counterparts. Finally, we illustrate this formal setting with the regularization of DT-MRI volumes in Section 5, and present our concluding remarks in Section 6.

²This paper develops some of the ideas previously sketched by the authors in [9].

2. GEOMETRIC SETTING

2.1. Basic principle. We start with a generic remark on the geometry of constrained flows. Let us assume a set of nonlinear constraints defines a submanifold \mathcal{M} in a linear space \mathcal{E} (the initial configuration space of the problem). If \mathbf{V} is a smooth vector field on \mathcal{E} , and the one-parameter family $t \mapsto \mathbf{X}(t)$ is solution of the evolution equation

$$\partial_t \mathbf{X} = \mathbf{V}(\mathbf{X}), \quad \mathbf{X}(0) \in \mathcal{E},$$

the following equivalence holds [1, 17]:

$$\mathbf{X}(t) \in \mathcal{M}, \quad \forall t > 0 \iff \mathbf{X}(0) \in \mathcal{M} \text{ and } \mathbf{V}(\mathbf{X}(t)) \in T_{\mathbf{X}(t)}\mathcal{M}, \quad \forall t > 0.$$

The main idea is that building a flow in the embedding space and satisfying the constraints amounts to characterize the tangent space at any given point on \mathcal{M} . We illustrate this principle in Fig. 1, with the example of a flow induced by a vector field on sphere embedded in \mathbb{R}^3 .

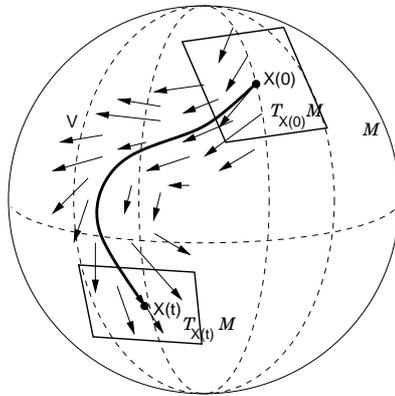


FIGURE 1. Integral curve of a vector field on S^2 .

An alternative approach is to directly define a constrained flow using a parametrization on \mathcal{M} . However, a parametrization is not unique, and using extrinsic coordinates in the embedding space avoids dealing with several coordinate charts (which are often necessary to cover \mathcal{M} entirely).

2.2. Constrained sets. Note that the previous remarks are valid, at least formally, for both finite and infinite-dimensional manifolds. Here, our embedding space is a linear set of matrix-valued functions $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$ where Ω is a bounded region in \mathbb{R}^k , and the set of real-valued $n \times p$ matrices is identified with $\mathbb{R}^{n \times p}$. For our applications, we are mostly interested in constraints acting point-wise on these mappings. In the simplest situation, we want them to take value into a target submanifold $\mathcal{N} \subset \mathbb{R}^{n \times p}$. This leads us to consider the constrained set

$$\mathcal{F}(\Omega, \mathcal{N}) = \{\mathbf{X} \in \mathcal{F}(\Omega, \mathbb{R}^{n \times p}) / \forall \mathbf{p} \in \Omega, \mathbf{X}(\mathbf{p}) \in \mathcal{N}\}.$$

Naturally, we can expect $\mathcal{F}(\Omega, \mathcal{N})$ to inherit most of its properties from the geometry of its codomain \mathcal{N} . We refer the interested reader to [41] to see how, with a suitable function-space topology, $\mathcal{F}(\Omega, \mathcal{N})$ could be equipped with infinite-dimensional

structure of Banach or Hilbert manifold³. Note that we will also consider a slightly more generic set of mappings to take into account spatially varying constraints.

Many constraints on matrices (orthogonality, rank and other spectral constraints) can be expressed in terms of Lie groups and homogeneous spaces. They are differentiable manifolds that present some nice algebraic properties. A Lie group is a manifold which has also a group structure (the group operation and its inverse are smooth). Its tangent space at any given point can be easily defined in terms of its tangent space at identity (its Lie algebra) by left- (or right-) translation. And a homogeneous space is a manifold on which a transitive Lie group action exists. We refer to differential geometry textbooks [6, 18, 26] for an introduction to this topic. In the following, we propose to explore some of the properties of these matrix manifolds to define and study three related sets of matrix-valued mappings of particular interest in image processing.

2.2.1. Orthogonality constraints. Numerous multi-valued datasets undergo, directly or indirectly, orthogonality constraints. We can first mention orientation fields (for instance obtained from optical flow algorithms), and chromaticity features of color images, which are both made of unit norm vectors. There are also camera orientation trajectories and orientation features of DT-MRI volumes, which are modeled as fields of orthogonal matrices. The suitable configuration space for this type of data is the set of mappings taking value into the Stiefel manifold $\text{St}(n, p)$. $\text{St}(n, p)$ is the set of matrices made of p orthonormal vectors of size n :

$$\text{St}(n, p) = \{\mathbf{X} \in \mathbb{R}^{n \times p} / \mathbf{X}^T \mathbf{X} = \mathbf{I}_p\}.$$

$\text{St}(n, p)$ is a homogeneous space (any element in $\mathbf{X} \in \text{St}(n, p)$ can be viewed as $\mathbf{X} = \mathbf{Q} \mathbf{I}_{n \times p}$, where \mathbf{Q} is an element of the Lie group of orthogonal matrices $\text{O}(n)$, and $\mathbf{I}_{n \times p}$ is an incomplete $n \times p$ identity matrix). $\text{O}(n)$ and the unit hypersphere S^{n-1} are special instances of $\text{St}(n, p)$, for $p = n$ and $p = 1$ respectively. We introduce here the simplified notations:

$$\mathcal{F}_{stiefel} = \mathcal{F}(\Omega, \text{St}(n, p)) \text{ and } \mathcal{F}_{ortho} = \mathcal{F}(\Omega, \text{O}(n)).$$

Note that a set of mappings taking value into a Lie group, such as \mathcal{F}_{ortho} , endowed a suitable topology, is an example of infinite-dimensional Lie group [22].

2.2.2. Prescribed rank/signature. We denote $\text{Rk}(r, n)$ the set of $n \times n$ real matrices of rank $r \leq n$. This set has a manifold structure which can be derived from a group action of the product $\text{GL}(n, \mathbb{R}) \times \text{GL}(n, \mathbb{R})$ (where $\text{GL}(n, \mathbb{R})$ is the Lie group of nonsingular $n \times n$ real matrices) [19, 26]. We will consider here fields of *symmetric* matrices with a prescribed rank. The corresponding set of mappings is defined as

$$\mathcal{F}_{rank} = \mathcal{F}(\Omega, \text{S}(r, n)),$$

where $\text{S}(r, n)$ is the subset of $\text{Rk}(r, n)$ given by

$$\text{S}(r, n) = \{\mathbf{X} \in \mathbb{R}^{n \times n} / \mathbf{X}^T = \mathbf{X} \text{ and } \text{Rank}(\mathbf{X}) = r\}.$$

The $(r + 1)$ connected components of $\text{S}(r, n)$ are made of matrices with identical signature (the difference between the number of positive and negative eigenvalues) [19]. Fields of symmetric positive (semi-)definite matrices, used to represent DT-MRI volumes in the following, can be modeled using \mathcal{F}_{rank} .

³Note that other manifolds of mappings have been explicitly introduced in the context of computer vision and image processing, such as the group of diffeomorphisms considered in [37].

2.2.3. *Prescribed eigenvalues.* Also in the context of DT-MRI data, we investigate a stronger constraint: Fields of matrices with prescribed eigenvalues. Our goal is to build *isospectral* flows acting on a field \mathbf{X}_0 of real symmetric matrices while preserving their eigenvalues through time. We recall that if two matrices \mathbf{A} and \mathbf{B} have the same spectrum, there exists a similarity transformation between them, i.e. $\exists \mathbf{Q} \in \text{GL}(n) / \mathbf{A} = \mathbf{Q}^{-1}\mathbf{B}\mathbf{Q}$ (furthermore, \mathbf{Q} is orthogonal if \mathbf{A} and \mathbf{B} are real symmetric) [20]. In this case, since we do not assume the matrix eigenvalues are the same at all points in Ω , the point-wise constraint we impose on the field varies spatially. The underlying set of functions is not of the form $\mathcal{F}(\Omega, \mathcal{N})$. Instead, by analogy with the finite-dimensional isospectral manifold presented in [19], we consider the transitive action induced by a field of orthonormal matrices. We define our constrained configuration space as

$$\mathcal{F}_{iso} = \{\mathbf{U}^T \mathbf{X}_0 \mathbf{U} / \mathbf{U} \in \mathcal{F}_{ortho}\}.$$

Here and in the sequel, we assume that we extend to $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$ the standard matrix operations: Transpose, matrix product and inverse, matrix exponential, etc. They just apply point-wise.

2.3. **Tangent spaces.** Now, in order to build PDEs acting on $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} and \mathcal{F}_{iso} , the next step is to identify their tangent spaces. The tangent space at a point \mathbf{X} on a manifold \mathcal{M} is generally defined as

$$T_{\mathbf{X}}\mathcal{M} = \{\mathbf{H} / \exists \text{ a path } \Gamma : (-\varepsilon, \varepsilon) \rightarrow \mathcal{M} \text{ such that } \Gamma(0) = \mathbf{X}, \dot{\Gamma}(0) = \mathbf{H}\}.$$

When \mathcal{M} is a manifold of mappings $\mathcal{F}(\Omega, \mathcal{N})$, its tangent space will be formally identified [41] to

$$T_{\mathbf{X}}\mathcal{F}(\Omega, \mathcal{N}) = \{\mathbf{V} \in \mathcal{F}(\Omega, \mathbb{R}^{n \times p}) / \forall \mathbf{p} \in \Omega, \mathbf{V}(\mathbf{p}) \in T_{\mathbf{X}(\mathbf{p})}\mathcal{N}\}.$$

2.3.1. *Orthogonality constraints.* We set $\mathcal{N} = \text{St}(n, p)$. A derivation of the constraint $\mathbf{X}^T \mathbf{X} = \mathbf{I}_p$ yields

$$T_{\mathbf{X}}\text{St}(n, p) = \{\mathbf{H} \in \mathbb{R}^{n \times p} / \mathbf{H}^T \mathbf{X} + \mathbf{X}^T \mathbf{H} = \mathbf{0}\}.$$

In practice, we also use the explicit parametrization proposed in [14]:

$$T_{\mathbf{X}}\text{St}(n, p) = \{\mathbf{X}\mathbf{A} + \mathbf{X}_{\perp}\mathbf{B} / \mathbf{A} \in \mathfrak{so}(p), \mathbf{B} \in \mathbb{R}^{(n-p) \times p}\},$$

where $\mathfrak{so}(p)$ is the set of $p \times p$ skew-symmetric matrices, and \mathbf{X}_{\perp} is the $n \times (n-p)$ matrix such that $\mathbf{X}\mathbf{X}^T + \mathbf{X}_{\perp}\mathbf{X}_{\perp}^T = \mathbf{I}_n$. Equipped with the commutator $[\mathbf{A}, \mathbf{B}] = \mathbf{A}\mathbf{B} - \mathbf{B}\mathbf{A}$ (Lie bracket), $\mathfrak{so}(n)$ corresponds to the Lie algebra of the orthogonal group $\text{O}(n)$. In the true orthogonal case: $\text{O}(n) = \text{St}(n, n)$, the tangent space reduces to $T_{\mathbf{X}}\text{O}(n) = \{\mathbf{X}\mathbf{A} / \mathbf{A} \in \mathfrak{so}(n)\}$. Finally, if we denote $\mathcal{F}_{skew} = \mathcal{F}(\Omega, \mathfrak{so}(p))$, we can write

$$T_{\mathbf{X}}\mathcal{F}_{stiefel} = \{\mathbf{X}\mathbf{A} + \mathbf{X}_{\perp}\mathbf{B} / \mathbf{A} \in \mathcal{F}_{skew}, \mathbf{B} \in \mathcal{F}(\Omega, \mathbb{R}^{(n-p) \times p})\}.$$

Thus, we can then expect evolution equations on $\mathcal{F}_{stiefel}$ to have the following form:

$$(2.1) \quad \partial_t \mathbf{X} = \mathbf{X}\mathbf{A}(\mathbf{X}) + \mathbf{X}_{\perp}\mathbf{B}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}, \mathbf{B}(\mathbf{X}) \in \mathcal{F}(\Omega, \mathbb{R}^{(n-p) \times p}).$$

In particular, this equation reduces to

$$\partial_t \mathbf{X} = \mathbf{X}\mathbf{A}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew},$$

on \mathcal{F}_{ortho} (i.e. when $p = n$).

We note here \mathbf{A} and \mathbf{B} as functions of \mathbf{X} , but \mathbf{A} and \mathbf{B} could be any functions (operators) taking respectively value in \mathcal{F}_{skew} and $\mathcal{F}(\Omega, \mathbb{R}^{(n-p) \times p})$. It is the choice

of \mathbf{A} and \mathbf{B} that yields specific PDEs acting on $\mathcal{F}_{stiefel}$. In our application, \mathbf{A} and \mathbf{B} will be differential operators (in fact diffusion operators) allowing the progressive smoothing of the original field \mathbf{X}_0 .

2.3.2. *Prescribed rank/signature.* One can show (see [19] for details) that $T_{\mathbf{X}}\text{Rk}(r, n) = \{\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{B} / \mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}\}$, and that the expression of the tangents spaces on $S(r, n)$ reduces to $T_{\mathbf{X}}S(r, n) = \{\mathbf{A}^T\mathbf{X} + \mathbf{X}\mathbf{A} / \mathbf{A} \in \mathbb{R}^{n \times n}\}$. By extension, $T_{\mathbf{X}}\mathcal{F}_{rank}$ follows directly:

$$T_{\mathbf{X}}\mathcal{F}_{rank} = \{\mathbf{A}^T\mathbf{X} + \mathbf{X}\mathbf{A} / \mathbf{A} \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n})\}.$$

The corresponding rank/signature preserving evolution equations thus satisfy

$$(2.2) \quad \partial_t \mathbf{X} = \mathbf{A}(\mathbf{X})^T \mathbf{X} + \mathbf{X}\mathbf{A}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n}).$$

2.3.3. *Prescribed eigenvalues.* In this case, the task of characterizing the tangent space is bit more difficult. Let $\mathbf{X} \in \mathcal{F}_{iso}$. By extension of the argument developed in [19] in the finite-dimensional case, we use the surjective linear map between $T_{\mathbf{Z}}\mathcal{F}_{ortho}$ and $T_{\mathbf{Z}^T\mathbf{X}\mathbf{Z}}\mathcal{F}_{iso}$ induced by the differential of the smooth map:

$$\sigma^{\mathbf{X}}(\mathbf{Z}) : \begin{array}{ccc} \mathcal{F}_{ortho} & \longrightarrow & \mathcal{F}_{iso} \\ \mathbf{Z} & \longmapsto & \mathbf{Z}^T\mathbf{X}\mathbf{Z}. \end{array}$$

In fact, its differential at point \mathbf{Z} in the direction \mathbf{H} in $T_{\mathbf{Z}}\mathcal{F}_{ortho}$ is

$$d\sigma_{\mathbf{Z}}^{\mathbf{X}}(\mathbf{H}) = \mathbf{H}^T\mathbf{X}\mathbf{Z} + \mathbf{Z}^T\mathbf{X}\mathbf{H},$$

which yields

$$\begin{aligned} T_{\mathbf{Z}^T\mathbf{X}\mathbf{Z}}\mathcal{F}_{iso} &= \{\mathbf{H}^T\mathbf{X}\mathbf{Z} + \mathbf{Z}^T\mathbf{X}\mathbf{H} / \mathbf{H} \in T_{\mathbf{Z}}\mathcal{F}_{ortho}\} \\ &= \{(\mathbf{Z}\mathbf{A})^T\mathbf{X}\mathbf{Z} + \mathbf{Z}^T\mathbf{X}\mathbf{Z}\mathbf{A} / \mathbf{A} \in \mathcal{F}_{skew}\} \\ &= \{(\mathbf{Z}^T\mathbf{X}\mathbf{Z})\mathbf{A} - \mathbf{A}(\mathbf{Z}^T\mathbf{X}\mathbf{Z}) / \mathbf{A} \in \mathcal{F}_{skew}\}. \end{aligned}$$

Since, $\forall \mathbf{Y} \in \mathcal{F}_{stiefel}$ there exists $\mathbf{Z} \in \mathcal{F}_{ortho}$ such that $\mathbf{Y} = \mathbf{Z}^T\mathbf{X}\mathbf{Z}$, we obtain

$$\begin{aligned} T_{\mathbf{Y}}\mathcal{F}_{iso} &= \{\mathbf{Y}\mathbf{A} - \mathbf{A}\mathbf{Y} / \mathbf{A} \in \mathcal{F}_{skew}\} \\ &= \{[\mathbf{Y}, \mathbf{A}] / \mathbf{A} \in \mathcal{F}_{skew}\}, \end{aligned}$$

where we use the Lie bracket notation $[\mathbf{A}, \mathbf{B}] = \mathbf{A}\mathbf{B} - \mathbf{B}\mathbf{A}$. Consequently, flows on \mathcal{F}_{iso} satisfy

$$(2.3) \quad \partial_t \mathbf{X} = [\mathbf{X}, \mathbf{A}(\mathbf{X})] / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}.$$

A summary of all these constrained flows is given in Table 1 below. But before we try to specialize them to perform image processing tasks, we propose to immediately address the problem of their numerical implementation.

$\mathcal{F}_{stiefel}$	$\partial_t \mathbf{X} = \mathbf{X}\mathbf{A}(\mathbf{X}) + \mathbf{X}_{\perp}\mathbf{B}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}, \mathbf{B}(\mathbf{X}) \in \mathcal{F}(\Omega, \mathbb{R}^{(n-p) \times p})$
\mathcal{F}_{ortho}	$\partial_t \mathbf{X} = \mathbf{X}\mathbf{A}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}$
\mathcal{F}_{rank}	$\partial_t \mathbf{X} = \mathbf{A}(\mathbf{X})^T\mathbf{X} + \mathbf{X}\mathbf{A}(\mathbf{X}) / \mathbf{A}(\mathbf{X}) \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n})$
\mathcal{F}_{iso}	$\partial_t \mathbf{X} = [\mathbf{X}, \mathbf{A}(\mathbf{X})] / \mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}$

TABLE 1. Generic form of the constrained flows.

3. STRUCTURE-PRESERVING INTEGRATORS

3.1. Numerical integration on manifolds. Previously, we looked at the intrinsic geometry of $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} and \mathcal{F}_{iso} using a representation in terms of extrinsic coordinates in $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$. As mentioned before, the main advantage of this approach is to avoid switching between local coordinate charts. However, the corresponding numerical implementation requires a special attention. In fact, if we use standard techniques for evolution PDEs, there is a risk of stepping out from the manifold after each iteration. The flow does not lie in a linear space, and one must adjust the integration method to accommodate the curved constrained geometry. The point-wise constraints considered in this paper only have consequences on the discretization in time of the equations. We assume standard finite-difference techniques are used for the spatial discretization⁴ of the differential operators that appear in the design of specific flows. We use a semi-discrete formulation in the following.

On an arbitrary manifold \mathcal{M} , the objective is to build a step-forward operator K_ε (ε being the time step), such that the discrete flow

$$\mathbf{X}_{k+1} = K_\varepsilon(\mathbf{X}_k), \quad \mathbf{X}_0 \in \mathcal{M},$$

provides a consistent approximation of the solution of the equation

$$\partial_t \mathbf{X} = \mathbf{V}(\mathbf{X}), \quad \mathbf{X}(0) \in \mathcal{M}, \quad \mathbf{V}(\mathbf{X}) \in T_{\mathbf{X}}\mathcal{M}.$$

To understand what consistency means in this case, we consider the general setting proposed by Chorin et al. [10] for numerical algorithms and evolution equations on (infinite-dimensional) manifolds. A consistent operator is defined as a mapping $K_\varepsilon : \mathcal{M} \rightarrow \mathcal{M}$, that provides at least a first order approximation in time of the continuous flow⁵. That is,

$$\begin{cases} K_\varepsilon : \mathcal{M} \rightarrow \mathcal{M}, \\ K_0(\mathbf{X}) = \mathbf{X}, \\ \partial_\varepsilon K_\varepsilon(\mathbf{X})|_{\varepsilon=0} = \mathbf{V}(\mathbf{X}). \end{cases}$$

When \mathcal{M} is linear, a wide range of operators are available. We naturally have the simple Euler step, defined by $K_\varepsilon(\mathbf{X}_k) = \mathbf{X}_k + \varepsilon \mathbf{V}(\mathbf{X}_k)$, or the Crank-Nicolson scheme, implicitly defined by $K_\varepsilon(\mathbf{X}_k) = \mathbf{X}_k + \frac{\varepsilon}{2}(\mathbf{V}(\mathbf{X}_k) + \mathbf{V}(K_\varepsilon(\mathbf{X}_k)))$. On a nonlinear manifold, the first idea is to re-project on \mathcal{M} the point given by a step-forward operator defined on its embedding space. This operation is not always well-defined and consistent. Hopefully, it is sometimes possible to incorporate more efficiently the geometry of \mathcal{M} in the design of K_ε . Exploring this idea, the following developments are inspired from existing geometric integration methods for ODEs on Lie groups and homogeneous spaces [12, 17, 21, 24], and are related to the optimization techniques developed in [7, 14, 19, 32]. Most of these methods rely on the existence of closed forms for geodesics (or exponential maps) on matrix manifolds, generally expressed in terms of the matrix exponential.

⁴Recall that Ω is assumed to be flat in this paper. We refer to [5] for the implementation of PDEs acting on images whose domain is an implicit surface.

⁵Provided the continuous equation admits a sufficiently regular solution (Lipschitz regularity of \mathbf{V} or semi-group properties) Chorin et al. show that $\lim_{n \rightarrow \infty} K_{t/n} \circ \dots \circ K_{t/n} = F_t$ where F_t is the evolution operator such that $\mathbf{X}(t) = F_t(\mathbf{X}_0)$.

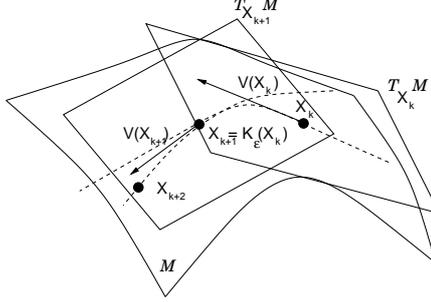


FIGURE 2. Two steps of a geometric numerical integrator on a nonlinear manifold \mathcal{M} .

Fig. 2 simply illustrates the action of a step-forward operator K_ϵ on a nonlinear manifold.

3.2. Geometric integrators for $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} , and \mathcal{F}_{iso} .

3.2.1. *Orthogonal flow.* We propose to define K_ϵ on $\mathcal{F}_{stiefel}$ from geodesic steps on $\text{St}(n, p)$. The equation of geodesics for the canonical metric on the Stiefel manifold has a closed form expressed in terms of the matrix exponential [14]. For a given point \mathbf{X} on $\text{St}(n, p)$, and a direction $\mathbf{H} = \mathbf{X}\mathbf{A} + \mathbf{X}_\perp\mathbf{B} \in T_{\mathbf{X}}\text{St}(n, p)$, this equation is given by

$$t \mapsto \text{Exp}_{\mathbf{X}}(t\mathbf{H}) = \begin{pmatrix} \mathbf{X} & \mathbf{X}_\perp \end{pmatrix} e^{t \begin{pmatrix} \mathbf{A} & -\mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{pmatrix}} \mathbf{I}_{n \times p}.$$

In practice, geodesics can be computed in $\mathcal{O}(np^2)$ flops with an algorithm given in [14]. We then build the operator

$$\forall \mathbf{p} \in \Omega, K_\epsilon(\mathbf{X}_k)(\mathbf{p}) = \begin{pmatrix} \mathbf{X}_k(\mathbf{p}) & \mathbf{X}_{k\perp}(\mathbf{p}) \end{pmatrix} e^{\epsilon \begin{pmatrix} \mathbf{A}(\mathbf{X}_k)(\mathbf{p}) & -\mathbf{B}(\mathbf{X}_k)(\mathbf{p})^T \\ \mathbf{B}(\mathbf{X}_k)(\mathbf{p}) & \mathbf{0} \end{pmatrix}} \mathbf{I}_{n \times p},$$

which provides a consistent approximation for a generic flow satisfying Eq. 2.1 on $\mathcal{F}_{stiefel}$. For equations acting on \mathcal{F}_{ortho} , the previous operator reduces to the simple expression

$$K_\epsilon(\mathbf{X}_k) = \mathbf{X}_k e^{\epsilon \mathbf{A}(\mathbf{X}_k)}.$$

A similar reduction occurs⁶ in the case $\mathcal{F}(\Omega, \text{S}^{n-1})$, for $\mathbf{H}_k = \mathbf{X}_{k\perp}\mathbf{B}(\mathbf{X}_k)$,

$$K_\epsilon(\mathbf{X}_k) = \begin{cases} \mathbf{X}_k \cos(\epsilon \|\mathbf{H}_k\|) + \frac{\mathbf{H}_k}{\|\mathbf{H}_k\|} \sin(\epsilon \|\mathbf{H}_k\|) & \text{if } \mathbf{H}_k \neq \mathbf{0}, \\ \mathbf{X}_k & \text{if } \mathbf{H}_k = \mathbf{0}. \end{cases}$$

This last result corresponds to the geodesic marching procedure already proposed in [8] for PDEs acting on direction fields.

3.2.2. *Rank/signature preserving flow.* In this case, we first need to introduce an auxiliary flow on $\mathcal{F}(\Omega, \text{GL}(n, \mathbb{R}))$. We identify the tangent space at a point \mathbf{Y} on $\text{GL}(n, \mathbb{R})$ with $T_{\mathbf{Y}}\text{GL}(n, \mathbb{R}) = \{\mathbf{Y}\mathbf{A} \mid \mathbf{A} \in \mathfrak{gl}(n, \mathbb{R}) = \mathbb{R}^{n \times n}\}$. Thus, the generic form of evolution equations on $\mathcal{F}(\Omega, \text{GL}(n, \mathbb{R}))$ is given by

$$\partial_t \mathbf{Y} = \mathbf{Y}\mathbf{A}(\mathbf{Y}) \mid \mathbf{A}(\mathbf{Y}) \in \mathcal{F}(\Omega, \mathbb{R}^{n \times n}).$$

⁶Note that \mathbf{B} is scalar in this case, and use the property $\exp \begin{pmatrix} 0 & -\alpha \\ \alpha & 0 \end{pmatrix} = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{pmatrix}$.

We also notice that the equation of the geodesic of tangent \mathbf{H} at the identity on $\text{GL}(n, \mathbb{R})$ is $t \mapsto e^{t\mathbf{H}}$ [18]. Using this result, and identifying any vector on $T_{\mathbf{Y}}\text{GL}(n, \mathbb{R})$ with a vector of $\mathfrak{gl}(n, \mathbb{R})$ (by left translation), we propose the following integrator L_ε on $\mathcal{F}(\Omega, \text{GL}(n, \mathbb{R}))$:

$$L_\varepsilon(\mathbf{Y}_k) = \mathbf{Y}_k e^{\varepsilon \mathbf{A}(\mathbf{Y}_k)}.$$

Once these tools are available, we introduce a constant field Λ of diagonal matrices $\text{diag}(\mathbf{I}_p, -\mathbf{I}_q, \mathbf{0})$ of size $n \times n$ (with p positive eigenvalues and q negative eigenvalues). We also define $\mathbf{X}(t) = \mathbf{Y}(t)^T \Lambda \mathbf{Y}(t)$ where $t \mapsto \mathbf{Y}(t) \in \mathcal{F}(\Omega, \text{GL}(n, \mathbb{R}))$ (note that such a decomposition exists for any family $t \mapsto \mathbf{X}(t) \in \mathcal{F}_{rank}$). If $t \mapsto \mathbf{Y}(t)$ is solution of

$$\partial_t \mathbf{Y} = \mathbf{Y} \mathbf{A}(\mathbf{X}) / \mathbf{Y}_0 \in \mathcal{F}(\Omega, \text{GL}(n, \mathbb{R})),$$

then $t \mapsto \mathbf{X}(t)$ satisfies the rank/signature preserving flow of Eq 2.2, with $\mathbf{X}_0 = \mathbf{Y}_0^T \Lambda \mathbf{Y}_0$. Looking at the discrete counterpart of this property yields a natural step-forward operator on \mathcal{F}_{rank} :

$$K_\varepsilon(\mathbf{X}_k) = L_\varepsilon(\mathbf{Y}_k)^T \Lambda L_\varepsilon(\mathbf{Y}_k),$$

which reduces to

$$\forall \mathbf{p} \in \Omega, K_\varepsilon(\mathbf{X}_k)(\mathbf{p}) = e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})^T} \mathbf{X}_k(\mathbf{p}) e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})}.$$

Note that any reference to the auxiliary field \mathbf{Y}_k disappears. We can verify the consistency of K_ε using well-known results of matrix calculus. Let $\mathbf{M}(\mathbf{p}) = e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})}$, the integration scheme becomes $\mathbf{X}_{k+1}(\mathbf{p}) = \mathbf{M}(\mathbf{p})^T \mathbf{X}_k(\mathbf{p}) \mathbf{M}(\mathbf{p})$. $\mathbf{M}(\mathbf{p})$ is a matrix exponential, thus nonsingular. Consequently, since congruence preserves the matrix inertia (the number of positive, negative and zero eigenvalues), $\mathbf{X}_{k+1}(\mathbf{p})$ and $\mathbf{X}_k(\mathbf{p})$ have the same rank and signature (Sylvester's law of inertia [20]). Furthermore, a direct derivation yields

$$\partial_\varepsilon K_\varepsilon(\mathbf{X})|_{\varepsilon=0} = \mathbf{A}(\mathbf{X})^T \mathbf{X} + \mathbf{X} \mathbf{A}(\mathbf{X}).$$

3.2.3. Isospectral flow. We use a similar mechanism for the isospectral flow. We consider the numerical approximation of an auxiliary problem on \mathcal{F}_{ortho} . Let $\mathbf{X}(t) = \mathbf{Y}(t)^T \mathbf{X}_0 \mathbf{Y}(t)$, such that $t \mapsto \mathbf{Y}(t) \in \mathcal{F}_{ortho}$ is solution of the orthogonal flow:

$$\partial_t \mathbf{Y} = \mathbf{Y} \mathbf{A}(\mathbf{X}), \quad \mathbf{Y}_0 = \mathbf{I}_n \iff \mathbf{Y}_{k+1} = L_\varepsilon(\mathbf{Y}_k), \quad \mathbf{Y}_0 = \mathbf{I}_n,$$

where $\mathbf{A}(\mathbf{X}) \in \mathcal{F}_{skew}$ and L_ε denotes the corresponding step-forward operator. A simple computation shows that $t \mapsto \mathbf{X}(t)$ is actually solution of an isospectral flow (Eq. 2.3). By analogy with the continuous equations, we propose to define the step-forward operator on \mathcal{F}_{iso} as

$$K_\varepsilon(\mathbf{X}_k) = L_\varepsilon(\mathbf{Y}_k)^T \mathbf{X}_0 L_\varepsilon(\mathbf{Y}_k),$$

which reduces to

$$\begin{aligned} \forall \mathbf{p} \in \Omega, K_\varepsilon(\mathbf{X}_k)(\mathbf{p}) &= e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})^T} \mathbf{X}_k(\mathbf{p}) e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})} \\ &= e^{-\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})} \mathbf{X}_k(\mathbf{p}) e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})}, \end{aligned}$$

since $\mathbf{A}(\mathbf{X}_k) \in \mathcal{F}_{skew}$. K_ε is consistent by construction: If we set $\mathbf{M}(\mathbf{p}) = e^{\varepsilon \mathbf{A}(\mathbf{X}_k)(\mathbf{p})}$, the integration scheme becomes $\mathbf{X}_{k+1}(\mathbf{p}) = \mathbf{M}(\mathbf{p})^{-1} \mathbf{X}_k(\mathbf{p}) \mathbf{M}(\mathbf{p}) = \mathbf{M}(\mathbf{p})^T \mathbf{X}_k(\mathbf{p}) \mathbf{M}(\mathbf{p})$. $\mathbf{M}(\mathbf{p})$ is nonsingular and orthogonal. Since similarity preserves the eigenvalues [20],

the discrete flow is isospectral. Moreover, the property of first order approximation in time follows from the Baker-Campbell-Hausdorff formula [1]:

$$e^{-\varepsilon\mathbf{A}}\mathbf{X}e^{\varepsilon\mathbf{A}} = \mathbf{X} + \varepsilon[\mathbf{X}, \mathbf{A}] + \frac{\varepsilon^2}{2}[[\mathbf{X}, \mathbf{A}], \mathbf{A}] + \dots$$

A summary of the previous numerical schemes is given in Table 2. One will notice that they all require the numerical evaluation of the matrix exponential. This problem is discussed in the appendix.

$\mathcal{F}_{stiefel}$	$\mathbf{K}_\varepsilon(\mathbf{X}_k)(\mathbf{p}) = \begin{pmatrix} \mathbf{X}_k(\mathbf{p}) & \mathbf{X}_{k_\perp}(\mathbf{p}) \end{pmatrix} e^{\varepsilon \begin{pmatrix} \mathbf{A}(\mathbf{X}_k)(\mathbf{p}) & -\mathbf{B}(\mathbf{X}_k)(\mathbf{p})^T \\ \mathbf{B}(\mathbf{X}_k)(\mathbf{p}) & \mathbf{0} \end{pmatrix}} \mathbf{I}_{n \times p}$
\mathcal{F}_{ortho}	$\mathbf{K}_\varepsilon(\mathbf{X}_k)(\mathbf{p}) = \mathbf{X}_k(\mathbf{p})e^{\varepsilon\mathbf{A}(\mathbf{X}_k)(\mathbf{p})}$
\mathcal{F}_{rank}	$\mathbf{K}_\varepsilon(\mathbf{X}_k)(\mathbf{p}) = e^{\varepsilon\mathbf{A}(\mathbf{X}_k)(\mathbf{p})^T} \mathbf{X}_k(\mathbf{p})e^{\varepsilon\mathbf{A}(\mathbf{X}_k)(\mathbf{p})}$
\mathcal{F}_{iso}	$\mathbf{K}_\varepsilon(\mathbf{X}_k)(\mathbf{p}) = e^{-\varepsilon\mathbf{A}(\mathbf{X}_k)(\mathbf{p})} \mathbf{X}_k(\mathbf{p})e^{\varepsilon\mathbf{A}(\mathbf{X}_k)(\mathbf{p})}$

TABLE 2. Geometric numerical integrators.

4. COST FUNCTIONAL MINIMIZATION

Using our previous developments, we now try to answer the following question: How to design a specific flow to achieve an image regularization or restoration task?

From the large literature on PDE-based techniques for scalar images (we refer to [3, 31, 44] for surveys on this topic), three strategies emerge: The equations come from local heuristics (geometry-driven diffusion), they are derived within an axiomatic framework (nonlinear scale-space), or they are obtained as minimizing flows of suitable cost functionals. The resulting PDEs usually include nonlinear diffusion operators. The idea is to allow a selective “edge-preserving” smoothing of the image.

In the matrix-valued case, a direct specialization of the first two approaches seems far from obvious, due to the specific form imposed by the constraints. For this reason, the cost functional minimization approach is preferred here. The goal is to build a minimizing flow for a functional providing some “image irregularity” measure. Assuming such a functional is given (its choice will be discussed later), we consider below the design of the corresponding flow.

4.1. Minimizing flows. Let us consider a cost functional $f : \mathcal{M} \rightarrow \mathbb{R}^+$ (which we assume sufficiently regular). From a suitable initial guess $\mathbf{X}(0) \in \mathcal{M}$ (the original image in the case of regularization problems), we build a minimizing flow by following the direction of steepest descent on \mathcal{M} . This direction is given, up to the choice of a metric, by the corresponding gradient vector field ∇f :

$$\partial_t \mathbf{X} = -\nabla f(\mathbf{X}), \quad \mathbf{X}(0) \in \mathcal{M}.$$

The gradient $\nabla f(\mathbf{X})$ is generally defined [1, 41] as the element of the tangent space $T_{\mathbf{X}}\mathcal{M}$ which satisfies⁷:

$$(4.1) \quad df_{\mathbf{X}}(\mathbf{H}) = g_{\mathbf{X}}(\nabla f(\mathbf{X}), \mathbf{H}), \quad \forall \mathbf{H} \in T_{\mathbf{X}}\mathcal{M},$$

where $g_{\mathbf{X}}(\cdot, \cdot)$ is an arbitrary Riemannian metric on \mathcal{M} , and $df_{\mathbf{X}}(\mathbf{H})$ denotes the first variation (first differential) of f in the direction \mathbf{H} . In practice, we can choose an arbitrary curve \mathbf{X}_{ε} on \mathcal{M} such that $\mathbf{X}_0 = \mathbf{X}$ and $\partial_{\varepsilon}\mathbf{X}_{\varepsilon}|_{\varepsilon=0} = \mathbf{H}$, and obtain the first variation as $df_{\mathbf{X}}(\mathbf{H}) = \partial_{\varepsilon}f(\mathbf{X}_{\varepsilon})|_{\varepsilon=0}$.

4.1.1. *Unconstrained gradient.* Given a functional f defined on the linear function space $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$, the idea is to express its first variation as follows:

$$(4.2) \quad df_{\mathbf{X}}(\mathbf{H}) = \int_{\Omega} \text{Trace}(\mathbf{G}(\mathbf{X})(\mathbf{p})^T \mathbf{H}(\mathbf{p})) \, d\mathbf{p}.$$

For this purpose, it will sometimes be necessary to assume that \mathbf{H} ($\in T_{\mathbf{X}}\mathcal{F}(\Omega, \mathbb{R}^{n \times p}) \simeq \mathcal{F}(\Omega, \mathbb{R}^{n \times p})$) vanishes on the domain boundary $\partial\Omega$, or impose homogeneous Neumann⁸ boundary conditions on the components of \mathbf{X} . Once \mathbf{G} has been identified, if we define $g_{\mathbf{X}}$ as a standard L^2 scalar product on $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$:

$$g_{\mathbf{X}}(\mathbf{U}, \mathbf{V}) = \langle \mathbf{U}, \mathbf{V} \rangle_{L^2} = \int_{\Omega} \text{Trace}(\mathbf{U}(\mathbf{p})^T \mathbf{V}(\mathbf{p})) \, d\mathbf{p},$$

we get $\nabla f = \mathbf{G}$. Note that in this case, the equality $\mathbf{G}(\mathbf{X}) = \mathbf{0}$ would correspond to the classical Euler-Lagrange equations of variational calculus.

4.1.2. *Natural constrained gradient.* By extension, we would like to obtain the natural gradient⁹ of the restriction of f to $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} and \mathcal{F}_{iso} . In all three cases, we need to identify the element $\nabla f(\mathbf{X})$ of $T_{\mathbf{X}}\mathcal{M}$ which satisfies Eq. 4.1, using the first variation of f when \mathbf{H} is restricted to $T_{\mathbf{X}}\mathcal{M}$ (now \mathbf{H} depends on \mathbf{X}). We also need to choose a suitable metric $g_{\mathbf{X}}$.

In fact, assuming the first variation of f is still given by Eq. 4.2, we can provide a systematic method to convert the unconstrained L^2 gradient of f on $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$ into the natural gradient of its restriction to $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} or \mathcal{F}_{iso} . Up to the change of metric, this operation is a projection of \mathbf{G} into the suitable tangent space¹⁰.

We choose to endow \mathcal{F}_{rank} , $\mathcal{F}_{stiefel}$ and \mathcal{F}_{iso} with metrics obtained by integrating over Ω existing canonical metrics on Lie groups and homogeneous spaces [14, 19, 33]. Our argument is that these manifolds also have an algebraic structure of group, or result from a group action, which can (and should) be taken into account through the choice of relevant metrics.

In each case we give below the expression of the corresponding gradient. We also detail its derivation in the orthogonal case.

⁷If $(T_{\mathbf{X}}\mathcal{M}, g_{\mathbf{X}})$ is a Hilbert space and $df_{\mathbf{X}}$ a continuous linear form on $T_{\mathbf{X}}\mathcal{M}$, the existence and uniqueness of $\nabla f(\mathbf{X})$ is a direct consequence of Riez representation theorem [1].

⁸ $\sum_k \partial_{\mathbf{p}_k} \mathbf{X}_{i,j}(\mathbf{p}) \cdot \mathbf{n}_k = 0, \quad \forall 1 \leq i \leq n, \quad 1 \leq j \leq p, \quad \forall \mathbf{p} \in \partial\Omega$, where \mathbf{n} is the unit outward normal vector on $\partial\Omega$.

⁹We borrow the terminology of *natural gradient* from Amari [2] who takes into account Riemannian metrics to build efficient minimizing flows in learning and information theory.

¹⁰This process corresponds to the use of Lagrange multipliers in classical variational calculus.

4.1.3. *Orthogonality constraint.* For two vectors $\mathbf{V}_1 = \mathbf{X}\mathbf{A}_1 + \mathbf{X}_\perp\mathbf{B}_1$, $\mathbf{V}_2 = \mathbf{X}\mathbf{A}_2 + \mathbf{X}_\perp\mathbf{B}_2 \in T_{\mathbf{X}}\mathcal{F}_{stiefel}$, we set

$$(4.3) \quad g_{\mathbf{X}}(\mathbf{V}_1, \mathbf{V}_2) = \int_{\Omega} \text{Trace}((\mathbf{V}_1(\mathbf{p}))^T (\mathbf{I}_n - \frac{1}{2}\mathbf{X}(\mathbf{p})\mathbf{X}(\mathbf{p})^T) \mathbf{V}_2(\mathbf{p})) \, d\mathbf{p} \\ = \frac{1}{2} \langle \mathbf{A}_1, \mathbf{A}_2 \rangle_{L^2} + \langle \mathbf{B}_1, \mathbf{B}_2 \rangle_{L^2},$$

using the canonical metric induced from the quotient space structure of $\text{St}(n, p)$. We refer to [14] for a justification of this choice. Now, given an arbitrary direction $\mathbf{H} = \mathbf{X}\mathbf{A} + \mathbf{X}_\perp\mathbf{B} \in \mathcal{F}_{stiefel}$, and using the properties of the trace operator, we can rewrite the first variation of f (Eq. 4.2) as

$$(4.4) \quad df_{\mathbf{X}}(\mathbf{X}\mathbf{A} + \mathbf{X}_\perp\mathbf{B}) = \frac{1}{2} \langle \{\mathbf{X}, \mathbf{G}(\mathbf{X})\}, \mathbf{A} \rangle_{L^2} + \langle \mathbf{X}_\perp^T \mathbf{G}(\mathbf{X}), \mathbf{B} \rangle_{L^2},$$

using the bracket notation $\{\mathbf{A}, \mathbf{B}\} = \mathbf{A}^T \mathbf{B} - \mathbf{B}^T \mathbf{A} \in \mathcal{F}_{skew}$. If we set $\nabla f(\mathbf{X}) = \mathbf{X}\mathbf{C} + \mathbf{X}_\perp\mathbf{D} \in T_{\mathbf{X}}\mathcal{F}_{stiefel}$, we can identify

$$(4.5) \quad g_{\mathbf{X}}(\nabla f(\mathbf{X}), \mathbf{X}\mathbf{A} + \mathbf{X}_\perp\mathbf{B}) = \frac{1}{2} \langle \mathbf{C}, \mathbf{A} \rangle_{L^2} + \langle \mathbf{D}, \mathbf{B} \rangle_{L^2}$$

with Eq. 4.4, for arbitrary $\mathbf{A} \in \mathcal{F}_{skew}$ and $\mathbf{B} \in \mathcal{F}(\Omega, \mathbb{R}^{(n-p) \times p})$. We then obtain

$$\begin{aligned} \nabla f(\mathbf{X}) &= \mathbf{X}\{\mathbf{X}, \mathbf{G}(\mathbf{X})\} + \mathbf{X}_\perp\mathbf{X}_\perp^T \mathbf{G}(\mathbf{X}) \\ &= \mathbf{X}\{\mathbf{X}, \mathbf{G}(\mathbf{X})\} + (\mathbf{I}_n - \mathbf{X}\mathbf{X}^T)\mathbf{G}(\mathbf{X}) \\ &= \mathbf{G} - \mathbf{X}\mathbf{G}^T \mathbf{X}. \end{aligned}$$

Thus, the natural gradient flow of f on $\mathcal{F}_{stiefel}$ is given by

$$\partial_t \mathbf{X} = \mathbf{X}\mathbf{G}^T \mathbf{X} - \mathbf{G}, \quad \mathbf{X}_0 \in \mathcal{F}_{stiefel}.$$

Note that the same equation was also derived in [39] using Lagrange multipliers.

4.1.4. *Prescribed rank/signature.* In this case, we extend to \mathcal{F}_{rank} the canonical metric on $\text{S}(r, n)$ discussed in [19]. The gradient computation, based on similar derivations also proposed in [19] (in the finite-dimensional case), leads to

$$\nabla f(\mathbf{X}) = (\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T)\mathbf{X}^2 + \mathbf{X}^2(\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T).$$

4.1.5. *Prescribed eigenvalues.* The same procedure, applied to \mathcal{F}_{iso} equipped with the metric inherited from its finite-dimensional counterpart, yields a *double-bracket* expression:

$$\nabla f(\mathbf{X}) = [\mathbf{X}, [\mathbf{X}, [\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T]]].$$

The corresponding gradient flows are summarized in Table 3.

$\mathcal{F}_{stiefel}$	$\partial_t \mathbf{X} = \mathbf{X}\mathbf{G}(\mathbf{X})^T \mathbf{X} - \mathbf{G}$
\mathcal{F}_{rank}	$\partial_t \mathbf{X} = -((\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T)\mathbf{X}^2 + \mathbf{X}^2(\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T))$
\mathcal{F}_{iso}	$\partial_t \mathbf{X} = [\mathbf{X}, [\mathbf{X}, -(\mathbf{G}(\mathbf{X}) + \mathbf{G}(\mathbf{X})^T)]]$

TABLE 3. Natural gradient flows.

4.2. Anisotropic regularization. As direct application of these results, we design a simple image-driven anisotropic regularization scheme.

4.2.1. ϕ -function formulation. Our goal is to build a regularizing flow that reduces the noise and the artifacts corrupting a matrix-valued mapping \mathbf{X}_0 (which belongs to $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} or \mathcal{F}_{iso}) while preserving its large global features (equivalent to the “edges” or object contours of gray-valued images). For a field $\mathbf{X} \in \mathcal{F}(\Omega, \mathbb{R}^{n \times p})$, we consider a functional relying on an increasing function $\phi : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ of the norm of its spatial variations:

$$(4.6) \quad f(\mathbf{X}) = \int_{\Omega} \phi(\|d\mathbf{X}(\mathbf{p})\|) \, d\mathbf{p},$$

with

$$\|d\mathbf{X}(\mathbf{p})\| = \left(\sum_{i=1}^k \text{Trace}(\partial_{\mathbf{p}_i} \mathbf{X}(\mathbf{p})^T \partial_{\mathbf{p}_i} \mathbf{X}(\mathbf{p})) \right)^{1/2}.$$

Regularization techniques derived from this type of functionals have proven their efficiency in the scalar case (see [3] and references therein). Computing the first variation of f , together with homogeneous Neumann boundary conditions on \mathbf{X} , yields:

$$(4.7) \quad \mathbf{G}(\mathbf{X})(\mathbf{p}) = - \sum_{i=1}^k \partial_{\mathbf{p}_i} \left(\frac{\phi'(\|d\mathbf{X}(\mathbf{p})\|)}{\|d\mathbf{X}(\mathbf{p})\|} \partial_{\mathbf{p}_i} \mathbf{X}(\mathbf{p}) \right).$$

For a suitable choice of ϕ , we can expect gradient flows based on this nonlinear differential operator to exhibit a robust anisotropic smoothing behavior (the convex function $\phi : s \mapsto 2\sqrt{1+s^2} - 2$ is used in the numerical experiments). Then, we consider a one-parameter family of matrix-valued mappings which is solution of

$$\partial_t \mathbf{X} = -\nabla f(\mathbf{X}), \quad \mathbf{X}(0) = \mathbf{X}_0,$$

the natural gradient flow of f on $\mathcal{F}_{stiefel}$, \mathcal{F}_{rank} or \mathcal{F}_{iso} . This initial value problem is simply obtained by replacing \mathbf{G} with the expression given in Eq. 4.7 in the corresponding equation of Table 3. In practice, unless a data-attachment term is added to f (in which case we look for the asymptotic solution $\mathbf{X}(+\infty)$), we set an arbitrary time $\tau > 0$ and define $\mathbf{X}(\tau)$ as solution of our regularization problem.

When we use the word *solution* in the previous discussion, we implicitly assume the well-posedness of our gradient flows (existence, uniqueness). However, it is important to notice that our derivations followed essentially from geometric arguments, and such a proof has *not* been established. To achieve this goal, additional assumptions on the regularity of our matrix-valued functions would be needed, and the properties of the functional should be carefully studied.

4.2.2. Toward intrinsic cost functionals. So far, we assumed the functional was defined on the embedding space $\mathcal{F}(\Omega, \mathbb{R}^{n \times p})$, regardless of the underlying constraints. Geometry and metrics of the constrained set of mappings were taken into account in the gradient flow computation and its numerical implementation, but not in the functional itself. As we will see in the next section, this approach already yields very satisfactory experimental results.

Of course, our framework would remain valid for more complex functionals that also include geometric and metric information. Energy functionals inspired from (p -)harmonic maps theory were suggested in [25, 35, 43]. For a mapping $\mathbf{X} \in$

$\mathcal{F}(\Omega, \mathcal{N})$, the corresponding energy would be similar to Eq. 4.6 with $\phi : s \mapsto s^p$, and a norm of the spatial variations induced by a Riemannian metric on \mathcal{N} . Along the same line of ideas, we can mention the Polyakov action functional and the Beltrami framework proposed in [23, 34]. These approaches could be extended to the matrix-valued mappings considered in this paper. In this perspective, we refer to [13, 40, 41] for existing theoretical works on harmonic maps into Lie groups and homogeneous spaces.

5. EXPERIMENTS

In this section, we illustrate the previous theoretical framework and experiment different ways of regularizing a field of symmetric positive definite matrices. For this purpose, we use successively the rank/signature preserving, orthogonal, and isospectral flows. The application of interest is the regularization of Diffusion Tensor MRI data (DT-MRI) in medical imaging. This technique allows to measure the motion of water molecules in the white matter of the brain. After estimation, we assume here that each voxel of a DT-MRI volume is modeled as a 3×3 symmetric positive definite matrix (the so-called *diffusion tensor*). Its eigenvalues give the diffusivities of water molecules along principal directions given by its eigenvectors. The idea is to reconstruct the underlying white matter fiber tracts by following at each voxel the direction given by the eigenvector of largest eigenvalue (the motion of water molecules, restricted by the axons, tends to follow the fiber tracts [46]). When regularizing this type of data, the objective is twofold:

- Remove the noise inherent to the acquisition/estimation process,
- Allow an easier retrieval of the global fiber bundle structures (to build models of cerebral connectivity).

Previous related works can be found in [11, 30, 38, 42, 46].

The initial configuration space of our problem is the set of mappings $\mathbf{X} : \Omega \mapsto \mathbb{P}(3)$, where $\mathbb{P}(3)$ denotes the set of 3×3 symmetric positive definite matrices. Note that $\mathbb{P}(n)$ is a convex half-cone [26]. Consequently, any regularizing PDE whose numerical implementation reduces to a positive linear combination of the original data will preserve the symmetry and positive definiteness of the field (a similar remark can be found in [45]). Here, we propose the alternative use of our geometric flows as a more efficient way to integrate the intrinsic properties of the underlying field. In particular, we want to emphasize the importance of the tensor orientations which are essential in the fiber reconstruction process.

In the following, we evaluate the qualitative behavior of the regularizing flows on a synthetic dataset and then apply our tools to a real DT-MRI volume.

5.1. Synthetic experiments. Let us consider a synthetic field $\mathbf{X} : \Omega \mapsto \mathbb{P}(3)$. A simple graphical representation is obtained in terms of ellipsoids whose radii and axis orientations are given respectively by the eigenvalues and the eigenvectors of the symmetric positive definite matrices (Fig. 3a). Then, we corrupt its eigenvectors¹¹ with noise to form a new field \mathbf{X}_0 (Fig. 3b).

¹¹A Gaussian noise was applied to the 6 independent coefficients of each matrix, followed by a spectral decomposition used to force the eigenvalues to their initial value.

1st method. Since $P(3)$ is one of the 4 connected components of $S(3,3)$ [19], we naturally propose to use the *rank/signature preserving flow* to regularize \mathbf{X}_0 by minimizing the cost functional f (Eq. 4.6) on \mathcal{F}_{rank} . Note that in all the experiments, we use $\phi : s \mapsto 2\sqrt{1+s^2} - 2$. We present the results after 40 iterations. The time step ε is always chosen empirically to ensure the stability of the discrete flow. In this case, the recovered field is clearly smoother (Fig. 3c) but presents an eigenvalue swelling effect. The information about the eigenvalues and the eigenvectors is mixed by the regularizing flow. There is a risk of losing the principal direction of the initial field.

2nd method. To cope with this issue, the first idea is to perform a spectral decomposition of \mathbf{X}_0 , such that $\forall \mathbf{p} \in \Omega$, $\mathbf{X}_0(\mathbf{p}) = \mathbf{U}(\mathbf{p})^T \text{diag}(\lambda_1(\mathbf{p}), \lambda_2(\mathbf{p}), \lambda_3(\mathbf{p}))\mathbf{U}(\mathbf{p})$, and regularize its eigenvalues (diffusivities) and eigenvector matrices (orientation feature) separately. As explained in [38], this approach requires an additional alignment step due to the non-uniqueness of the spectral decomposition (eigenvectors of similar orientation with opposite directions) which can create artificial “discontinuities” disturbing the regularization of \mathbf{U} . Assuming this procedure was applied, \mathbf{U} ($\in \mathcal{F}_{ortho}$) is regularized by minimizing f on \mathcal{F}_{ortho} with the corresponding *orthogonal flow*. In this case, the initial field \mathbf{X} is almost perfectly recovered (Fig. 3d), including the sharp variation (the “edge”) observed on its orientations.

3rd method. A simple alternative to the previous method is the application of the *isospectral flow*. It allows us to directly regularize \mathbf{X}_0 (by minimizing f on \mathcal{F}_{iso}) while preserving the initial diffusivities. In this case, there is no need for an explicit spectral decomposition and a realignment procedure. The result is also very satisfactory (Fig. 3b). Note that despite similar results, the second and third methods are quite different. The *edge preserving* property of the regularization functional depends respectively on the spatial variations of the orientation field \mathbf{U} and \mathbf{X}_0 .

5.2. DT-MRI regularization. Now, we test our regularization PDEs on a $128 \times 128 \times 56$ DT-MRI image of a human brain¹² (Fig. 4). To make visualization easier, our results are presented on a small region of a slice of the initial volume (white square area of Fig. 4a). The left side of each figure shows the DT-MRI data as a field of ellipsoids, while the right side corresponds to streamlines of the principal direction field. The results highlight properties already observed with the synthetic experiments:

- The rank/signature preserving flow (Fig. 4b) tends to blend the orientation and diffusivity features (eigenvalue swelling effect). We quickly lose the structure of the underlying network of fiber tracts.
- The orthogonal flow applied to the field of eigenvectors (Fig. 4c) performs an efficient selective smoothing of the orientation feature, but requires a spectral decomposition and a realignment step. The fiber tracts network is simplified but keeps its main structures.
- The isospectral flow (Fig. 4d) exhibits a quite similar behavior. Moreover it does not require an explicit tensor decomposition, thus reducing the computational burden.

¹²The authors would like to thank J.-F. Mangin and J.-B. Poline (SHFJ-CEA) for the DT-MRI data (this work was partially supported by ARC MC2). We also thank R. Fournier for his visualization tool *TensView*.

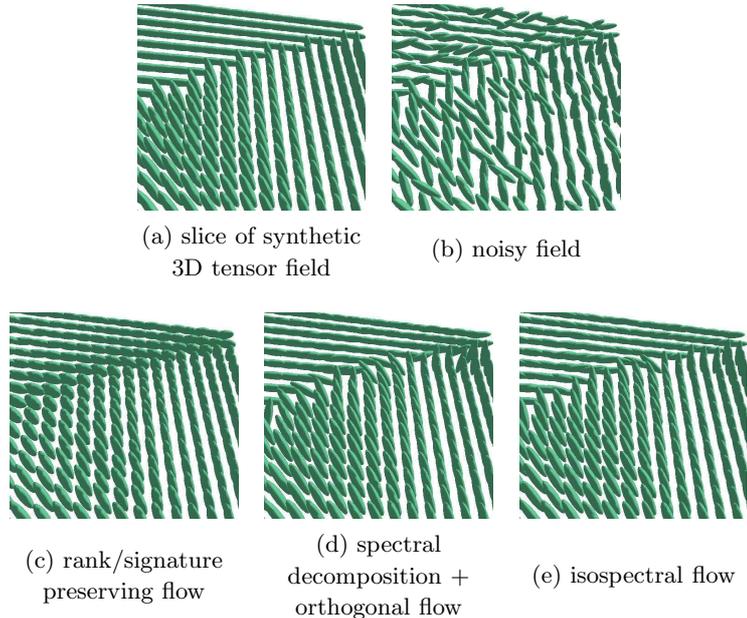


FIGURE 3. Synthetic experiments.

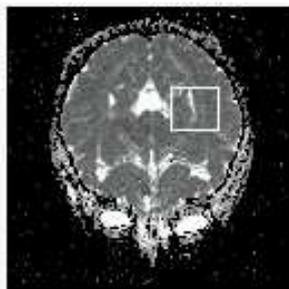
In Fig. 5, we show the action of two of the previous methods on a part of the dataset that corresponds to the *corpus callosum*, a bundle of nerve fibers that connects the two hemispheres of the brain. In this particular region, the fiber tracts are known to lie almost entirely on axial planes. The isospectral regularization approach naturally enhances this property, thanks to a coherent tensor orientation smoothing.

A complete validation and physiological interpretation of these results remains to be done, but these preliminary experiments seem to show that our regularizing flows reduce the level of noise and yield a smoother and more coherent model of the fiber tracts structures.

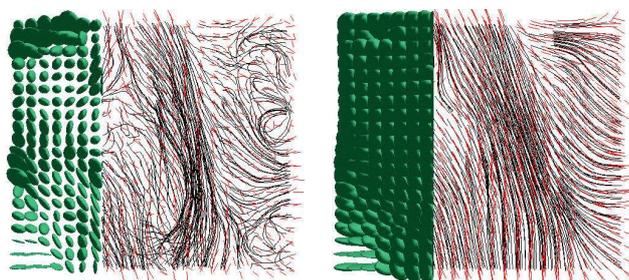
6. CONCLUSION

Our objective was to introduce a set of tools and a systematic approach to design and implement regularizing PDEs acting on constrained multi-valued images. The suitable flows were obtained by considering the geometry of constrained sets of matrix-valued mappings, which inherit most of their properties from finite dimensional Lie groups and homogeneous spaces. The corresponding structure-preserving numerical methods were also proposed, based on existing geometric integration techniques on Riemannian manifolds. The efficiency and versatility of this formalism was demonstrated on the problem of DT-MRI regularization.

We believe that our geometric approach gives new perspectives on some of the ideas previously developed in the field of PDE-based methods for multi-valued images [8, 23, 29, 43, 38, 39]. At this point, our derivations are essentially formal, leaving out numerous technicalities that arise when working with infinite-dimensional sets of mappings (involving both functional analysis and differential geometry).

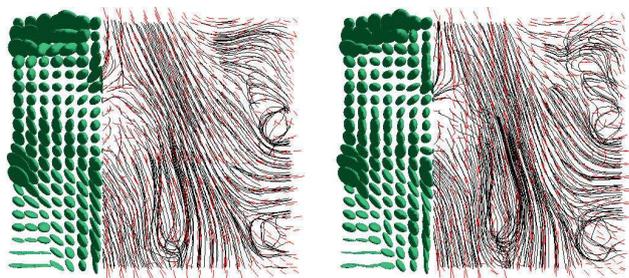


(a) slice of a DT-MRI of the brain (mean diffusivity)



(b) local tensors/streamlines corresponding to the white square area

(c) rank/signature preserving flow



(d) spectral decomposition + orthogonal flow

(e) isospectral flow

FIGURE 4. DT-MRI regularization.

A rigorous study of both the well-posedness of the resulting equations and the convergence of the geometric integrators shall form the next step of this work.

7. APPENDIX

Approximation of the matrix exponential. The numerical schemes proposed in this paper require the evaluation of matrix exponentials. The truncated power series representation $e^{\mathbf{B}} \simeq \sum_{k=0}^p \mathbf{B}^k/k!$ is hardly tractable in practice due to its slow convergence. Instead, we use the rational (Padé) approximation with scaling

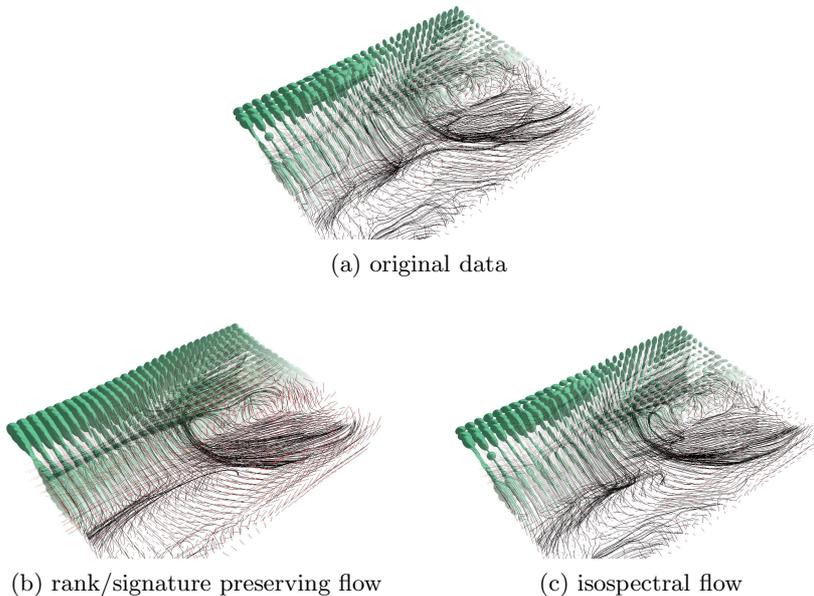


FIGURE 5. DT-MRI regularization (corpus callosum area).

and squaring proposed in [16]. This method relies on a two-parameter family of approximants $e^{\mathbf{B}} \simeq \mathcal{R}_{p,q}(\mathbf{B})$, with $\mathcal{R}_{p,q}(\mathbf{B}) = \mathcal{D}_{p,q}(\mathbf{B})^{-1} \mathcal{N}_{p,q}(\mathbf{B})$ such that

$$\mathcal{D}_{p,q}(\mathbf{B}) = \sum_{k=0}^p \frac{(p+q-k)!p!}{(p+q)!k!(p-k)!} \mathbf{B}^k, \quad \mathcal{N}_{p,q}(\mathbf{B}) = \sum_{k=0}^q \frac{(p+q-k)!p!}{(p+q)!k!(q-k)!} (-\mathbf{B})^k.$$

When dealing with skew-symmetric matrices, other techniques are considered (offering a lower computational cost). In particular, we use a reformulation in terms of trigonometric functions (Rodrigues' formula) or replace the exponential with the Cayley map, as detailed below.

Exponential of skew-symmetric matrices and Rodrigues' formula. Rodrigues' formula [1] gives a closed form of the matrix exponential of 3×3 skew-symmetric matrices in terms of trigonometric functions. We give here its generalization to $n \geq 3$ proposed by Gallier and Xu in [15]:

Theorem (Generalized Rodrigues' formula). *Given any non null skew-symmetric $n \times n$ ($n \geq 3$) matrix \mathbf{A} , if $\{i\theta_1, -i\theta_1, \dots, i\theta_p, -i\theta_p\}$ is the set of distinct eigenvalues of \mathbf{A} , where each $i\theta_1$ (and $-i\theta_1$) has multiplicity $k_j \geq 1$, there are p unique skew-symmetric matrices $\mathbf{A}_1, \dots, \mathbf{A}_p$ such that $\forall 1 \leq i, j \leq p$,*

$$\begin{aligned} \mathbf{A} &= \theta_1 \mathbf{A}_1 + \dots + \theta_p \mathbf{A}_p \\ \mathbf{A}_i \mathbf{A}_j &= \mathbf{A}_j \mathbf{A}_i = \mathbf{0} \quad (i \neq j) \\ \mathbf{A}_i^3 &= -\mathbf{A}_i \end{aligned}$$

Furthermore,

$$e^{\mathbf{A}} = \mathbf{I}_n + \sum_{i=1}^p (\sin \theta_i \mathbf{A}_i + (1 - \cos \theta_i) \mathbf{A}_i^2).$$

In particular, for $\mathbf{A} = \begin{pmatrix} 0 & -c & b \\ c & 0 & -a \\ -b & a & 0 \end{pmatrix}$, we have $p = 1$, $\theta_1 = \sqrt{a^2 + b^2 + c^2}$, and $\mathbf{A}_1 = \mathbf{A}/\theta_1$.

Cayley Map. The Cayley map (or Cayley transform) [17, 21] is the mapping

$$\text{Cay} : \mathbf{A} \mapsto (\mathbf{I}_n - \rho\mathbf{A})^{-1}(\mathbf{I}_n + \rho\mathbf{A}),$$

where $\rho \in \mathbb{R} \setminus \{0\}$ is an arbitrary constant. The Cayley map is well-defined, i.e. $\mathbf{I}_n - \rho\mathbf{A}$ is nonsingular, for all skew-symmetric matrix \mathbf{A} (since skew-symmetric matrices have complex eigenvalues), and is a computationally efficient alternative to matrix exponentials in some of our numerical schemes. This mapping shares two key properties with the matrix exponential. First, for a $n \times n$ skew-symmetric matrix \mathbf{A} , $\text{Cay}(\mathbf{A}) \in \text{SO}(n)$ (this follows from the skew-symmetry of \mathbf{A} and the commutativity of the product $(\mathbf{I}_n + \rho\mathbf{A})(\mathbf{I}_n - \rho\mathbf{A})$). Moreover, $\partial_\varepsilon \text{Cay}(\varepsilon\mathbf{A})|_{\varepsilon=0} = \partial_\varepsilon e^{\varepsilon\mathbf{A}}|_{\varepsilon=0} = \mathbf{A}$ for $\rho = 1/2$ (this result is obtained using the chain rule of derivation and the differential of the matrix inverse $d(\mathbf{X}^{-1})(\mathbf{X}) \cdot \mathbf{H} = -\mathbf{X}^{-1}\mathbf{H}\mathbf{X}^{-1}$). These results show that one can use Cayley maps instead of exponentials in the orthogonal and isospectral step-forward operators. Their consistency with the continuous flows is preserved. Note that for $\rho = 1/2$, the Cayley map is equal to the Padé approximant $\mathcal{R}_{1,1}(\mathbf{A})$.

REFERENCES

1. R. Abraham, J.E. Marsden, and T.S. Ratiu, *Manifolds, tensors, analysis, and applications*, Springer, 1988, Draft 3rd edition 2001 from http://www.cds.caltech.edu/~marsden/bib_src/mta/Book/.
2. S.-I. Amari, *Natural gradient works efficiently in learning*, Neural Computation **10** (1998), 251–276.
3. G. Aubert and P. Kornprobst, *Mathematical problems in image processing*, Applied Mathematical Sciences, vol. 147, Springer, 2001.
4. F. Barbaresco, *Spatial denoising of statistical parameters estimation by Beltrami diffusion on embedding Siegel space*, PSIP'2003 - Physics in Signal & Image Processing (Grenoble, France), January 2003.
5. M. Bertalmio, G. Sapiro, L-T. Cheng, and S. Osher, *Variational problems and PDE's on implicit surfaces*, in Paragios [28].
6. W.M. Boothby, *An introduction to differentiable manifolds and riemannian geometry*, 2nd ed., Academic Press, 1986.
7. C.A. Botsaris, *Constrained optimization along geodesics*, Journal of Mathematical Analysis and Applications **79** (1981), 295–306.
8. T. Chan and J. Shen, *Variational restoration of non-flat image features: models and algorithms*, SIAM J. Appl. Math. **61** (2000), no. 4, 1338–1361.
9. C. Chef'd'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras, *Constrained flows of matrix-valued functions: Application to diffusion tensor regularization*, Springer, May 2002, pp. 251–265.
10. A. Chorin, T. Hugues, M. McCracken, and J. Marsden, *Product formulas and numerical algorithms*, Communications on Pure and Applied Mathematics **31** (1978), 205–256.
11. O. Coulon, D.C. Alexander, and S.R. Arridge, *A regularization scheme for diffusion tensor magnetic resonance images*, International Conference on Information Processing in Medical Imaging (Davis, USA), 2001, pp. 92–105.
12. P.E. Crouch and R. Grossman, *Numerical integration of ordinary differential equations on manifolds*, J. Nonlinear Sci. **3** (1993), 1–33.
13. Y.-J. Dai, M. Shoji, and H. Urakawa, *Harmonic maps into Lie groups and homogeneous spaces*, Differential Geometry and its Applications **7** (1997), 143–160.
14. A. Edelman, T.A. Arias, and S.T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl. **20** (1998), no. 2, 303–353.

15. J. Gallier and D. Xu, *Computing exponentials of skew symmetric matrices and logarithms of orthogonal matrices*, To appear in the International Journal of Robotics and Automation, 2002.
16. G. Golub and C. Van Loan, *Matrix computations*, 2nd ed., North Oxford Academic, 1986.
17. E. Hairer, C. Lubich, and G. Wanner, *Geometric numerical integration: Structure preserving algorithms for ordinary differential equations*, Springer Series in Computational Mathematics, vol. 31, Springer, 2002.
18. S. Helgason, *Differential geometry, Lie groups, and symmetric spaces*, Academic Press, 1978.
19. U. Helmke and J.B. Moore, *Optimization and dynamical systems*, Communications and Control Engineering Series, Springer, 1994.
20. R. Horn and C. Johnson, *Matrix analysis*, Cambridge University Press, 1985.
21. A. Iserles, H.Z. Munthe-Kaas, S.P. Nørsett, and A. Zanna, *Lie-group methods*, Acta Numerica (2000), 215–365.
22. V. Kac (ed.), *Infinite dimensional Lie groups with applications*, Mathematical Sciences Research Institute Publications, vol. 4, Springer, 1985.
23. R. Kimmel and N. Sochen, *Orientation diffusion or how to comb a porcupine?*, Special issue on PDEs in Image Processing, Computer Vision, and Computer Graphics, Journal of Visual Communication and Image Representation **13** (2002), 238–248.
24. D. Lewis and P.J. Olver, *Geometric integration algorithms on homogeneous manifolds*, Preprint, 2001.
25. F. Mémoli, G. Sapiro, and S. Osher, *Solving variational problems and partial differential equations mapping into general target manifolds*, Computational and applied mathematics report, UCLA, 2002.
26. R. Mneimné and F. Testard, *Introduction à la théorie des groupes de Lie classiques*, Hermann, 1986.
27. S. Osher and R. Fedkiw, *Level set methods and dynamic implicit surfaces*, Springer, 2002.
28. N. Paragios (ed.), *Proceedings of the IEEE Workshop on Variational and Level Set Methods in Computer Vision*, Vancouver, Canada, IEEE Computer Society, July 2001.
29. P. Perona, *Orientation diffusion*, IEEE Transactions on Image Processing **7** (1998), no. 3, 457–467.
30. C. Poupon, *Détection des faisceaux de fibres de la substance blanche pour l'étude de la connectivité anatomique cérébrale*, Ph.D. thesis, ENST, 1999.
31. G. Sapiro, *Geometric partial differential equations and image analysis*, Cambridge University Press, 2001.
32. S.T. Smith, *Geometric optimization for adaptative filtering*, Ph.D. thesis, Harvard University, 1993.
33. ———, *Optimization techniques on Riemannian manifolds*, Fields Institute Communications, vol. 3, pp. 113–146, American Mathematical Society, Providence, RI, 1994.
34. N. Sochen, R. Kimmel, and R. Malladi, *A general framework for low level vision*, IEEE Transactions on Image Processing **17** (1998), no. 3, 310–318.
35. B. Tang, G. Sapiro, and V. Caselles, *Diffusion of general data on non-flat manifolds via harmonic maps theory: The direction diffusion case*, International Journal of Computer Vision **36** (2000), no. 2, 149–161.
36. B. ter Haar Romeny (ed.), *Geometry driven diffusion in computer vision*, Computational Imaging and Vision, Kluwer Academic Publishers, 1994.
37. A. Trounev, *Diffeomorphisms groups and pattern matching in image analysis*, International Journal of Computer Vision **28** (1998), no. 2, 213–221.
38. D. Tschumperlé and R. Deriche, *Diffusion tensor regularization with constraints preservation*, Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2001), IEEE Computer Society, 2001.
39. ———, *Orthonormal vector sets regularization with PDE's and applications*, International Journal of Computer Vision **50** (2002), no. 3, 237–252.
40. K. Uhlenbeck, *Harmonic maps into Lie groups (classical solutions of the chiral model)*, J. Differential Geometry **30** (1989), 1–50.
41. H. Urakawa, *Calculus of variations and harmonic maps*, Translations of Mathematical Monographs, no. 132, American Mathematical Society, 1993.
42. B. Vemuri, Y. Chen, M. Rao, T. McGraw, T. Mareci, and Z. Wang, *Fiber tract mapping from diffusion tensor MRI*, in Paragios [28].

43. L.A. Vese and S.J. Osher, *Numerical methods for p -harmonic flows and applications to image processing*, Computational and applied mathematics report, UCLA, 2001.
44. J. Weickert, *Anisotropic diffusion in image processing*, Teubner-Verlag, 1998.
45. J. Weickert and T. Brox, *Diffusion and regularization of vector- and matrix-valued images*, Preprint 58, Universität des Saarlandes, 2002.
46. C.-F. Westin, S.E. Maier, H. Mamata, A. Nabavi, F.A. Jolesz, and R. Kikinis, *Processing and visualization for diffusion tensor MRI*, Medical Image Analysis **6** (2002), no. 2, 93–108.

ODYSSÉE LAB., INRIA SOPHIA-ANTIPOLIS, FRANCE

E-mail address: `Christophe.Chef_d_Hotel@sophia.inria.fr`