J Math Imaging Vis 24: 349–358, 2006

DOI: 10.1007/s10851-005-3633-z

Some Differential-Geometric Remarks on a Method for Minimizing Constrained Functionals of Matrix-Valued Functions

DAVID GROISSER
Department of Mathematics, University of Florida, Gainesville, FL 32611-8105, USA
groisser@math.ufl.edu

Published online: 28 January 2006

Abstract. In [3] an approach is given for minimizing certain functionals on certain spaces $\mathcal{N} = \text{Maps}(\Omega, N)$, where Ω is a domain in some Euclidean space and N is a space of square matrices satisfying some extra condition(s), e.g. symmetry and positive-definiteness. The approach has the advantage that in the associated algorithm, the preservation of constraints is built in automatically. One practical use of such an algorithm its its application to diffusion-tensor imaging, which in recent years has been shown to be a very fruitful approach to certain problems in medical imaging. The method in [3] is motivated by differential-geometric considerations, some of which are discussed briefly in [3] and in greater detail in [4]. We describe here certain geometric aspects of this approach that are not readily apparent in [3] or [4]. We also discuss what one can and cannot hope to achieve by this approach.

Keywords: diffusion tensor imaging, matrix flow, constrained flow

1. Introduction

In [3] an approach is given for minimizing certain functionals on certain spaces $\mathcal{N} = \text{Maps}(\Omega, N)$, where Ω is a domain in some Euclidean space and N is a space of square matrices satisfying some extra condition(s), e.g. symmetry and positive-definiteness. ("Map" in this paper means "smooth map" unless otherwise specified.) The approach has the advantage that in the associated algorithm, the preservation of constraints is built in automatically; one does not have to worry about "stepping off" the constraint manifold and projecting back onto it.

One practical use of such an algorithm is its application to diffusion-tensor imaging, which in recent years has been a very fruitful approach to certain problems in medical imaging [5–9]. In this method one uses the diffusion tensor field, a positive-definite symmetric 3×3 matrix varying from point to point in the region of anatomical interest, to help determine the connections between tissues in the region; microstructural connections between a point and its neighbors are encoded by the direction corresponding to the largest eigenvalue at that point. To obtain this information one must restore (or estimate) the "true" diffusion tensor field from usually noisy measurements. Often this estimate is accomplished by a variational approach [10, 11] in which the object one can vary is a 3×3 matrix of functions, subject to the constraints that at each point the matrix be symmetric and strictly positive-definite.

In [3] Chefd'hotel et al. describe a variational method (a constrained flow) motivated by differential-geometric considerations, some of which are discussed briefly in [3] and in greater detail in [4]. The approach of Chefd'hotel et al. applies to many different constraints; in [3] the examples of orthogonal constraints and isospectral constraints ([3, Sections 2.3 and 2.4]) are considered in a addition to the symmetric positive-definite constraint ([3, Sections 2.2 and 2.4]). In the present paper we discuss certain geometric aspects of this approach that are not readily apparent in [3] or [4]. We also discuss what one can and cannot hope to achieve by this approach.

When Ω and N have positive dimension, Maps(Ω, N) is an infinite-dimensional object. There are several technical issues related to infinite-dimensionality that can divert attention from the basic underlying geometric principles in the method of [3], so, for geometric clarity, at first we will take $\Omega = \{\text{point}\}$. Of course, when Ω is a single point, $\mathcal{N} = \text{Maps}(\Omega, N)$ can be naturally identified with N itself.

Henceforth let N be a (finite-dimensional) manifold and V a continuous vector field on N. (Note: in this paper, "manifold" always means "manifold without boundary.") Consider the equation for the flow of V:

$$\frac{dx}{dt} = V(x(t)), \quad x(0) = x_0. \tag{1.1}$$

The equilibrium solutions of (1.1) are precisely the zeroes of V, and given any solution of (1.1) for which $\lim_{t\to\infty} x(t)$ exists, the limit is a zero of V.

Suppose now that $\mathbf{L}$ is a section of $\text{Aut}(TN)$; i.e. for each $x \in N, \mathbf{L}_x := \mathbf{L}(x)$ is an invertible linear transformation from T_xN to itself, varying continuously with x. Then the assignment $x \mapsto \mathbf{L}_x(V(x))$ defines a new continuous vector field $\mathbf{L}(V)$, and we can consider the following modification of (1.1):

$$\frac{dx}{dt} = \mathbf{L}(V)(x(t)) \ . \tag{1.2}$$

Because $\mathbf{L}$ is invertible at every point, the zeroes of $\mathbf{L}(V)$ are exactly the zeroes of V. Hence (1.2) has the same equilibrium solutions as (1.1), and for any solution of (1.2) that converges as $t \to \infty$, the limit is again a zero of V.

An important special case of (1.1) occurs when we are given a C^1 function ("energy") $E : N \to \mathbf{R}$ and N is equipped with a Riemannian metric g_N. In this case we can consider the negative-gradient flow

$$\frac{dx}{dt} = -(\text{grad}\, E)|_{x(t)}, \quad x(0) = x_0, \tag{1.3}$$

or, as in (1.2), a related equation of the form

$$\frac{dx}{dt} = -\mathbf{L}(\text{grad}\, E)|_{x(t)}, \quad x(0) = x_0, \tag{1.4}$$

for some well-chosen $\mathbf{L}$, and attempt to locate critical points of E by following the flow as $t \to \infty$. The equilibria of (1.3) and (1.4) are identical and correspond to these critical points.

The replacement of (1.3) with (1.4), in an infinite-dimensional setting, is what underlies the method presented in [3]. Although the motivation in [3] is to handle certain constrained variational problems, we will see below that the method amounts to passing from (1.3) to (1.4) by "factoring" a variational problem through another manifold, which leads to (1.4) with a very particular form for $\mathbf{L}$. The operators $\mathbf{L}_x$ that arise from the method in [3] are not merely invertible, but symmetric and *positive-definite* with respect to g_N, so that the energy is still a decreasing function along the trajectories of (1.4); in fact we can recast (1.4) as the negative-gradient flow of E on N with respect to a modified metric g'_N, defined by

$$g'_N(X, Y) = g_N(X, (\mathbf{L}_x)^{-1}Y), \quad X, Y \in T_xN. \tag{1.5}$$

Below, we show explicitly where the operators $\mathbf{L}_x$ come from (giving systematically the *change of metric* referred to in [3] and [4]), and compute them in the situations relevant to [3].

We use differential-geometric notation and terminology in this paper; e.g. if M is a manifold then T_pM is the tangent space of M at p, and if Φ is a map with domain M, the derivative of Φ at p is denoted Φ_{*p}. For readers who are not at home in this language, we have included a glossary in the Appendix.

2. The Geometric Setup

While the setup in [3] uses only operators $\mathbf{L}$ in (1.4) that arise from a change of metric on N, one of the motivating examples leads to a more general idea that we consider first.

Assume that (N, g_N) is Riemannian manifold and that $E_N : N \to \mathbf{R}$ is a C^1 function. Suppose we are also given a second Riemannian manifold (M, g_M) and a submersion $\pi : M \to N$, i.e. a smooth map whose derivative $\pi_{*p} : T_pM \to T_{\pi(p)}N$ is surjective for all $p \in M$. We assume further that π is surjective, an assumption that is redundant if M is compact and N connected (because submersions map open sets to open sets, by the Implicit Function Theorem). If π is one-to-one, we can regard it as simply providing a reparametrization of N. More generally—e.g. if $\dim(M) > \dim(N)$—we can think of π as an "over-parametrization" of N, in which each point $x \in N$ is represented by the manifold $\pi^{-1}(x)$ of dimension $\dim(M) - \dim(N)$. If we replace the surjectivity of the derivative with a slightly stronger technical hypothesis, "local triviality", as we will do for simplicity henceforth, our assumptions above amount simply to the statement that M is a fiber bundle over N with projection map π (see [12, 16], or our Appendix).

Let $E_M = \pi^*E_N = E_N \circ \pi$ denote the pullback of E_N to M. Because M, N are Riemannian, the gradient vector-fields $\mathrm{grad}_M E_M$, $\mathrm{grad}_N E_N$ are defined on M, N respectively. To express the explicit relation of $\mathrm{grad}_M E_M$ to $\mathrm{grad}_N E_N$ we recall the definition of the *adjoint* of a linear map between two inner-product spaces. If W_1, W_2 are vector spaces with inner products $(\cdot,\cdot)_1, (\cdot,\cdot)_2$ respectively, and $T : W_1 \to W_2$ is a linear map, then we say that a linear map $S : W_2 \to W_1$ is adjoint to T if for all $w_1 \in W_1, w_2 \in W_2$ we have $(w_2, Tw_1)_2 = (Sw_2, w_1)_1$. When adjoints exist, as they always do when W_1, W_2 are finite-dimensional, they are necessarily unique. Hence in these cases we speak of *the* adjoint of T, which we denote $T^\dagger$.

Of interest to us are the linear maps $\pi_{*p} : T_pM \to T_{\pi(p)}N$ (the derivative of π at p) and their adjoints $\pi_{*p}^\dagger : T_{\pi(p)}N \to T_pM$. Below we let $\langle\cdot,\cdot\rangle$ denote the pairing between a vector space and its dual.

Lemma 2.1. *For all $p \in M$ we have*

$$(\mathrm{grad}_M E_M)|_p = \pi_{*p}^\dagger((\mathrm{grad}_N E_N)|_{\pi(p)}). \tag{2.1}$$

Proof: It suffices to check that both sides of (2.1) have the same inner product with all $Y \in T_pM$. But

$$\begin{aligned} g_M(\mathrm{grad}_M E_M; Y)_p &= \langle d(\pi^* E_N)|_p, Y\rangle \\ &= \langle (\pi^* dE_N)_P, Y\rangle = \langle dE_N, \pi_{*p}Y\rangle \\ &= g_N(\mathrm{grad}_N E_N, \pi_{*p}Y)_{\pi(p)} \\ &= g_M(\pi_{*p}^\dagger(\mathrm{grad}_N E_N), Y)_p. \end{aligned}$$

□

Note that for each $p \in M$, since π_{*p} is assumed surjective, the adjoint $\pi_{*p}^\dagger$ is injective. Also observe that (2.1) implies that $(\mathrm{grad}_M E_M)|_p$ is orthogonal to the fiber $\pi^{-1}(\pi(p))$, since the tangent space to the fiber at p is exactly $\ker(\pi_{*p})$. The injectivity of $\pi_{*p}^\dagger$ immediately implies the following.

Corollary 2.2. *$\mathrm{grad}_M E_M$ vanishes at $p_0 \in M$ if and only if $\mathrm{grad}_N E_N$ vanishes at $\pi(p_0) \in N$. Hence p_0 is an equilibrium solution of the negative-gradient flow of E_M on M if and only if $\pi(p_0)$ is an equilibrium solution of the negative-gradient flow of E_N on N.*

Thus if the negative-gradient flow of E_M converges, the limit of a trajectory will project to a critical point of E_N.

In the method used in [3], one does not explicitly look at the flow on M; rather one projects it back down to N—which can only be done under certain circumstances described shortly—and obtains a new flow on N whose trajectories one then attempts to follow.

Note that in general, the flow of $-\mathrm{grad}_M E_M$ does not project to a well-defined flow on N unless π is one-to-one, because trajectories of $-\mathrm{grad}_M E_M$ starting at two different pre-images of $x_0 \in N$ may project down to two different curves in N. There are two ways of dealing with this issue: (i) ensure, via further assumptions, that the vector field $\tilde{V} = -\mathrm{grad}_M E_M$ is *projectable*, i.e. that if $\pi(p_1) = \pi(p_2)$ then $\pi_{*p_1}(\tilde{V}_{p_1}) = \pi_{*p_2}(\tilde{V}_{p_2})$, or (ii) for each $x_0 \in N$, single out a pre-image $s(x_0) \in \pi^{-1}(x_0)$, and project the trajectory of $-\mathrm{grad}_M E_M$ starting at $s(x_0)$ back down to N. The second approach is not of practical utility. One would want $s(x_0)$ to depend continuously on x_0; i.e. for s to be a continuous section of the fiber bundle (M, N, π). Such sections often do not exist; for principal fiber bundles, a category that includes all the examples considered in Section 3 of this paper (corresponding to [3, Sections 2.2–2.4]) they exist only when the bundle is topologically trivial (a Cartesian product $N \times F$)—and in this case, one might as well replace M with the image $M' := s(N)$ and work with the one-to-one map $\pi|_{M'} : M' \to N$ (with a certain induced metric on M'). Thus we are led to approach (i), which is what underlies the method in [3, Sections 2.2 and 2.4] (in [3, Section 2.3], the map π is one-to-one from the start, and approaches (i) and (ii) above coincide).

To ensure that the vector field $-\mathrm{grad}_M E_M$ is projectable, we henceforth make the simplifying assumption (which covers all cases considered in [3]) that M is a Lie group G equipped with a smooth, transitive right-action on N; we denote the action by $(x, h) \mapsto x \cdot h$, where $x \in N$, $h \in G$. With these data, N is a (*right*) *homogeneous space* for G. We select a point $x_b \in N$ and define $\pi : G \to N$ by $\pi(h) = x_b \cdot h$. Recall that a Riemannian metric g_G on a Lie group G is called *left-invariant* if $L_h^* g_G = g_G$ for all $h \in G$, where $L_h : G \to G$ denotes left-multiplication by h.

Lemma 2.3. *In the setting above, if the metric g_G is left-invariant, then the vector field $\tilde{V} := -\mathrm{grad}_G E_G$ is projectable (where $E_G = \pi^* E_N = E_N \circ \pi$). The vector field V on N to which $\tilde{V}$ projects satisfies*

$$V_x = -\pi_{h*}\pi_{h*}^\dagger((\mathrm{grad}_N E_N)|_x) \quad \forall x \in N, h \in \pi^{-1}(x). \tag{2.2}$$

Proof: In general, if (M, g) is a Riemannian manifold and $\Phi : M \to M$ is an isometry, one has

$(\Phi_*)^{-1}(\mathrm{grad}_M f) = \mathrm{grad}_M(\Phi^* f)$. Hence in our situation, for all $h \in G$ we have $(L_h^{-1})_* \mathrm{grad}_G E_G = \mathrm{grad}_G(L_h^* E_G)$.

Suppose now that $\pi(h_1) = \pi(h_2)$. Then $h := h_2 h_1^{-1} \in \mathrm{Stab}(x_b)$ the stabilizer of x_b. But for such h we have $\pi \circ L_h = \pi$, implying that $L_h^* E_G = L_h^* \pi^* E_N = (\pi \circ L_h)^* E_N = \pi^* E_N = E_G$, and hence $(L_h^{-1})_* \mathrm{grad}_G E_G = \mathrm{grad}_G E_G$. In other words $\tilde{V}_{h_2} = \tilde{V}_{hh_1} = (L_h)_* \tilde{V}_{h_1}$. But then $\pi_{*h_2} \tilde{V}_{h_2} = (\pi \circ L_h)_* \tilde{V}_{h_1} = \pi_{*h_1} \tilde{V}_{h_1}$. Thus $\tilde{V}$ is projectable. Equation (2.2) now follows from (2.1). □

Thus the operator $\mathbf{L}$ we will use in (1.4) is given at $x \in N$ by

$$\mathbf{L}_x = \pi_{h*}\pi_{h*}^\dagger, \quad \forall h \in \pi^{-1}(x)\,. \tag{2.3}$$

As noted earlier, the vector field $\mathbf{L}(\mathrm{grad}_N E_N)$ is simply the gradient of E_N with respect to the metric g'_N in (1.5). This metric is just the natural "Riemannian submersion" metric on N induced by the left-invariant metric g_G on G, expressed in a way that is computable from an arbitrary metric g_N on N.

To relate these ideas to what is done in [3] and [4], assume now that we are initially given an energy function E_{amb}, defined on some Euclidean space $\mathbf{R}^m$ that we wish to minimize over certain submanifolds $N \subset \mathbf{R}^m$. Let g_{Euc} denote the standard Riemannian metric on $\mathbf{R}^m$ and let $\iota : N \to \mathbf{R}^m$. denote the inclusion map. Then N inherits the Riemannian metric $g_N = \iota^* g_{\mathrm{Euc}}$ and the restricted energy functional $E_N = \iota^* E_{\mathrm{amb}}$. For each $x \in \mathbf{R}^m$, there is a natural isomorphism $T_x\mathbf{R}^m \cong \mathbf{R}^m$ that allows us to identify vector fields on subsets of $\mathbf{R}^m$ with vector-valued functions on these subsets. For each $y \in \mathbf{R}^m$, we let j_y^{amb} denote the canonical isomorphism $T_y\mathbf{R}^m \to \mathbf{R}^m$; similarly for $x \in N$, we let $j_x^N : T_x N \to \mathbf{R}^m$ denote the corresponding identification of $T_x N$ with a subspace of $\mathbf{R}^m$. We can then write $(g_N)_x = (j_x^N)^* g_{\mathrm{std}}$, where g_{std} is a fixed inner product on $\mathbf{R}^m$, and we have

$$j_x^N = j_x^{\mathrm{amb}} \circ \iota_{*x} \tag{2.4}$$

for all $x \in N$. The vector field $\mathbf{L}(\mathrm{grad}_N E_N)$ in (1.4) is then identified with the vector-valued function $j^N \circ \mathbf{L}(\mathrm{grad}_N E_N)$. Furthermore, by the same argument as in Lemma 2.1 one has

$$\begin{aligned}(\mathrm{grad}_N E_N)_x &= (\mathrm{grad}_N \iota^* E_{\mathrm{amb}})_x \\ &= (\iota_{*x})^\dagger((\mathrm{grad}_{\mathrm{Euc}} E_{\mathrm{amb}})|_x), \quad \forall\, x \in N.\end{aligned} \tag{2.5}$$

Note that each map j_X^{amb} is an isometry, and hence its adjoint equals its inverse; thus from (2.4) we have $(i_{*x})^\dagger \circ (j_x^{\mathrm{amb}})^{-1} = (j_x^N)^\dagger$. Thus if we define the "ambient gradient" Z^{amb} (an $\mathbf{R}^m$-valued function on N) by

$$Z^{\mathrm{amb}}(x) = j_x^{\mathrm{amb}}((\mathrm{grad}_{\mathrm{Euc}} E_{\mathrm{amb}})|_x)\,, \tag{2.6}$$

then the vector field $\mathbf{L}(\mathrm{grad}_N E_N)$ on N is identified with the vector-valued function on N given by

$$\begin{aligned} x \mapsto\ & j_x^N(\mathbf{L}(\mathrm{grad}_N E_N)|_x) \\ &= j_x^N \circ \pi_{*h} \circ (\pi_{*h})^\dagger \circ (i_{*x})^\dagger((\mathrm{grad}_{\mathrm{Edu}} E_{\mathrm{amb}})|_x) \\ &= j_x \circ \pi_{*h} \circ (\pi_{*h})^\dagger \circ \left(j_x^N\right)^\dagger (Z^{\mathrm{amb}}(x)) \\ &= \left(j_x^N \circ \pi_{*h}\right) \circ \left(j_x^N \circ \pi_{*h}\right)^\dagger (Z^{\mathrm{amb}}(x)), \quad h \in \pi^{-1}(x). \end{aligned} \tag{2.7}$$

3. Examples

We consider three examples that relate to Sections 2.2, 2.3, and 2.4 of [3]. In all of these the appropriate $\mathbf{R}^m$ is a vector subspace of the space $\mathrm{Mat}(n)$ of $n \times n$ real matrices (sometimes the whole space). The standard inner product on $\mathrm{Mat}(n)$ can be written in the form

$$g_{\mathrm{std}}(Y, Z)|_X = \mathrm{tr}(Y^T Z) \tag{3.1}$$

where the superscript T denotes transpose. The inner product we use on vector subspaces of $\mathrm{Mat}(n)$ is just the restriction of (3.1).

Example 1. We take $\mathbf{R}^m = \mathrm{Sym}(n) :=$ {symmetric $n \times n$ matrices}, $N = \mathrm{Sym}^+(n) :=$ {positive-definite symmetric $n \times n$ matrices}, and $G = GL(n, \mathbf{R})$, the group of invertible $n \times n$ matrices. The map $\pi : G \to N$ defined by

$$\pi(A) = A^T A \tag{3.2}$$

is a surjective submersion, For $A \in G$ and $X \in N$, let $j_A^G : T_A G \to \mathrm{Mat}(n)$ and $j_X^N = j_X^{\mathrm{amb}} : T_X N \to \mathrm{Sym}(n)$ denote the natural identifications of tangent spaces of the submanifolds $G \subset \mathrm{Mat}(n)$, $N \subset \mathrm{Sym}(n)$ with $\mathrm{Mat}(n)$, $\mathrm{Sym}(n)$ respectively. Since G and N are open subsets of the indicated vector spaces, the maps j_A^G and j_X^N are all isomorphisms. We will also make use of the isomorphisms $\ell_A : T_I G \to T_A G$ (where I is the identity) defined by

$$\ell_A(Y) = (L_A)_* Y = \frac{d}{dt}(Ae^{tY})|_{t=0}, \qquad A \in G, \quad Y \in T_I G. \tag{3.3}$$

(Here the matrix exponential is used.) Thus we have the simple identity $j^G_A(\ell_A(Y)) = AY$. Since $\ell_A : T_I G \to T_A G$ is an isomorphism for each $A \in G$, we can define a left-invariant Riemannian metric g_G on G by

$$g_G(\ell_A(Y), \ell_A(Z))|_A = g_G(Y, Z)|_I := g_{\rm std}(Y, Z) = {\rm tr}(Y^T Z). \quad (3.4)$$

Writing $X = \pi(A) = A^T A$, we have

$$j^N_X \circ \pi_{*A}(\ell_A(Y)) = j^N_X \left(\frac{d}{dt}(\pi(Ae^{tY}))|_{t=0}\right) = Y^T A^T A + A^T AY = Y^T X + XY. \quad (3.5)$$

To compute the adjoint map, let $W \in {\rm Sym}(n)$ and define $W' \in T_I G$ by $(j^N_X \circ \pi_{*A})^\dagger(W) = \ell_A(W')$. Then for all $Y \in T_I G$ we have

$$\begin{aligned} & g_G(\ell_A(Y), \ell_A(W')) \\ &= g_G\big(\ell_A(Y), \big(j^N_X \circ \pi_{*A}\big)^\dagger(W)\big) \\ &= g_{\rm std}\big(\big(j^N_X \circ \pi_{*A} \circ \ell_A\big)(Y), W\big) \\ &= g_{\rm std}(Y^T X + XY, W) \\ &= {\rm tr}((Y^T X + XY)W) \\ &= {\rm tr}(Y^T(XW + (WX)^T)) \\ &= {\rm tr}(Y^T(2XW)) \text{ since } X \text{ and } W \text{ are symmetric} \\ &= g_G(\ell_A(Y), \ell_A(2XW)), \end{aligned}$$

implying $W' = 2XW$. Hence

$$\big(j^N_A \circ \pi_{*A}\big)^\dagger(W) = \ell_A(2XW). \quad (3.6)$$

Thus, from (2.7) we have

$$\begin{aligned} & j^N_X(\mathbf{L}({\rm grad}_N E_N)|X) \\ &= \big(j^N_X \circ \pi_{*A}\big) \circ \big(j^N_X \circ \pi_{*A}\big)^\dagger(Z^{\rm amb}(X)) \\ &= (2XZ^{\rm amb}(X))^T X + X(2XZ^{\rm amb}(X)) \\ &= 2(Z^{\rm amb}(X)X^2 + X^2 Z^{\rm amb}(X)). \quad (3.7) \end{aligned}$$

Example 2. We take $\mathbf{R}^m = {\rm Mat}(n)$, $G = N = O(n)$ (the orthogonal group), π = identity map. In this case $j^G_X = j^N_X$; however, unlike in Example 1 these maps are distinct from the canonical isomorphism $j^{\rm amb}_X; T_X{\rm Mat}(n) \to {\rm Mat}(n)$. The left-invariant metric on $O(n)$ is defined just as in (3.4), but in this case $T_I G$ is the space $\mathfrak{so}(n)$ of antisymmetric $n \times n$ matrices, so that at each $X \in O(n)$ we have

$$g_G(\ell_X(Y), \ell_X(Z))|_X = -{\rm tr}(YZ). \quad (3.8)$$

Again $j^G_X(\ell_X(Y)) = XY$, and $j^G_X(T_X N)$—the tangent space to G at X, viewed as a subspace of ${\rm Mat}(n)$—is simply $\{XY | Y \in \mathfrak{so}(n)\}$. Since $X^T X = I$, we have $g_{\rm std}(XY, XZ) = {\rm tr}((XY)^T XZ) = -{\rm tr}(YZ)$, and hence the Euclidean metric on N coincides with the left-invariant metric g_G. Thus both π_{*X} and its adjoint are identity maps, and one can easily show that for all $W \in {\rm Mat}(n)$, $(j^N_X)^\dagger(W) = \ell_X \circ {\rm proj}(X^T W)$, where the orthogonal projection ${\rm proj} : {\rm Mat}(n) \to \mathfrak{so}(n)$ is the map carrying a matrix B to its antisymmetric part $(B - B^T)/2$. Thus,

$$j_X(\mathbf{L}_X(Z)) = j_X j^\dagger_X(Z^{\rm amb}(X)) = X{\rm proj}(X^T Z^{\rm amb}(X)) = \frac{1}{2}(Z^{\rm amb} - X(Z^{\rm amb}(X))^T X). \quad (3.9)$$

Example 3. Fix a matrix $Q \in {\rm Sym}(n)$. Take $\mathbf{R}^m = {\rm Sym}(n)$, $G = O(n)$, and N the orbit of Q under the map π defined by $\pi(A) = A^T QA$. We use the same metric g_G on $O(n)$ as in Example 2. Again setting $X = \pi(A)$, this time we find

$$\big(j^N_X \circ \pi_{*A}\big)(\ell_A(Y)) = [X, Y] := XY - YX, \quad Y \in \mathfrak{so}(n). \quad (3.10)$$

As in Example 1, let $W \in {\rm Sym}(n) = j^{\rm amb}_X(T_X{\rm Sym}(n))$ and define $W' \in T_I G$ by $(j^N_X \circ \pi_{*A})^\dagger(W) = \ell_A(W')$. For all $Y \in T_I G$ we then compute

$$\begin{aligned} -{\rm tr}(YW') &= g_G(\ell_A(Y), \ell_A(W')) \\ &= g_G\big(\ell_A(Y), \big(j^N_X \circ \pi_{*A}\big)^\dagger(W)\big) \\ &= g_{\rm std}\big(\big(j^N_X \circ \pi_{*A} \circ \ell_A\big)(Y), W\big) \\ &= g_{\rm std}([X, Y], W) \\ &= {\rm tr}([X, Y], W) \quad \text{since} [X, Y] \text{ is symmetric}, \\ &= -{\rm tr}(Y[X, W]) \ . \end{aligned}$$

Since X and W are symmetric, $[X, W]$ is antisymmetric. Thus $[X, W] \in \mathfrak{so}(n)$ and $g_G(Y, W)|_I = g_G(Y, [X, W])|_I \; \forall Y \in \mathfrak{so}(n)$, so $W' = [X, W]$,

$$\big(j^N_X \circ \pi_{*A}\big)^\dagger(W) = \ell_A = ([X, W]), \quad (3.11)$$

and

$$j^N_X(\mathbf{L}({\rm grad}_N E_N)|_X) = [X, [Z^{\rm amb}(X)]] \ . \quad (3.12)$$

4. Generalization to Mapping-Spaces

When we replace N by $\mathcal{N} = \text{Maps}(\Omega, N)$, where Ω is a bounded open domain in (say) $\mathbf{R}^2$ or $\mathbf{R}^3$, certain geometrical aspects remain essentially unchanged, while analytical aspects can become more troublesome. To endow $\mathcal{N}$ with the structure of a Banach manifold one generally completes the space of smooth maps in a suitable Sobolev norm; henceforth in this context "map" means an element of such a completion. Given a map $X : \Omega \to N$, the tangent space $T_X\mathcal{N}$ is naturally identified with $\{Y \in \text{Maps}(\Omega, TN) \mid Y(p) \in T_{X(p)}N \ \forall p \in \Omega\}$ (the space of sections of the pulled-back tangent bundle X^*TN). Given a Riemannian metric g on N, we define an inner product on $T_X\mathcal{N}$ by

$$g_{\mathcal{N}}(Y, Z)|_X = \int_\Omega g_N(Y(p), Z(p))|_{X(p)} dp. \quad (4.1)$$

This is the standard "L^2 metric" on $\mathcal{N}$, a weak Riemannian metric ("weak" because the tangent spaces are not complete in this inner product). Given a smooth function $\mathcal{E} : \mathcal{N} \to \mathbf{R}$, we write $V = \text{grad}_{\mathcal{N}}\mathcal{E}$ (the "L^2-gradient" of $\mathcal{E}$) if $\langle d\mathcal{E}, Y\rangle_x = g_{\mathcal{N}}(V, Y)$ for all $X \in \text{domain}(V)$ and all C^∞ compactly supported $Y \in T_X\mathcal{N}$ (where $\langle dE, Y\rangle)_X$ is the variation of $\mathcal{E}$ at X in the direction Y). If G is as in Section 2 then we obtain an L^2 metric on $\mathcal{G} := \text{Maps}(\Omega, G)$, and it is not hard to check that the constructions in the previous section carry over pointwise to this setting, giving us a map $\pi^\Omega : \mathcal{G} \to \mathcal{N}$, and linear maps $\pi^\Omega_{*h} : T_h\mathcal{G} \to T_{\pi\circ h}\mathcal{N}$, and their adjoints (where now $h \in \mathcal{G}$). In particular the operators $\mathbf{L}_X : T_X\mathcal{N} \to T_X\mathcal{N}$ are given pointwise in terms of the finite-dimensional operators constructed in Section 2:

$$(\mathbf{L}_X(Y))(p) = \mathbf{L}_{X(p)}(Y(p)) \in T_{X(p)}N \ . \quad (4.2)$$

Thus, in the analogs of Examples 1–3 with N replaced by $\mathcal{N}$, the vector fields $j^{\mathcal{N}}(\mathbf{L}(\text{grad}_{\mathcal{N}}\mathcal{E}_{\mathcal{N}}))$ are still given pointwise by Eq. (3.7), (3.9), and (3.12), with Z^{amb} the L^2-gradient of a functional $E_{\text{amb}} : \text{Maps}(\Omega, \text{Mat}(n)) \to \mathbf{R}$. In this way we obtain Eq. (9)–(10) and (14) of [3], as well as the un-numbered equation following (12) in that paper.[1]

5. Advantages of This Approach

Let $E_{\text{amb}} : \mathbf{R}^m \to \mathbf{R}$ be a function we wish to minimize on a submanifold N, or more generally a function $\mathcal{E}_{\text{amb}} : \text{Maps}(\Omega, \mathbf{R}^m) \to \mathbf{R}$ to minimize over a subset $\mathcal{N}$ that is a submanifold of $\text{Maps}(\Omega, \mathbf{R}^m)$, in some suitable sense. *Any* Riemannian metric on N determines a gradient-flow equation for minimizing $E|_N$ and $\mathcal{E}_{\text{amb}}|_{\mathcal{N}}$ in particular this is true of the restriction to N of the Euclidean metric. The question arises as to why one would want to consider a change of metric on N, bringing the gradient-flow equation into the form (1.4).

The motivations suggested in [3] are twofold. One is that in the examples considered, the submanifolds N carry natural metrics other than the restricted Euclidean metric. It is certainly reasonable to consider the gradient-flow equation for any natural metric.

The second motivation mentioned in [3] concerns a numerical scheme for approximating the flow Eq. (1.1), replacing it with an iterative scheme of the form

$$x(t_{n+1}) = \exp_{x(t)}(V(t_n)\Delta t_n) \quad (5.1)$$

where exp is the Riemannian exponential map on N, which for submersions $G \to N$ is expressible in terms of matrix exponentials. One of the advantages of this scheme over an ordinary Euler scheme $x(t_{n+1}) = x(t_n) + V(t_n)\Delta t_n$ is that there is no danger of stepping off the constraint manifold N. Thus, if it is not too costly to compute the matrix exponentials, there is time saved in each iteration by not having to "project" back to N.

A third potential advantage of this approach is that it may happen that if one is searching for a minimum of E by solving (1.3) numerically in a given example, the particular nature of grad E may lead to undesirable features. In such instances it is reasonable to try to circumvent the numerical problems by instead solving an equation of the form (1.4) for some well-chosen $\mathbf{L}$, since the equilibria of (1.3) and (1.4) are identical. It may happen that the space N has more than one natural metric, for example, and that the metric that leads to the best numerical behavior of the discretized negative-gradient flow is not the most obvious metric. The examples given in [3] and [4] illustrate that the gradient-flows that one obtains for certain functionals on $\text{Maps}(\Omega, \mathbf{R}^m)$ appear to be very nicely behaved. However, *a priori* there is no mathematical reason to expect faster convergence with one natural metric than with another.

6. Potential Pitfalls of This Approach

When faced with a problem for which the restricted Euclidean gradient-flow equation for E on N (or the Euclidean L^2 gradient-flow for $\mathcal{E}$ on $\mathcal{N}$) is ill-behaved, the intended advantages of the scheme above may be illusory, and perhaps even misleading. This is especially true when N is an open subset of Euclidean space, as in Example 1 of Section 3. The reasons are already present in the finite-dimensional situation $\Omega = \{\text{point}\}$, so we concentrate on that case, and remark on further complications in the infinite-dimensional case only at the end of this section.

First we recall some generalities concerning when and why the gradient-flow approach works. In (1.1), if the vector field V is C^1 (or, more generally, locally uniformly Lipschitz), then local existence and uniqueness of solutions to (1.1) are guaranteed. Long-time existence of solutions to (1.1) is a problem only when V is not sufficiently regular or when N is not compact. In (1.2), if $\mathbf{L}$ is C^1 (or at least locally uniformly Lipschitz), we again have the same existence/uniqueness behavior as in (1.1).

This negative-gradient flow preserves all sub-level sets $N^c := \{x \in N | E(x) \leq c\}, c \in \mathbf{R}$. If these sub-level sets are compact, the flow $x(t)$ will exist for all $t \geq 0$ and all initial conditions x_0. Compactness ensures the existence of at least one critical point of E, a global minimum, and if the critical points of E are isolated then every flow line will converge to a critical point as $t \to \infty$. If E has only one local minimum point $x_{\min}$ (necessarily the global minimum), then the flow will converge to $x_{\min}$, with probability 1 in the space of initial conditions N^c(for any c greater than or equal to the minimum value of E). Thus in this situation, one can reliably locate a global minimum by following the flow from a randomly chosen initial condition. Below, we will refer this as *topologically guaranteed convergence*. Note that if N is compact, then automatically so are the sub-level sets of any continuous function E.

Even in this finite-dimensional setting, several problems can arise with the flow when the sub-level sets N^c of E are not compact: first, there is no guarantee that a minimum of E exists; second, one cannot guarantee long-time existence of the gradient flow for all initial conditions, and third, even when one has long-time existence starting at a given x_0, the trajectory $x(t)$ may fail to converge by permanently exiting any fixed compact subset of N in finite time.

In the setting of Example 1, the manifold N is noncompact, and for functions E whose minimization is commonly considered, the sub-level sets are also noncompact. Thus a given function E on N may fail to have a minimum, and we do not have topologically-guaranteed convergence of the flow. A change of metric on N, or even a more general transformation of the flow equation as in (1.4), cannot remedy this situation, nor can following a flow on the "intermediate" manifold M that projects to the desired flow on N. It is the function E itself, not the metric on N (or factorizability through an intermediate manifold), that determines the obstacles to *whether* the flow converges; changing the metric can only affect the *rate* of convergence, not turn a divergent flow into a convergent one. If the flow (1.3) *does* converge with, say, a Euclidean metric, convergence may or may not be more rapid after a change of metric. Thus, while it is worth trying a new metric to see whether convergence-rate is improved, *a priori* there is no reason to expect more rapid convergence with the flow (1.4) (using 3.12) than for the numerically simpler Euclidean gradient flow.

Returning to the general setting, another issue concerns the approximation of the gradient flow by (5.1), with $V = -\text{grad}_N E$, in place of (for example) a naïve Euler scheme, in which one would have to compensate for stepping off the constraint manifold at each iteration. It should be noted that a scheme of the form (5.1), with *any* complete metric on N, achieves this constraint-preservation. In particular this is true for the restricted Euclidean metric (when it is complete, as in Examples 2 and 3 of Section 3), not just for other natural metrics on N such as the ones in Section 3. On a *closed* constraint manifold N that happens to be a homogeneous space for a matrix group, the chief advantage to using the metrics discussed in Section 3 (as compared with other replacements for the Euclidean metric) is that the exponential map is very easily computed.

However, when N is not closed, then the apparent advantage above may be illusory. We consider again the situation of Example 1, in which N is an open cone in a Euclidean space. The non-compactness of N may allow a given function E not to have a minimum, and this will be reflected in the non-convergence of the gradient flow, whether we use the restricted Euclidean gradient or the one in (3.7). The Euclidean metric is incomplete, so there are tangent vectors for which the exponential map is not defined, and one may not be able to use (5.1). However, this fact is a relevant feature of the problem, rather than a defect of the Euclidean gradient flow. Using the complete metric from Example 1 will allow us to define the scheme (5.1) perfectly well, but its properties may be misleading. For example, a non-convergent trajectory may slowly approach the boundary of N as $t \to \infty$, instead of reaching the boundary in finite time as it might were the Euclidean metric used. Given a trajectory that is well-approximated by (5.1), this reduction in speed may make it appear that the iterates are converging to a nonexistent critical point in N, rather than slowly diverging to the boundary of N. Thus, such an attempt to "pull the flow back" from crossing the boundary can succeed only at the expense of making it falsely appear that a divergent trajectory converges. A numerical scheme that approximates (1.4) by an iterative procedure generating a sequence $\{x_n = x(t_n)\}$ by some procedure (such as the one in (5.1)) that forces x_n to satisfy the constraint for all n, cannot both converge and yet faithfully represent the differential equation along non-convergent trajectories.

As an illustration of this phenomenon, consider a punctured unit disk D^* in $\mathbf{R}^2$ and let Z be the unit outward radial vector field. Starting at any point of D^*, the flow of Z reaches the boundary circle in finite time. If we replace Z at each point x by $\mathbf{L}_x(Z)$ for some positive-definite linear operator $\mathbf{L}_x$ defined at each point, possibly approaching zero at the boundary, we may slow down the flow, and will change its trajectories ($\mathbf{L}$ may rotate Z), but we cannot make the flow converge in D^*. However, by slowing down the flow enough, we may make it *appear* to converge.

Furthermore, note that even given a complete metric on a noncompact manifold N, not every gradient vector-field will have integral curves defined for all time. Thus it is possible to have a function E for which the gradient flow, with respect to the metric g_N, reaches the boundary in finite time (more precisely, for which the gradient flow, starting from some point, does not exist beyond some finite time). This again is a *significant* feature of the function E. However, because the metric is complete, the scheme (5.1) will be defined for *all* n and all step sizes Δt_n. Thus in this instance (5.1) cannot be a good approximation to the gradient flow.

It should also be noted that when N is an open subset of the ambient Euclidean space, one loses one of the chief motivations for replacing a naïve Euler step by another scheme; since N is flat there is no danger that a small Euler step will take one out of the constraint manifold. Mathematically, constraining a function to take its values in a closed submanifold (necessarily of positive codimension) is very different from constraining it to take its values in an open subset (necessarily of codimension zero) of Euclidean space.

When the finite-dimensional submanifold $N \subset \mathbf{R}^m$ is replaced by the infinite-dimensional manifold $\mathcal{N} = \text{Maps}(\Omega, N)$ where $\dim(\Omega) > 0$, the space $\mathcal{N}$ is never even locally compact (even if N is compact), so all the problems related to non-compactness in the finite-dimensional situation can arise. Additionally, in such situations, $\text{grad}_{\mathcal{N}}\mathcal{E}_{\mathcal{N}}$ may be defined only on a dense subset of $\mathcal{N}$, and may not have the regularity properties necessary to ensure even short-time existence of the gradient flow.

Although manifolds $\text{Maps}(\Omega, N)$ as above are never locally compact, nonetheless one may expect the gradient flow to have better convergence properties if N is compact than otherwise. This may underlie the empirical findings in [3, Section 3] and [4, Section 5.2] for diffusion-tensor MRI regularization. Starting with the same initial data, flows based on the three choices of N in Examples 1–3 of 3 of this paper (Sections 2.2–2.4 of [3]) were followed. The results were better for cases 2 and 3, in which N is compact (in the case of choice 3, potentially a different compact manifold for each point in Ω), than in case 1, in which N is noncompact.

Appendix: A Glossary of Some Differential-Geometric Terms

Precise and self-contained definitions of many of the terms below would require many pages. Instead, in these instances we have given basic ideas instead of actual definitions, at the cost of some precision. For precise treatments, among the more accessible accounts for non-experts are [2, 15], and [1, chapter 1], which consider only finite-dimensional spaces; [13, chapter 9], written in sufficient generality to handle finite and certain infinite-dimensional spaces; and [14, chapter 1], which focuses on infinite-dimensional spaces.

Underlined terms are defined elsewhere in the glossary.

The **cotangent space** T_p^*N at a point p of a manifold N is the dual of the tangent space: $T_p^*N = (T_pN)^*$. Its elements are called covectors.

The **derivative** Φ_{*p} at p of a smooth map Φ from a manifold M to a manifold N is the linear map $T_pM \to T_{\Phi(p)N}$ given by generalized directional derivatives: $\Phi_{*p}(v) = \frac{d}{dt}(\gamma)|_{t=0}$ where γ is any curve in M with $\gamma(0) = p$ and $\gamma'(0) = v$. When $N = \mathbf{R}^m$ we compose with the natural isomorphism $j_{\Phi(p)} : T_{\Phi(p)}\mathbf{R}^m \to \mathbf{R}^m$ and define the **differential** $d\Phi|_p = j_{\Phi(p)} \circ \Phi_{*p} : T_pM \to \mathbf{R}^m$.

A **diffeomorphism** between open sets in a (topological) vector space or between manifolds is a smooth bijective map whose inverse is smooth.

The **dual space** V^* of a real (topological) vector space V is the space of (continuous) linear functions $V \to \mathbf{R}$. If V is finite-dimensional, "continuous" is redundant. If $\alpha \in V^*$ and $v \in V$, we define $\langle \alpha, v \rangle := \alpha(v)$. The function $\langle \cdot, \cdot \rangle : V^* \times V \to \mathbf{R}$ is called the **dual pairing**.

A **fiber bundle** M **over** N is a family of manifolds (called fibers), all diffeomorphic to some fixed manifold F, smoothly parametrized by another manifold N, and together filling out a manifold M. More formally, a fiber bundle is a triple (M, N, π) where the *total space* M and *base space* N are manifolds, $\pi : M \to N$ is a smooth surjective map (the *projection map* of the bundle), such that for small enough open sets $U \subset N$, $\pi^{-1}(U)$ is diffeomorphic to $U \times F$ via a map carrying each fiber $\pi^{-1}(p)$ to $\{p\} \times F$. (The last condition, called *local triviality*, is guaranteed if the derivative of π is surjective at each point and either (i) F is compact or (ii) M is a Lie group for which N is a homogeneous space.) Example: If $M = G$ is a Lie group and H is a closed subgroup, then the space of left-cosets G/H is the base-space of a fiber bundle with total space G and projection-map $\pi : q \mapsto qH$. A **vector bundle** is a fiber bundle whose fibers are vector spaces; examples are the **tangent bundle** $TN := \bigcup_{p \in N} T_pN$ and *cotangent bundle* $T^*N := \bigcup_{p \in N} T_p^*N$ of a manifold N. A **principal fiber bundle** is a fiber bundle in which the fibers are homogeneous spaces for a Lie group H acting on itself by right-translations (in particular each fiber is diffeomorphic to H).

The **gradient** grad ψ, when it exists, of a smooth real-valued function ψ on a Riemannian manifold (N, g), is the (automatically unique) vector field for which $g_p(\text{grad } \psi|_p, v) = \langle d\psi_p, v\rangle \ \forall p \in N, v \in T_pN$. The gradient of ψ always exists if N is finite-dimensional; if N is infinite-dimensional, grad ψ may or may not exist.

A **homogeneous space** for a Lie group G is a manifold N on which G acts smoothly and transitively (any point of N can be mapped to any other point by some element of G). In such a situation, if p_0 in N and H is the stabilizer of p_0, then there is a diffeomorphism $N \cong G/H$ (the left- or right-coset space accordingly as G acts from the left or right on N) commuting with the G-action, and thus exhibiting G as a principal fiber bundle over N with fiber H.

An **isometry** from a Riemannian manifold (M, g_M) to a Riemannian manifold (N, g_N) is a smooth map Φ whose derivative preserves inner products (equivalently, for which $\Phi^* g_N = g_M$).

A **Lie group** is a manifold equipped with a group-structure for which all the group-operations are smooth. All of the classical matrix-groups (e.g. the orthogonal groups $O(n)$) are Lie groups.

"lower-star" notation e.g. "F_*": see pullbacks and push-forwards.

A **manifold** is a space N that is locally, but not necessarily globally, topologically equivalent to an open set in some fixed topological vector space V, e.g. a finite- or infinite-dimensional Banach space (in the latter case we call N a **Banach manifold**). The local pieces are required to fit together smoothly (via local diffeomorphisms). This makes the local topology and notion of differentiable functions on N the same as those of V. Note that a (topological) vector space is a special case of a manifold.

A **1-form** η on a manifold N is an assignment of a covector $\eta_p \in T_p^*N$ to each p in N. In the differential-geometry literature, definitions usually require η to be smooth, or at least continuous. Example: the differential of a real-valued function on N.

pullbacks and **push-forwards** are objects induced on one manifold from an object of the same type on another manifold via a map Φ from one manifold to the other. (Note: this terminology is *completely unrelated* to the phrase "pull the flow back" in Section 6.) For concreteness, let $\Phi : M \to N$ be the map. Differential geometers use upper-stars to indicate pullbacks—objects $\Phi^*\xi$ on M constructed from Φ and an object ξ on N—and lower-stars to indicate push-forwards of objects defined on M to objects defined on N. Example 1: If $p \in M$ and $v \in T_pM$, the derivative Φ_{*p} pushes-forward vectors in T_pM to vectors in $T_{\Phi(p)}N$. Example 2: If ψ is a real-valued function on N, its pullpack to M is defined by $\Phi^*\psi := \psi \circ \Phi : M \to \mathbf{R}$. Example 3: If ψ is a 1-form on N, its pullback to M is the 1-form defined by $\langle (\Phi^*\psi)|_p, v\rangle = \langle \psi_{\Phi(p)}, \Phi_{*p}v\rangle$ for all $p \in M, v \in T_pM$. Example 4: If g is a Riemannian metric on N, its pullback to M is the Riemannian metric defined by $(\Phi^*g)_p(v, w) = g_{\Phi(p)}(\Phi_{*p}v, \Phi_{*p}w)$ for all $p \in M, v, w \in T_pM$. Example 5: If E is a vector bundle over N, the pulled-back bundle Φ^*E is defined by declaring the fiber over $p \in M$ to be the fiber of E over $\Phi(p) \in N$.

A **Riemannian manifold** is a pair (N, g) where N is a manifold and g is a (weak) Riemannian metric on N. An important special case is Euclidean space with the standard inner product on vectors based at the same point.

A (weak) **Riemannian metric** g on a manifold N is an assignment g_p of an inner product on each tangent space T_pN, smoothly varying with p. "Weak" is a technical term that applies only if N is infinite-dimensional.

A **section** s of a fiber bundle (M, N, π) is an assignment of an element of $\pi^{-1}(p) \subset M$ to each $p \in N$. In the literature, definitions usually require s to be smooth, or at least continuous. Example: a vector field (respectively, 1-form) on N is a section of the tangent bundle TN (resp., cotangent bundle T^*N). All vector bundles have infinitely many continuous sections, but many fiber bundles have *no* continuous sections (there are topological obstructions).

A **smooth** map is one that is C^k (i.e., has continuous derivatives up through order k) for some author-dependent $k \geq 1$ (usually $k = 1$ or ∞). In this paper, $k = 1$ suffices; often in the differential-geometry literature, "smooth" means "C^∞".

The **tangent space** T_pN at a point of a manifold N is the vector space of infinitesimal variations of p within N; its elements are called *tangent vectors*. When N is a submanifold of Euclidean space, T_pN can be viewed as the dim(N)-dimensional hyperplane tangent to N at p, with p treated as the origin.

A **topological vector space** V is a vector space equipped with a topology for which the vector-space operations are continuous. A norm on V determines a topology, and if V is finite-dimensional, all norms yield the same topology (the *norm topology*); thus in the finite-dimensional case, "vector space" is taken to mean "topological vector space with the norm topology". For infinite-dimensional vector spaces, different norms can yield different topologies, and some important topologies do not come from *any* norm.

"**upper-star**" **notation**, e.g. "F^*": see pullbacks and push-forwards.

vector bundle: see fiber bundle.

A **vector field** X on a manifold N is an assignment of a tangent vector $X_p \in T_pN$ to each p in N. In the differential-geometry literature, definitions usually require X to be smooth, or at least continuous.

Note

1. The cited equations in [3] simplify to ours, modulo an unimportant overall factor-of-2 discrepancy between the definition of the metric used in [3, Sections 2.3 and 2.4] and our definition, because the matrix-valued functions called G (our Z^{amb}) in [3, Sections 2.2 and 2.4] are symmetric on the domain of the flow.

References

1. R. Abraham and J.E. Marsden, *Foundations of Mechanics*, 2nd ed. Benjamin/Cummings Publ. Co.: Reading, 1978.
2. W.M. Boothby, *An Introduction to Differentiable Manifolds and Riemannian Geometry*, 2nd ed. Academic Press: San Diego, 1986.
3. C. Chefd'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras, "Constrained flows of matrix-valued functions: application to diffusion tensor regularization," *Proceedings of International ECCV 2002 Workshop (Copenhagen)*, Springer-Verlag: Berlin, 2002, pp. 251–265.
4. C. Chefd'hotel, D. Tschumperlé, R. Deriche, and O. Faugeras, "Regularizing flows for constrained matrix-valued images," *J. Math. Imaging and Vision*, Vol. 20, pp. 147–162, 2004.
5. P.J. Basser and C. Pierpaoli, "Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI," *J. Magnetic Resonance* Ser. B, Vol. 111, pp. 209–219, 1996.
6. R.I. Shrager and P.J. Basser, "Anisotropically weighted MRI," *Magn. Reson. Med.*, Vol. 40, pp. 160–165, 1998.
7. C.R. Tench, P.S. Morgan, M. Wilson, and L.D. Blumhardt, "White matter mapping using diffusion tensor MRI," *Magn. Reson. Med.*, Vol. 47, pp. 967–972, 2002.
8. C.-F. Westin, S.E. Maier, H. Mamata, A. Nabavi, F.A. Jolesz, and R. Kikinis, "Processing and visualization for diffusion tensor MRI," *Med. Image Anal.*, Vol. 6, pp. 93–108, 2002.
9. D. Le Bihan, "Methods and applications of diffusion MRI," *Magnetic Resonance Imaging and Spectroscopy in Medicine and Biology*, I.R. Young (Ed.), John Wiley and Sons, 2000.
10. D. Tschumperlé and R. Deriche, "Diffusion tensor regularization with constraints preservation," in *Proceedings of the 1st IEEE Workshop on Variational and Level Set Methods (VLSM'01)*, July 2001.
11. 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 (VLSM'01)*, July 2001.
12. S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Vol. 1, John Wiley & Sons, 1963.
13. L. Loomis and S. Sternberg, *Advanced Calculus*, Addison-Wesley, Reading, Massachusetts, 1968.
14. J. Marsden, *Applications of Global Analysis in Mathematical Physics*, Publish or Perish, Inc.: Boston, 1974.
15. M. Spivak, *A Comprehensive Introduction to Differential Geometry*, Publish or Perish, Inc.: Boston, 1979.
16. N. Steenrod, *The Topology of Fiber Bundles*, Princeton Mathematical Series Vol. 14, Princeton University Press: Princeton, 1951.

David Groisser is an Associate Professor of Mathematics at the University of Florida. Dr. Groisser conducts research in pure and applied differential geometry. His current interests include the geometry of shape spaces and applications of differential geometry to image-analysis.