



Geodesics, Parallel Transport & One-parameter Subgroups for Diffeomorphic Image Registration

Marco Lorenzi, Xavier Pennec

► To cite this version:

Marco Lorenzi, Xavier Pennec. Geodesics, Parallel Transport & One-parameter Subgroups for Diffeomorphic Image Registration. International Journal of Computer Vision, 2013, 105 (2), pp.111-127. 10.1007/s11263-012-0598-4 . hal-00813835

HAL Id: hal-00813835

<https://inria.hal.science/hal-00813835v1>

Submitted on 2 May 2013

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

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

Geodesics, Parallel Transport & One-parameter Subgroups for Diffeomorphic Image Registration

Marco Lorenzi · Xavier Pennec

the date of receipt and acceptance should be inserted later

Abstract Computational anatomy aims at developing models to understand the anatomical variability of organs and tissues. A widely used and validated instrument for comparing the anatomy in medical images is non-linear diffeomorphic registration which is based on a rich mathematical background. For instance, the “large deformation diffeomorphic metric mapping” (LDDMM) framework defines a Riemannian setting by providing a right invariant metric on the tangent spaces, and solves the registration problem by computing geodesics parametrized by time-varying velocity fields. A simpler alternative based on Stationary Velocity Fields (SVF) has been proposed, using the one-parameter subgroups from Lie groups theory. In spite of its better computational efficiency, the geometrical setting of the SVF is more vague, especially regarding the relationship between one-parameter subgroups and geodesics.

In this work, we detail the properties of finite dimensional Lie groups that highlight the geometric foundations of one-parameter subgroups. We show that one can define a proper underlying geometric structure (an affine manifold) based on the canonical Cartan connections, for which one-parameter subgroups and their translations are geodesics. This geometric structure is perfectly compatible with all the group operations (left, right composition and inversion), contrarily to left- (or right-) invariant Riemannian metrics. Moreover, we derive closed-form expressions for the parallel transport.

Then, we investigate the generalization of such properties to infinite dimensional Lie groups. We suggest that some of the theoretical objections might actually be ruled out by the practical implementation of both the LDDMM and the SVF frameworks for image registration. This leads us to a more practical study comparing the parameterization (initial velocity field) of metric and Cartan geodesics in the specific optimization context of longitudinal and inter-subject image registration. Our experimental results suggest that stationarity is a good approximation for longitudinal deformations, while metric geodesics notably differ from stationary ones for inter-subject registration, which involves much larger and non-physical deformations. Then, we turn to the practical comparison of five parallel transport techniques along one-parameter subgroups. Our results point out the fundamental role played by the numerical implementation, which may hide the theoretical differences between the different schemes. Interestingly, even if the parallel transport generally depend on the path used, an experiment comparing the Cartan parallel transport along the one-parameter subgroup and the LDDMM (metric) geodesics from inter-subject registration suggests that our parallel transport methods are not so sensitive to the path.

1 Introduction

One of the main objectives of computational anatomy is to develop suitable statistical models on several subjects for the study of the anatomical variability of organs and tissues. In particular, longitudinal observations from time series of images are an important source of information for understanding the developmental processes and the dynamics of pathologies. Thus, a reliable

Marco Lorenzi
Asclepios Research Group, INRIA Sophia Antipolis, France
LENITEM, IRCCS San Giovanni di Dio, Fatebenefratelli,
Brescia, Italy
E-mail: marco.lorenzi@inria.fr

Xavier Pennec
Asclepios Research Group, INRIA Sophia Antipolis, France
E-mail: xavier.pennec@inria.fr

method for comparing different longitudinal trajectories is required, in order to develop population-based longitudinal models.

Non-rigid registration is an instrument for the detection of anatomical changes on medical images, and it has been widely applied on different clinical contexts for the definition of population-based anatomical atlases [38, 26, 9]. However, in case of longitudinal data, the optimal method for comparing *deformation trajectories* across different subjects is still unknown. In fact, the methods for integrating the subtle inter-subject changes into the group-wise analysis have an important impact on the accuracy and reliability of the subsequent statistical results. The aim is to preserve as much as possible the biological informations carried on by the different subjects, while allowing a precise comparison in a common geometrical space.

Among the different techniques proposed for the comparison of longitudinal trajectories ([33, 5, 10]), the *parallel transport* represents a promising method which relies on a solid mathematical background. Basically, it consists in transporting the infinitesimal deformation vector across different points by preserving its properties with respect to the space geometry, such as the parallelism.

Parallel transport has been introduced in medical imaging within the LDDMM setting [40, 42]. LDDMM solves the image registration problem by using a Riemannian framework in which the deformations are parametrized as *diffeomorphisms* living in a suitable space, once provided an opportune right-invariant metric [27]. The registration problem is solved by computing the deformation that best matches the images with a penalization on the energy of the trajectory in the space of diffeomorphisms. The solution is given by the endpoint of a geodesic parametrized by a time-varying velocity field. This endpoint can also be parametrized by the Riemannian exponential map of the initial velocity field (or its metric dual, the momentum). The setting allows the computation of the parallel transport along geodesics at the cost of a computationally intensive scheme, and this limitation often prevents the application on high resolution images or large datasets. Moreover, this assumes that both longitudinal deformations and inter-subject transformations live in the same space of diffeomorphisms, which may have very different characteristics.

A simplified solution to the diffeomorphic registration problem was introduced with the stationary velocity field (SVF) setting [2]. In this case, the diffeomorphisms are one-parameter subgroups parameterized by *stationary* velocity fields through the Lie group exponential. This restriction allows an efficient numerical

scheme for the computation of the deformation but it does apparently not rely on any geometric assumption on the underlying space. This implies that some important mathematical properties are not guaranteed, for instance whether the one-parameter subgroups are still geodesics or if the space is complete. In spite of this lack of knowledge, the framework was found very efficient and reliable in many application in different contexts [25, 23, 36]. For instance, a framework based on the Schild's Ladder has been proposed for the evaluation of the parallel transport with the SVF in [22].

In this paper we investigate the relationship between Lie groups and affine geometry and we highlight many interesting properties that provide the SVF setting with part of the geometrical solidity required. In Section 2 we present the relevant properties of the finite dimensional Lie groups and the relationship with the Riemannian setting for the definition of the geodesics and the parallel transport. Section 3 is dedicated to a discussion on the extension of the Lie group theory in the infinite dimensional case. In Section 4, we study the differences between the registration based on the one-parameter subgroups and on the Riemannian metric on specific registration problems. Finally, Section 5 focuses on the evaluation of Cartan's parallel transport of deformation vectors in the image registration context.

2 Differential Geometry on Lie Group

This section references the conceptual basis for the definition of the parallel transport along the one-parameter subgroups. Many details on differential geometry on Lie groups can be found in classical books like [11, 7, 13]. However, most results of this section are more easily found in the more modern (and quite comprehensive) presentation of differential geometry and Lie groups of Postnikov [32].

2.1 Basics of Lie Groups

A Lie group \mathbb{G} is a smooth manifold provided with an identity element id , a smooth associative composition rule $(g, h) \in \mathbb{G} \times \mathbb{G} \mapsto gh \in \mathbb{G}$ and a smooth inversion rule $g \mapsto g^{-1}$ which are both compatible with the differential manifold structure. As such, we have a tangent space $T_g\mathbb{G}$ at each point $g \in \mathbb{G}$. A vector field \mathbf{X} is a smooth function that maps a tangent vector $\mathbf{X}|_g$ to each point g of the manifold. The set of vector fields (the tangent bundle) is denoted $T\mathbb{G}$. Vector fields can be viewed as the directional (or Lie) derivative of a scalar function ϕ along the vector field at each point: $\partial_{\mathbf{X}}\phi|_g = \frac{\partial\phi(g+t\mathbf{X}|_g)}{\partial t}$. Composing directional derivatives

$\partial_{\mathbf{X}}\partial_{\mathbf{Y}}\phi$ leads in general to a second order derivation. However, we can remove the second order terms by subtracting $\partial_{\mathbf{Y}}\partial_{\mathbf{X}}\phi$ (this can be checked by writing these expression in a local coordinate system). We obtain the Lie bracket that acts as an internal multiplication in the algebra of vector fields:

$$[\mathbf{X}, \mathbf{Y}](\phi) = \partial_{\mathbf{X}}\partial_{\mathbf{Y}}\phi - \partial_{\mathbf{Y}}\partial_{\mathbf{X}}\phi.$$

Given a group element $a \in \mathbb{G}$, we call left translation L_a the composition with the fixed element a on the left: $L_a : g \in \mathbb{G} \mapsto ag \in \mathbb{G}$. The differential DL_a of the left translation maps the tangent space $T_g\mathbb{G}$ to the tangent space $T_{ag}\mathbb{G}$. We say that a vector field $\mathbf{X} \in T(\mathbb{G})$ is left invariant if it remains unchanged under the action of the left translation: $DL_a\mathbf{X}_g = \mathbf{X}|_{ag}$. The sub-algebra of left-invariant vector fields is closed under the Lie bracket and is called the Lie algebra \mathfrak{g} of the Lie group. Since a left-invariant vector field is uniquely determined by its value at identity through the one-to-one map $\tilde{\mathbf{X}}|_g = DL_g\mathbf{X}$, the Lie algebra can be identified to the tangent space at the identity $T_{id}\mathbb{G}$. One should notice that any smooth vector field can be written as a linear combination of left-invariant vector fields with smooth *functional* coefficients.

Left-invariant vector fields are *complete* in the sense that their flow φ_t is defined for all time. Moreover, this flow is such that $\varphi_t(g) = g\varphi_t(id)$ by left invariance. The map $X \mapsto \varphi_1(id)$ of \mathfrak{g} into \mathbb{G} is called *Lie group exponential* and denoted by \exp . In particular, the group exponential defines the *one-parameter subgroup* associated to the vector X and has the following properties:

- $\varphi_t(id) = \exp(tX)$, for each $t \in \mathbb{R}$;
- $\exp((t+s)X) = \exp(tX)\exp(sX)$, for each $t, s \in \mathbb{R}$.

In finite dimension, it can be shown that the Lie group exponential is a diffeomorphism from a neighborhood of 0 in \mathfrak{g} to a neighborhood of id in \mathbb{G} .

For each tangent vector $X \in \mathfrak{g}$, the one parameter subgroup $\exp(tX)$ is a curve that starts from identity with this tangent vector. One could question if this curve could be seen as a geodesic like in Riemannian manifolds. To answer this question, we first need to define what are geodesics. In a Euclidean space, straight lines are curves which have the same tangent vector at all times. In a manifold, tangent vectors at different times belong to different tangent spaces. When one wants to compare tangent vectors at different points, one needs to define a specific mapping between their tangent spaces: this is the notion of parallel transport. There is generally no way to define globally a linear operator $\Pi_g^h : T_g\mathbb{G} \rightarrow T_h\mathbb{G}$ which is consistent with composition (i.e. $\Pi_g^h \circ \Pi_f^g = \Pi_f^h$). However, specifying the parallel transport for infinitesimal displacements allows

integrating along a path, thus resulting into a parallel transport that depend on the path. This specification of the parallel transport for infinitesimal displacements is called the (affine) connection.

2.2 Affine Connection Spaces

An *affine connection* on \mathbb{G} is an operator which assigns to each $\mathbf{X} \in T(\mathbb{G})$ a linear mapping $\nabla_{\mathbf{X}} : T(\mathbb{G}) \rightarrow T(\mathbb{G})$ such that, for each vector field $\mathbf{X}, \mathbf{Y} \in T(\mathbb{G})$, and smooth function $f, g \in C^\infty(\mathbb{G}, \mathbb{R})$

$$\nabla_{f\mathbf{X}+g\mathbf{Y}} = f\nabla_{\mathbf{X}} + g\nabla_{\mathbf{Y}} \quad (\text{Linearity}); \quad (1)$$

$$\nabla_{\mathbf{X}}(f\mathbf{Y}) = f\nabla_{\mathbf{X}}(\mathbf{Y}) + (\mathbf{X}f)\mathbf{Y} \quad (\text{Leibniz rule}). \quad (2)$$

The *affine connection* is therefore a derivation on the tangent space which infinitesimally maps tangent vectors from one tangent plane to another. Once the manifold provided with a connection, it is possible to generalize to the manifolds the notion of “straight lines”: a vector field \mathbf{X} is *parallel transported* along a curve $\gamma(t)$ if $\nabla_{\dot{\gamma}(t)}\mathbf{X} = 0$ for each t . Thus, a path $\gamma(t)$ on \mathbb{G} is said to be straight or *geodesic* if $\nabla_{\dot{\gamma}}\dot{\gamma} = 0$.

In a local coordinate system, the geodesic equation is a second order differential equation. Thus, given a point $p \in \mathbb{G}$ and a vector $X \in T_p\mathbb{G}$, there exist a unique geodesic $\gamma(t, p, X)$ such that at the instant $t = 0$ passes through p with velocity X [32]. We define therefore the *Affine exponential* as the application $\exp : \mathbb{G} \times T(\mathbb{G}) \rightarrow \mathbb{G}$ given by $\exp_p(X) = \gamma(1, p, X)$.

If, as in the Euclidean case, we want to associate to the straight lines the property of minimizing the distance between points, we need to provide the group \mathbb{G} with a Riemannian manifold structure, i.e. with a metric operator g on the tangent space. In this case there is a unique connection, called *Levi-Civita connection*, which, for each $\mathbf{X}, \mathbf{Y}, \mathbf{Z} \in T(\mathbb{G})$:

- Preserves the metric, i.e. the parallel transport along any curve connecting f to g is an isometry:

$$g(\mathbf{X}, \mathbf{Y})_g = g(\Pi_g^f\mathbf{X}, \Pi_g^f\mathbf{Y})_f.$$

- Is torsion free:

$$\nabla_{\mathbf{X}}\mathbf{Y} - \nabla_{\mathbf{Y}}\mathbf{X} = [\mathbf{X}, \mathbf{Y}],$$

thus the parallel transport is symmetric with respect to the Lie bracket.

By choosing the Levi-Civita connection of a given Riemannian metric, the affine geodesics are the length minimizing paths (i.e. classical Riemannian geodesics). However, given a general affine connection, there may not exist any Riemannian metric for which affine geodesics are length minimizing.

2.3 From Affine Geodesic to One-Parameter Subgroups

Given an affine connection ∇ and a vector X on $T_{id}\mathbb{G}$, we can therefore define two curves on \mathbb{G} passing through id and having X as tangent vector, one given by the Lie group exponential \exp and the other given by the affine exponential \exp_{id} . When do they coincide?

The connection ∇ on \mathbb{G} is *left-invariant* if, for each left translation L_a ($a \in \mathbb{G}$) and any vector fields \mathbf{X} and \mathbf{Y} , we have $\nabla_{DL_a\mathbf{X}}(DL_a\mathbf{Y}) = DL_a\nabla_{\mathbf{X}}(\mathbf{Y})$. Using two left invariant vector fields $\tilde{\mathbf{X}}, \tilde{\mathbf{Y}} \in \mathfrak{g}$ generated by the tangent vectors $X, Y \in T_{id}\mathbb{G}$, we see that $\nabla_{\tilde{\mathbf{X}}}\tilde{\mathbf{Y}}$ is itself a left-invariant vector field generated by its value at identity. Since a connection is completely determined by its action on the left-invariant vector fields (we can recover the connection on arbitrary vector fields using Eq. (1,2) from their decomposition on the Lie Algebra), we conclude that each left-invariant connection ∇ is uniquely determined by a product α (symmetric bilinear operator) on $T_{id}\mathbb{G}$ through

$$\alpha(X, Y) = \nabla_{\tilde{\mathbf{X}}}\tilde{\mathbf{Y}}\Big|_{id}.$$

Notice that such a product can be uniquely decomposed into a commutative part $\alpha' = \frac{1}{2}(\alpha(X, Y) + \alpha(Y, X))$ and a skew symmetric part $\alpha'' = \frac{1}{2}(\alpha(X, Y) - \alpha(Y, X))$. The symmetric part specifies the geodesics (i.e. the parallel transport of a vector along its own direction) while the skew-symmetric part specifies the torsion which governs the parallel transport of a vector along a different direction (the rotation around the direction of the curve if we have a metric connection with torsion).

Following [32], a left-invariant connection ∇ on a Lie group \mathbb{G} is a *Cartan connection* if, for any tangent vector X at the identity, the one-parameter subgroups and the affine geodesics coincide, i.e. $\exp(tX) = \exp_{id}(t, id, X)$. We can see that a Cartan connection satisfies $\alpha(X, X) = 0$ or, equivalently, is purely skew-symmetric.

The one-dimensional family of connections generated by $\alpha(X, Y) = \lambda[X, Y]$ obviously satisfy this skew-symmetry condition. Moreover, the connections of this family are also invariant by right translation [31], thus invariant by inversion also since they are already left invariant. This makes them particularly interesting since they are fully compatible with all the group operations.

Among this family, three connections have special curvature or symmetric properties and are called the canonical Cartan-Schouten connections [8]. The zero curvature connections given by $\lambda = 0, 1$ (with torsion $T = -[\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}]$ and $T = [\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}]$ respectively on left invariant vector fields) are called *left* and *right* Cartan connections. The choice of $\lambda = 1/2$ leads to average the

left and right Cartan connections. It is called the *symmetric* (or *mean*) Cartan connection. It is torsion-free, but has curvature $R(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}})\tilde{\mathbf{Z}} = -\frac{1}{4}[[\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}], \tilde{\mathbf{Z}}]$.

As a summary, the three canonical Cartan connections of a Lie group are (for two left-invariant vector fields):

$$\begin{aligned} \nabla_{\tilde{\mathbf{X}}}\tilde{\mathbf{Y}} &= 0 && \text{Left (Torsion, Flat);} \\ \nabla_{\tilde{\mathbf{X}}}\tilde{\mathbf{Y}} &= \frac{1}{2}[\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}] && \text{Symmetric (Torsion-Free, Curved);} \\ \nabla_{\tilde{\mathbf{X}}}\tilde{\mathbf{Y}} &= [\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}] && \text{Right (Torsion, Flat).} \end{aligned}$$

Since the three canonical Cartan connections only differ by torsion, they share the same affine geodesics which are the left and right translations of one parameter subgroups. In the following, we call them *group geodesics*. However, the parallel transport of general vectors along these group geodesics is specific to each connection as we will see below.

2.4 Left and Right Invariant Riemannian Metrics on Lie Groups

Given a metric $\langle X, Y \rangle$ on the tangent space at identity of a group, one can propagate this metric to all tangent spaces using left (resp. right) translation to obtain a left- (resp. right-) invariant Riemannian metric on the group. In the left-invariant case we have $\langle DL_a X, DL_a Y \rangle_a = \langle X, Y \rangle$ and one can show [20] that the Levi-Civita connection is the left-invariant connection generated by the product

$$\alpha(X, Y) = \frac{1}{2}[X, Y] - \frac{1}{2}(ad^*(X, Y) + ad^*(Y, X)),$$

if the operator ad^* verifying $\langle ad^*(Y, X), Z \rangle = \langle [X, Z], Y \rangle$ for all $X, Y, Z \in \mathfrak{g}$ is well defined. A similar formula can be established for right-invariant metrics using the algebra of right-invariant vector fields.

We clearly see that this left-invariant Levi-Civita connection has a symmetric part which makes it differ from the Cartan symmetric connection $\alpha(X, Y) = \frac{1}{2}[X, Y]$. In fact, the quantity $ad^*(X, X)$ specifies the rate at which a left invariant geodesic and a one parameter subgroup starting from the identity with the same tangent vector X deviates from each-other. More generally, the condition $ad^*(X, X) = 0$ for all $X \in \mathfrak{g}$ turns out to be a necessary and sufficient condition to have a bi-invariant metric [32]. It is important to notice that geodesics of the left- and right-invariant metrics differ in general as there do not exist bi-invariant metrics even for simple groups like the Euclidean motions [31]. However, right invariant geodesics can be obtained

from the left invariant one through inversion: if $\phi(t)$ is a left invariant geodesic joining identity to the transformation ϕ_1 , then $\phi^{-1}(t)$ is a right-invariant geodesic joining identity to ϕ_1^{-1} .

2.5 Parallel Transport on Cartan Connections

For the left Cartan connection, the unique fields that are covariantly constant are the *left-invariant* vector fields, and the parallel transport is induced by the differential of the left translation [32], i.e. $\Pi^L : T_p\mathbb{G} \rightarrow T_q\mathbb{G}$ is defined as

$$\Pi^L(X) = DL_{qp^{-1}}X. \quad (3)$$

One can see that the parallel transport is actually independent of the path, which is due to the fact that the curvature is null: we are in a space with absolute parallelism. Similarly, the *right-invariant* vector fields are covariantly constant with respect to the right invariant connection only. As above, the parallel transport is given by the differential of the right translation

$$\Pi^R(X) = DR_{p^{-1}q}X, \quad (4)$$

and we have an absolute parallelism as well.

Finally, the parallel transport for the symmetric Cartan connection is given by the infinitesimal alternation of the left and right transports. However, as there is curvature, it depends on the path: it can be shown [13] that the parallel transport of X along the geodesic $\exp(tY)$ is:

$$\Pi^S(X) = DL_{\exp(\frac{1}{2}Y)}DR_{\exp(\frac{1}{2}Y)}X. \quad (5)$$

2.6 The Schild's Ladder: a Parallel Transport Scheme for Symmetric Connections

A more general method for parallel transport was introduced in [29] after Schild's similar constructions [21]. The Schild's Ladder infinitesimally transports a vector along a given curve through the construction of geodesic parallelograms (Figure 1). Since this scheme relies only on geodesics, it is only valid for symmetric connections with no torsion [19], in which case it provides a first order approximation of the parallel transport. Although the Schild's Ladder was known in gravitation theory for 40 years, it was apparently not used in practice as a computational tool before it was turned into an algorithm for the parallel transport of deformation vectors in [22].




Fig. 1 The Schild's Ladder. Given a curve C , the vector \mathbf{A} on P_0 is transported to P_1 in two steps: 1) compute the geodesic connecting P_2 and P_1 and define the mid-point P_3 , 2) compute the geodesic from P_0 to P_3 and prolongate twice to reach P_4 . The tangent vector \mathbf{A}' to the curve connecting P_1 and P_4 is the parallel transport of \mathbf{A} .

3 A Glimpse of Lie Group Theory in Infinite Dimension

In the previous Section, we derived the equivalence of one-parameter subgroups and the affine geodesics of the canonical Cartan connections in a finite dimensional Lie group. In order to use such a framework for diffeomorphisms, we have to generalize the theory to infinite dimensions. However, defining infinite dimensional Lie groups is raising much more difficulties. This is in fact the reason why Lie himself restricted to finite dimensions. The theory was developed since the 70ies and is now an active field of research. We refer the reader to the recent books [18,41] for more details on this theory and to [34] for a good overview of the problems and applications.

3.1 Infinite Dimensional Lie Groups

The basic construction scheme is to consider an infinite dimensional manifold endowed with smooth group operations. Such a Lie group is locally diffeomorphic to an infinite-dimensional vector space which can be a Fréchet space (a locally convex space which is complete with respect to a translation invariant distance), a Banach space (where the distance comes from a norm) or a Hilbert space (where the norm is derived from a scalar product). We talk about Fréchet, Banach or Hilbert Lie groups, respectively. Extending differential calculus from R^n to Banach and Hilbert spaces is straightforward, but this is not so simple for Fréchet spaces. In particular, the dual of a Fréchet space need not be Fréchet, which means that some extra care must be taken when

defining differential forms. Moreover, some important theorems such as the inverse function theorem hold for Banach spaces but not necessarily for Fréchet spaces.

For instance, the set $Diff^k(\mathcal{M})$ of C^k diffeomorphisms of a compact manifold \mathcal{M} is a Banach manifold and the set of Sobolev H^s diffeomorphisms $Diff^s(\mathcal{M})$ is a Hilbert manifold (if $s > \dim \mathcal{M}/2$). However, these are no-classical "Lie groups" since one loses derivatives when differentiating the composition and inversion maps. To obtain the complete smoothness of the composition and inversion maps, one has to go to infinity, but the Banach structure is lost in the process [34, p.12] and we are left with $Diff^\infty(\mathcal{M})$ being only a Fréchet Lie group. Some additional structure can be obtained by considering the sequence of $Diff^k(\mathcal{M})$ spaces as a succession of dense inclusions as k goes to infinity: this is the Inverse Limit of Banach (ILB)-Lie group setting. Likewise, the succession of dense inclusions of Sobolev H^s diffeomorphisms give rise to the Inverse Limit of Hilbert (ILH)-Lie group setting.

3.2 General Groups of Diffeomorphisms

As the diffeomorphisms groups considered are Fréchet but not Banach, the usual setting of infinite dimensional Lie groups is the general framework of Fréchet manifolds. This implies that many of the important properties which are true in finite dimension do not hold any more for general infinite dimensional Lie groups [35].

First, there is no implicit or inverse function theorem (except Nash-Moser type theorems.) This implies for instance that the log-map (the inverse of the exponential map) may not be smooth even if the differential of the exponential map is the identity.

Second, the exponential map is not in general a diffeomorphism from a neighborhood of zero in the Lie algebra onto a neighborhood of the identity in the group. This means that it cannot be used as a local chart to work on the manifold. For instance in $Diff^s(\mathcal{M})$, in every neighborhood of the identity there exists diffeomorphisms which are not the exponential of an H^s vector field. A classical example of the non-surjectivity of the exponential map is the following function in $\text{Diff}(\mathbb{S}^1)$ [28]:

$$f_{n,\epsilon}(\theta) = \theta + \pi/n + \epsilon \sin^2(n\theta). \quad (6)$$

This function can be chosen as close as we want to the identity by opportunely dimensioning ϵ and θ . However, it can be shown that it cannot be reached by any one-parameter subgroup, and therefore the Lie group exponential is not a local diffeomorphism of $\text{Diff}(\mathbb{S}^1)$.

This example is quite instructive and shows that this theoretical problem might actually be a very prac-

tical advantage: the norm of the k -th derivative of $f_{n,\epsilon}$ is exploding when k is going to infinity, which shows that we would rather want to exclude this type of diffeomorphisms from the group under consideration.

3.3 Theoretical Background of Diffeomorphic Image Registration

In the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework [41], a different construction is leading to a more restricted subgroup of diffeomorphisms which is more rational from the computational point of view. One first chooses a Hilbert norm on the Lie Algebra which turn it into an admissible Hilbert space. Admissible means that it can be embedded into the space of vector fields which are bounded and vanishing at infinity, as well as all the first order derivatives. Typically, this is a Sobolev norm of a sufficiently high order. Then, one restricts to the subgroup of diffeomorphisms generated by the flow of integrable sequences of such vector fields for a finite time. To provide this group with a Riemannian structure, a right invariant metric is chosen. The reason for choosing right translation is that it is simply a composition which does not involve a differential operator as for the left translation. One can show that the group provided with this right-invariant metric is a complete metric space: the choice of the norm on the Lie algebra is specifying the subgroup of diffeomorphisms which are reachable, i.e. which are at a finite distance.

3.4 The Stationary Velocity Fields (SVF) framework

In [2], Arsigny proposed to parameterize deformations by the exponential (i.e. the flow) of stationary velocity fields [2]. The fact that the flow in an autonomous ODE allows us to generalize efficient algorithms such as the scaling and squaring algorithm: given an initial approximation $\exp(\delta Y) = id + \delta Y$, the exponential of a SVF Y can be efficiently and simply computed by recursive compositions:

$$\exp(Y) = \exp\left(\frac{Y}{2}\right) \circ \exp\left(\frac{Y}{2}\right) = \left(\exp\left(\frac{Y}{2^n}\right)\right)^{2^n}.$$

A second algorithm is at the heart of the efficiency of the optimization algorithms with SVFs: the Baker-Campbell-Hausdorff (BCH) formula [4] tells us how to approximate the log of the composition:

$$\begin{aligned} BCH(X, \delta Y) &= \log(\exp(X) \circ \exp(\delta Y)) \\ &= X + \delta Y + \frac{1}{2}[X, \delta Y] + \frac{1}{12}[X, [X, \delta Y]] + \dots \end{aligned}$$

In order to have a well-posed space of deformations, we need to specify on which space is modeled the Lie algebra, as previously. This is the role of the regularization term of the SVF registration algorithms [39, 15] or of the spline parameterization of the SVF in [3, 30]: this restricts the Lie algebra to the sub-algebra of sufficiently regular velocity fields. The subgroup of diffeomorphisms considered is then generated by the flow of these stationary velocity fields and their finite composition. So far, the theoretical framework is very similar to the LDDMM setting and we can see that the diffeomorphisms generated by the one-parameter subgroups (the exponential of SVFs) all belong to the group considered in the LDDMM setting, provided that we model the Lie algebra on the same admissible Hilbert space. As in finite dimension, the affine geodesics of the Cartan connections (group geodesics) are metric-free (the Hilbert metric is only used to specify the space on which is modeled the Lie Algebra) and generally differ from the Riemannian geodesics of LDDMM.

When modeling the Lie algebra on a reproducing kernel Hilbert space (RKHS) with a real analytic kernel (typically a Gaussian Kernel), the Lie algebra is stable under the bracket, which is necessary for the classical Lie group theory. If we could show that the adjoint operator is uniformly bounded with a sufficiently small bound, then we would expect the BCH series to be convergent. Having such a BCH-Lie group would already be sufficient to justify the optimization steps performed in the SVF registration framework.

However, it is well known that the subgroup of diffeomorphisms generated by this Lie algebra is significantly larger than what is covered by the group exponential. Indeed, although our affine connection space is geodesically complete (all geodesics can be continued for all time without hitting a boundary), there is no Hopf-Rinow theorem which state that any two points can be joined by a geodesic (metric completeness). Thus, in general, not all the elements of the group \mathbb{G} may be reached by the one-parameter subgroups. An example in finite dimension is given by $SL(2)$.

However, this might not necessarily results into a problem in the image registration context since we are not interested in recovering “all” the possible diffeomorphisms, but only those which lead to admissible anatomical transformations. For instance, the diffeomorphism on the circle defined above at Eq. (6) cannot be reached by any one-parameter subgroup of \mathbb{S}^1 . However, since

$$\lim_{k \rightarrow \infty} \|f_{n,\epsilon}\|_{H^k} \rightarrow \infty,$$

this function is not well behaved from the regularity point of view, which is a critical feature when dealing with image registration.

In practice, we have a spatial discretization of the SVF (and of the deformations) on a grid, and the temporal discretization of the time varying velocity fields by a fixed number of time steps. This intrinsically limits the frequency of the deformation below a kind of “Nyquist” threshold, which prevents these diffeomorphisms to be reached anyway both by the SVF and by the “discrete” LDDMM frameworks. Therefore, it seems more importance to understand the impact of using stationary velocity fields in registration from the *practical* point of view, than from the theoretical point of view, because we will have *in fine* to deal with the unavoidable numerical implementation and relative approximation issues.

4 Practical Differences between Metric and Group Geodesics in Registration

In this Section, we investigate the practical differences between metric and group geodesics, by applying the SVF and LDDMM image registration settings on common data. In this case, we are interested in the differences between the estimated diffeomorphisms and the related tangent parametrization.

We investigate two different scenarios: *longitudinal* and *inter-subject* registration problems. In the former case the registration aims at retrieving the small differences occurring on the same subject after two subsequent imaging session, most likely due to precise biological processes. In the latter case, the differences are given by the large variability among different subjects, with potential structural and topological changes. In this case, even if there might not exist real one-to-one correspondences, the diffeomorphic constraint is important in term of smoothness and reliability of the represented deformation. In the following the SVF and LDDMM settings will be provided by the Log-Demons and the AtlasWerk registration algorithms.

4.1 Log-Demons

The symmetric log-domain diffeomorphic Demons (or Log-Demons) [39] estimates a diffeomorphic deformation parametrized by a SVF. The Log-Demons alternates the estimation of unconstrained correspondences (encoded through the exponential of a SVF v_c) that optimize the image similarity measure, and the estimation of the transformation parameters (a SVF v) that best explains the correspondence field using a penalized least-squares approach. The regularity criterion is an infinite order differential quadratic forms (Q_k) of the

velocity field [6],

$$\begin{aligned} \text{Reg}(v) &= \|v_c - v\|^2 + \sum_{k=1}^{\infty} \frac{Q_k(v)}{\sigma_t^{2k}} \\ &\simeq \|v_c - v\|^2 + \frac{\sigma^2}{2} \|\nabla v\|^2 + \frac{\sigma^4}{4} \|\nabla^2 v\|^2 + \dots, \end{aligned}$$

where we chose $\sigma_t^2 = 2/\sigma^2$ so that the optimum is explicitly obtained through a Gaussian convolution of smoothing parameter σ [25]. The Lie group exponential is efficiently implemented by the scaling and squaring computational scheme. An additional (so-called fluid) regularization step is often performed when updating the correspondences in addition to the above elastic-like penalization.

In the following experiments we used the regularization parameters $\sigma_{fluid} = 0.5$ and $\sigma_{el} = 1.5$, maximum update step length $\sigma_x = 2$ voxels, with a multi-resolution scheme of 100 and 50 iterations at coarser and finer level respectively. This choice leads to a reasonable compromise between image matching and smoothness of the deformation in both longitudinal and inter-subject settings, as already tested in several previous experiments. The average computational time for the registration of a couple of brain images (image resolution 182x218x182, voxel size 1x1x1) is 25 minutes on a AMD Opteron dual core 2000Mhz.

4.2 AtlasWerks

The AtlasWerks suite [1] is an implementation of the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework. The algorithm minimizes the sum of an image similarity metric and an elastic regularization term which measures the energy of the trajectory deforming one image into the other. This formulation ensures that the optimum trajectory is a geodesic evolution equation, which means that the path is completely defined by the initial tangent vector or, equivalently, by the initial momentum [17,16]. The energy of the geodesic is measured by the norm of the initial vector $\|v_0\|^2 = \|\mathbf{L}v_0\|_{L^2}^2$ (or equivalently by its integral along the geodesic path). The differential operator $\mathbf{L} = \gamma - \alpha \nabla^2 - \beta \nabla(\nabla \cdot)$ is controlled by three parameters α, β and γ responsible respectively for the smoothness and the compressibility of the deformation, and trade-off between matching and regularity terms. This energy corresponds to a Sobolev H^2 norm on the Lie Algebra, which is weaker than the one imposed by the Log-Demons regularization: roughly, the normalized weight α/γ of the first order term could be compared to the weight $\sigma^2/2$ of the Log-Demons while β/γ should be compared to zero. The same comparison is also valid for




Fig. 2 Intra-subject registration. Tangent representation for the intra-subject longitudinal deformation estimated by A) the SVF, and B) the LDDMM (initial tangent vector $v(0)$) settings. C,C') LDDMM time varying velocity field $v(t_i)$ at time points $t_i = 0.4$ and 0.8 , and D) associated magnitude measured by the L^2 norm.

the second order term. Higher order terms are not penalized by AtlasWerks while they are increasingly penalized with the Log-Demons.

The AtlasWerks suite requires to set-up different registration parameters and, in order to obtain results compatible with those given by the Demons registration, we finally chose the ones which lead to the maximum SSD similarity with the Demons results (not shown). Therefore, the following experiments were performed by setting the fluid registration parameters as follow: $\alpha = 1.5, \beta = 0.01, \gamma = 3$ and $\sigma = 15$, with a multi-resolution scheme of 100 and 50 iterations and the geodesic path computed on 5 time points. These parameters seems of the same order than the Log-Demons one ($\alpha/\gamma = 0.5$ versus $\sigma^2/2 = 1$, $\beta/\gamma = 0.003$ versus 0). In the following experiments on brain images the average computational time was of 150 minutes.




Fig. 3 Deformation trajectory for the inter-subject warp estimated by the LDDMM (top sequence), and the SVF (bottom sequence).

4.3 Longitudinal Registration

We chose the one year follow-up brain images from 5 subjects affected by Alzheimer’s disease from a clinical study presented elsewhere [12]. The MR images were acquired with a 1.0 Tesla Philips Gyroscan. The T1-weighted scan was acquired in the sagittal plane with a gradient echo techniques as follows: TR= 20 ms, TE= 5 ms, flip angle= 30°, field of view= 220 mm, acquisition matrix= 256x256, slice thickness= 1.3 mm. After a standard preprocessing pipeline consisting in affine registration and histogram matching, the images were diffeomorphically registered with the LDDMM and the SVF algorithms.

In Figure 2, top row, we can see an example of the different tangent representations for the estimated deformations. The two methods retrieve similar deformation patterns, both mapping to the cortical areas, hippocampus and white matter.

In this example related to small longitudinal changes, the time-varying representation does not exhibit relevant changes in time concerning the direction of the deformation vectors (Figure 2, row C,C’), and the magnitude (Table 2), here evaluated as the L^2 norm of the velocity fields. In fact, as shown in Figure 2, row D, and in Table 2, the variation of the time-varying velocity fields are minimal in terms of location and magnitude of the forces.

In Table 1 we can see that both Demons and LDDMM provide similar results for the matching in terms of the SSD. Interestingly, the similarity between the deformed images obtained with the two registration methods is higher than the similarity of the deformed images with the target. This example suggests that, in the small deformation setting, both stationary and time-

varying parameterizations lead to similar results, and that the LDDMM velocity field stays barely constant during the evolution.

4.4 Inter-subject Registration

In this experiment, the two methods were employed to register a set of 10 T1 brain scans from the LPBA40 dataset of the Laboratory of Neuro Imaging (LONI) at UCLA¹[37]. A reference subject was chosen among the dataset, and the other brain images underwent affine registration and histogram matching prior to non-rigid registration with the SVF and the LDDMM frameworks.

In Figure 4 we can observe an example of the estimated tangent representation of the deformations. Contrarily to the intra-subject case, the variation in orientation and magnitude of the time-varying velocity fields is here more pronounced (rows C-C’, D), and the same result is quantified by the increased non-stationarity of the LDDMM velocities (Table 2). Moreover, the velocity field appears more localized than for the SVF. This could come from the regularization of the higher differential terms in the Log-Demons for the SVF. However, we also know that the optimal momentum for LDDMM has to be aligned with the gradient of the image and is thus localized on the edges. This second reason is probably more important although it remains to be quantified.

Table 3 (first row) shows the average L^2 distance between the displacement fields obtained with the two different frameworks. As already described before, the

¹ <http://www.loni.ucla.edu/Atlases/LPBA40>

SVF and LDDMM displacements are closer in the longitudinal setting than in inter-subject. Figure 3 shows the estimated evolution from the source to the target image under the two registration settings. As can be seen for instance for the deformation on the ventricles, the two methods perform here differently. However, as in the experiments on the longitudinal registration, the two methods provide similar results in terms of resulting SSD (Table 1).

Clearly these experiments do not aim to provide comparisons in terms of accuracy and precision which notably depend on the opportune tuning of the registration parameters. However it is interesting to notice the very similar result for the intensity matching in spite of the rather different parameterizations estimated by the two methods.




Fig. 4 Inter-subject registration. Tangent representation for the inter-subject deformation estimated by A) the SVF, and B) the LDDMM (initial tangent vector $v(0)$) settings. C, C') LDDMM time varying velocity field $v(t_i)$ at time points $t_i = 0.4$ and 0.8 , and D) associated magnitude measured by the L^2 norm. We notice the variation in magnitude and location of the time varying velocity.

Table 1 Average SSD (standard deviation) for the different tests in the longitudinal and the inter-subject registration settings. The SVF method performs very similarly to the LDDMM, as also shown by the low SSD between the resulting warped images (Row 2). Moreover, the exponential given by the “scaling and squaring” numerical scheme of the initial LDDMM tangent vector generally provides lower matchings with the target images, though the resulting warped images are still consistent with the LDDMM ones, especially in the longitudinal setting (Rows 4-5).

	Longitudinal	Inter-Subject
SVF-Target	256(37)	355(39)
SVF-LDDMM	139(102)	122(70)
LDDMM-Target	284(90)	372(35)
ScSq-LDDMM	162(105)	368(62)
ScSq-Target	514(111)	950(98)

Table 2 Evaluation of the non-stationarity of the LDDMM velocity fields. Average (standard deviation) change in the L^2 norm $\frac{\int_{\Omega} \|v(0) - v(t)\|_{L^2}^2}{\int_{\Omega} \|v(0)\|_{L^2}^2}$ for the sampled LDDMM time varying velocity fields in the inter-subject and longitudinal settings. As qualitatively shown in Figure 2 and 4, the LDDMM evolution is more stationary in the longitudinal setting.

Relative distance from $v(0)$		
	Longitudinal setting	Inter-Subject setting
$v(0.2)$	0.026 (0.020)	0.007 (0.013)
$v(0.4)$	0.077 (0.088)	0.046 (0.027)
$v(0.6)$	0.13 (0.125)	0.283 (0.213)
$v(0.8)$	0.208 (0.160)	1.6 (1.23)

Table 3 Differences between LDDM and SVF displacements. Relative average L^2 distance (standard deviation) $\frac{\int_{\Omega} \|\varphi_{LDDMM} - \varphi_{SVF}\|_{L^2}^2}{\int_{\Omega} \|\varphi_{LDDMM}\|_{L^2}^2}$ between the displacement fields obtained in the different settings: inter-subject vs longitudinal registration, and Lie group vs Riemannian exponentials on the LDDMM initial tangent vector. As for the SSD on the resampled images (Table 1), the SVF method performs more similarly to the LDDMM in the longitudinal setting. Moreover, in this setting the “scaling and squaring” Lie group exponential is closer to the Riemannian one.

	Longitudinal	Inter-Subject
SVF-LDDMM	674 (4.66)	5692 (1322)
ScSq-LDDMM	1.93 (1.63)	39.3 (58.8)

4.5 Lie Group vs Riemannian Exponential of the Same Initial Vector Field

Finally, we investigated the practical differences between the Lie group exponential and the Riemannian one, by comparing the relative results in terms of vector field exponentiation. For this purpose, the Lie group exponential implemented by the scaling and squaring algorithm was applied on the initial vector field $v(0)$ provided by the LDDMM, and the resulting warped image




Fig. 5 Lie group exponential (scaling and squaring) applied to the LDDMM initial vector for the intra and the inter-subject registration. From left to right: Target image, A) resulting warped source image with the scaling and squaring, B) warped source image with the LDDMM, and differences between C1) LDDMM and Scaling and Squaring warped images, C2) LDDMM warped and target images and C3) scaling and squaring warped and target images. We note that although the scaling and squaring provides worse matching with respect to the target, the resulting warped image is still close to the LDDMM one.

was compared to the LDDMM one, which is obtained from the geodesic evolution equation.

Figure 5 shows the result for both intra and inter-subject registration. In the longitudinal case (top row), the Lie group exponential generates results close to the LDDMM ones, as shown by the low intensity differences between the respective warped images. This is confirmed at the group level (Table 1), where the average SSD between the two exponential methods is lower than the one resulting from the LDDMM registration (162 vs 284), even though the Lie group exponential led to increased average intensity mismatch with respect to the target (514). In the inter-subject case the differences are more pronounced (Figure 5, bottom row): the Lie group exponential of the initial tangent vector produces a deformation that differs from the original LDDMM one (for instance in the ventricles and around the cortex). This is reflected at the group level by the resulting similarities (Table 1), where the average SSD between the exponential methods is comparable to the one of the inter-subject registration LDDMM-Target (368 vs 372), and the mismatch Lie Group Exponential-Target is sensitively higher (950). The result is confirmed when considering the relative difference between the L^2 norm of the displacements given by two exponentials (Table 3).

This experiment is a supporting argument for the stationary nature of the longitudinal deformations, where the scaling and squaring of the tangent vector lead to a satisfactory description of the morphological changes. Thus, for small deformations, the metric geodesics seems

to correspond to the one-parameter subgroups. A rather different scenario is given by the inter-subject registration problem, where the evolution of the momentum for the geodesic evolution cannot be neglected and produces a different representation of the deformation from the one from the one-parameter subgroup.

5 Cartan Parallel Transport along Group Geodesics in Practice

In group of diffeomorphisms, the left and right translations are respectively $L_g \exp(X) = g \circ \exp(X)$, and $R_g \exp(X) = \exp(X) \circ g$. Their first order Taylor expansion leads to

$$DL_g(X) = Dg \cdot X \quad DR_g(X) = X \circ g,$$

where $Dg(x) = \partial g(x)/\partial x$ is the usual Jacobian matrix. We can therefore provide an *explicit closed form formula* for the parallel transport with respect to the canonical Cartan connections. In particular, if X is a vector field to be transported along the one-parameter subgroup $\exp(tY)$ we have:

$$\Pi_Y^L(X) = D\exp(Y) \cdot X, \quad (7)$$

$$\Pi_Y^R(X) = X \circ \exp(Y), \quad (8)$$

$$\Pi_Y^S(X) = D\exp\left(\frac{Y}{2}\right) \cdot \left(X \circ \exp\left(\frac{Y}{2}\right)\right). \quad (9)$$

5.1 Computing the Jacobian of the Deformation

From the computational point of view, we notice that among the three transport methods, Π^R requires the simple resampling of the velocity field by the transformation, while both Π^L and Π^S involve the computation of the Jacobian matrix of the exponential. The presence of these first order differential terms is raising numerical accuracy problems when the Jacobian matrix is computed by finite differences of the displacement field sampled on the image grid. In the case of large deformations, the displacement field is undergoing high frequency changes and its sampling in the finite difference scheme is notably error prone.

We can alleviate this numerical problem by taking advantage of the properties of the one-parameter subgroups. Rather than computing directly the Jacobian $D\exp(Y)$ using finite differences on the final displacement field, we can derive a recursive scheme from the following property:

$$\exp\left(\frac{nY}{N}\right) = \exp\left(\frac{(n-1)Y}{N}\right) \circ \exp\left(\frac{Y}{N}\right)$$

Starting from the approximation $D\exp(Y/N) \simeq \text{Id} + \frac{1}{N}DY$ for a suitable scaling factor N , we have the recursive formula (for $n = 2$ to N):

$$D\exp\left(\frac{nY}{N}\right) = D\exp\left(\frac{(n-1)Y}{N}\right) \Big|_{\exp(\frac{Y}{N})} \cdot D\exp\left(\frac{Y}{N}\right) \quad (10)$$

Thanks to this recursive scheme, the Jacobian is computed by finite differences on a low frequency displacement field, which sampling does not raise numerical issues, and the matrix valued image of Jacobian matrices is recursively resampled and multiplied along the one-parameter subgroup. Thus, the scheme avoids the sampling of a high frequency field, at the cost of multiple interpolations. Here, interpolating derivatives (Jacobian matrices) instead of displacements is giving a very important gain in numerical accuracy. Although the smoothness of the interpolation scheme (trilinear, spline, etc.) may be thought of as having an important impact on the final computational accuracy, we use simple trilinear interpolation in the sequel.

5.2 Schild's Ladder Implementation for SVF

The Schild's Ladder scheme was introduced in the SVF setting in [22] to parallel transport longitudinal trajectories along inter-subject deformations. When applied

to the symmetric Cartan connection, we can take advantage of the symmetry properties of the parallelogram to replace the computation of the geodesics with the composition of group exponentials, and the initial tangent vector of the resulting geodesic can be efficiently approximated with the BCH formula. This leads to a computationally efficient and numerically stable method where the parallel transport of a vector X along the trajectory $\exp(tY)$ is

$$\Pi_{Schild}^Y(X) = X + t[Y, X] + \frac{t^2}{2}[Y, [Y, X]] + O(t^3),$$

with $[Y, X] = DY \cdot X - DX \cdot Y$. Notice that here we take the spatial gradient of the SVFs Y and X while the previous transport schemes differentiate the exponential $\exp(Y)$. This might make an important numerical difference.

5.3 Synthetic Experiment on a Simplified Geometry

Morphometric studies often investigate pathological phenomena described by the loss of matter, which is modeled by compressible deformations and quantified by scalar indices of volume change, such as the Jacobian determinant or its logarithm. Therefore it is important to preserve these measures when normalizing to a reference space in group-wise studies of deformation trajectories. In this section we propose a simple example aimed at testing the ability of the proposed Cartan parallel transport techniques to preserve the atrophy trajectory simulated on a simplified geometry.

A synthetic progression of longitudinal atrophy was simulated on a 3D gray matter sphere S_0 enclosing




Fig. 6 Synthetic example: Intra and inter-subject variations from the sphere source space to the ellipsoid target space with related deformations.

Table 4 Average measures of changes on the gray matter layer. Top-row (Source Space): changes measured on the reference sphere at each time point 1 – 4. Bottom-rows: changes measured from the transported longitudinal deformations on the ellipsoid. For the conjugate action it was not possible to compute the L^2 Norm of the associated stationary velocity field, since it acts on deformation fields.

	L^2 Norm				Log Jacobian			
	1	2	3	4	1	2	3	4
Source Space	2.97	9.85	22.68	44.62	-4.77	-9.54	-14.76	-19.14
Π^L	3.02	9.57	22.14	42.32	-5	-9.82	-14.88	-20.43
Π^R	2.94	10	22.81	44.58	-4.70	-9.36	-14.51	-19.18
Π^S	3.3	11.17	25.7	50.37	-5.74	-11.2	-17.13	-23.65
Schild’s Ladder	3.65	10.74	24.3	51.49	-4.83	-9.86	-14.65	-19.11
Conjugate	/	/	/	/	-2.6	-5.5	-9.18	-13.93
	Jacobian				Elastic energy			
	1	2	3	4	1	2	3	4
Source Space	0.68	0.47	0.35	0.37	3.47	3.93	4.5	5.23
Π^L	0.69	0.51	0.43	0.45	3.51	4.01	4.67	5.53
Π^R	0.69	0.49	0.36	0.37	3.49	3.9	4.44	5.15
Π^S	0.67	0.50	0.42	0.48	3.58	4.2	4.99	6.05
Schild’s Ladder	0.71	0.51	0.45	0.49	3.57	4.14	4.84	6.21
Conjugate	0.8	0.63	0.47	0.32	3.43	3.83	4.36	5.04

a white/black matter region. The atrophy was simulated by decreasing the gray layer thickness on four subsequent time points to generate the sequence S_i , $i = 1 \dots 4$ (Figure 7). The longitudinal trajectories of deformation fields $\exp(X^i)$ were then evaluated by registering the images to the baseline with the Log-Demons algorithm [39]. The sequence of deformations $\exp(X^i)$ was then transported on a target ellipsoidal geometry E_0 along the inter-subject deformation $\exp(Y)$ such that $\exp(Y) * S_0 = E_0$ (Figure 6). The transport methods that we tested were:

- Right Cartan connection (right translation) Π^R ;
- Left Cartan connection (left translation) Π^L and Cartan symmetric connection Π^S implemented with the recursive scheme;
- Conjugate action
 $Conj(\exp(X^i)) = \exp(Y)\exp(X^i)\exp(Y)^{-1}$;
- the Schild’s Ladder, which operates along the “diagonal” inter-subject deformations $\exp(Y^i)$ such that $\exp(Y^i) * S_i = E_0$ (Figure 6)

The methods were quantitatively assessed by evaluating the features of interest in the ellipsoidal gray layer: the average L^2 norm of the transported SVF, the Jacobian determinant, log-Jacobian determinant and elastic energy of the associated deformation fields. Since we are interested in preserving the interesting features of the transported trajectories, the transported quantities were compared to the original values in the reference sphere space. Moreover, the stability of the methods was tested by checking the scalar spatial maps associated to the features (this involves resampling).

Table 4 shows the accuracy of the transport methods in the preservation of the measure of changes in the

gray matter layer. Among the different methods, the transport Π^R was the most accurate in preserving the average measures, while the Schild’s Ladder performed better on the Log-Jacobian.

From the inspection of the related scalar log-Jacobian maps (Figure 7), the transport Π^L is the less stable and leads to noisy maps. Moreover, we notice that the areas of expansions does not fit the boundary of the ellipsoid. On the other hand, the transport Π^R leads to smooth maps of changes, consistent with the target geometry, while the transport Π^S lies “in between”, as one could reasonably expect. The Schild’s Ladder leads to smooth maps as well, although the inner spherical shape seems corrupted for higher deformations. This could explain the lower performance on the quantitative measurements for the time points 3 and 4. Finally, the log-Jacobian maps associated to the conjugate actions are smooth but fail to preserve the target ellipsoidal geometry, especially for the higher deformations.

5.4 Real Longitudinal Changes in Alzheimer’s Disease

In this section the parallel transports Π^L, Π^R, Π^S , and the Schild’s Ladder were applied to an example of real longitudinal brain atrophy. Baseline and follow-up scans from a subject affected by Alzheimer’s disease were registered with the Symmetric Log-Demons algorithm. Further details on the image acquisition protocol are given in section 4.3. The longitudinal atrophy encoded by the SVF was then transported along the inter-subject registration on a different reference anatomical space (Figure 8).

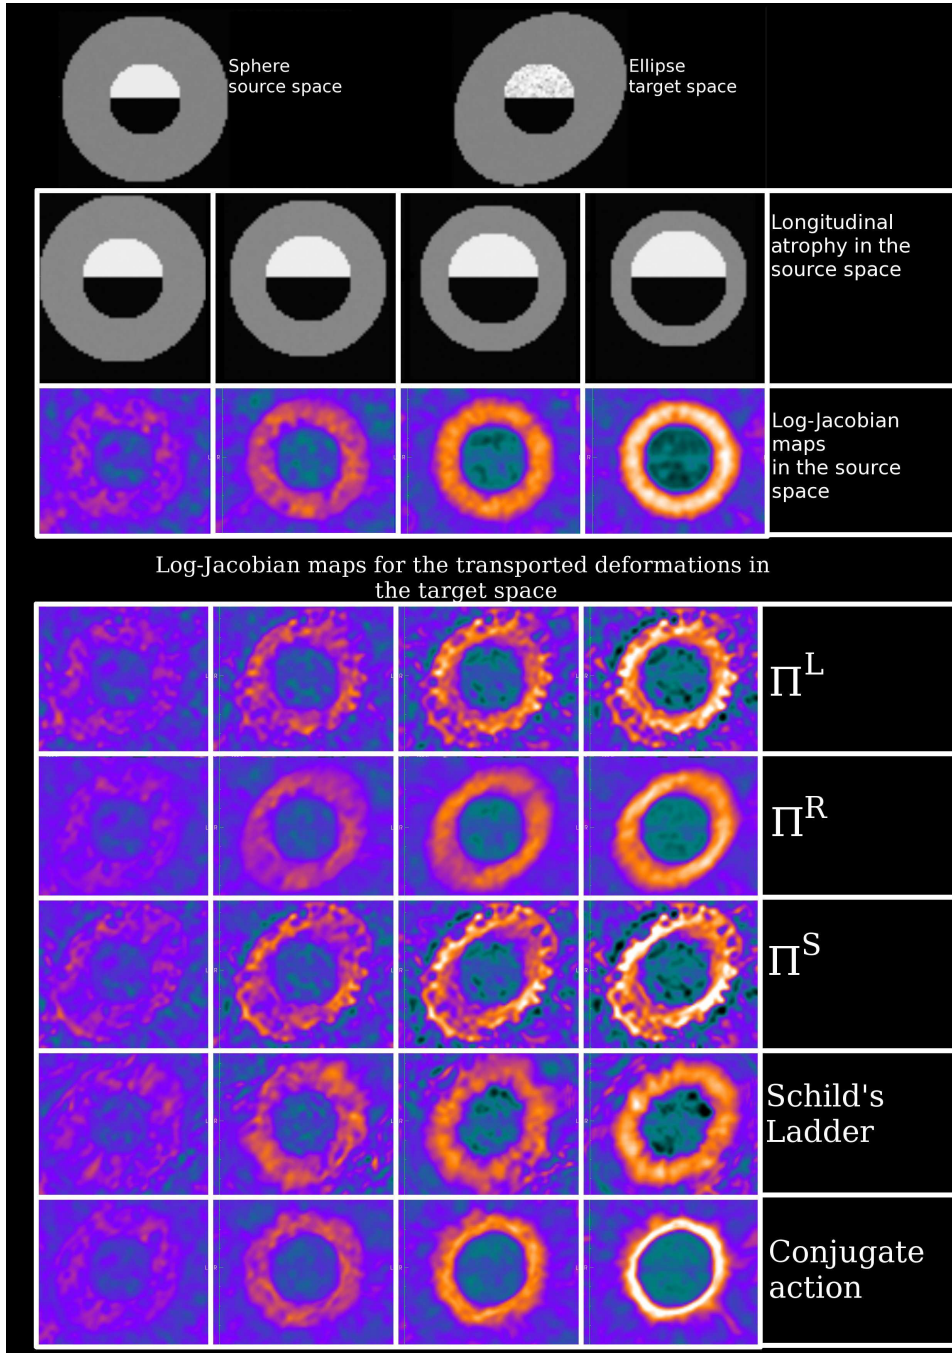


Fig. 7 Top row: Spherical source and ellipsoidal target geometrical references. From top to bottom: Longitudinal atrophy sequence in the spherical space, associated log-Jacobian determinant scalar maps, and log-Jacobian determinant maps associated to the different methods of transport.

As in the synthetic experiment, all the above methods provide results that are consistent with the original trajectory. As already seen in the synthetic case, the transport Π^L seems to introduce some noise which lead to increased velocities (for instance in the posterior cingulate and on the left side of the cortex), which indicates that numerical issues still play a central role.

Interestingly, we notice in the posterior part of the ventricles in the source space a rotational movement pointing to the medial axis of the brain which is captured by Π^L , Π^S , and by the Schild's Ladder, but which seems missing in Π^R . This observation suggests that the simple resampling of the vector field might provide only a partial representation of the deformation, and is poten-




Fig. 8 Parallel transport of longitudinal atrophy. Longitudinal ventricular expansion in the source space and related parallel transport given by the different methods.

tially missing some features of interest for a subsequent analysis.

5.5 Cartan Parallel Transport Along Metric vs Group Geodesics

Finally, we investigate the differences in transporting a longitudinal trajectory along different paths, namely the LDDMM metric geodesics and the one-parameter subgroups that best register two images. We restrict our study to the parallel transport of the Cartan connections. The LDDMM parallel transport could be computed by the equations of Jacobi fields [40] or could be approximated using the Schild's ladder algorithm by explicitly computing metric geodesics at each step of the ladder. However, the computational complexity is much worse because the simplifications used for the group exponential with the BCH in [22] are not valid any more for Riemannian geodesics. Thus, we focus here on the comparison of the different type of Cartan parallel transport along group and metric geodesics without comparing LDDMM versus Cartan parallel transports.

From the geometrical point of view, parallel transporting a vector along different trajectories generally leads to different parallel vectors (Figure 9). However, the similarity of the deformation trajectories evaluated by different methods could lead to according results. Moreover, the practical application to the image registration might hide the effect of these theoretical issues. Evaluating the impact of the choice of the trajectory is

therefore of interest to understand the effect of different parameterizations and to generalize the transport schemes to diverse registration settings.

We use here the Synthetic example of the Paragraph 5.3. The spherical source space was registered to the target ellipse with the Demons and the AtlasWerks methods. The related tangent parameterization define two different paths in the space of diffeomorphisms, respectively the one-parameter subgroup associated to the SVF, and the metric geodesic associated to the time varying velocity field $v(t)$. In the LDDMM case, the time-varying trajectory is sampled in time to give a series of n stationary vectors $v_i = v(t_i)$, $i = 1, \dots, n$, and the final deformation φ is generated from the composition:

$$\varphi = \exp\left(\frac{v_{n-1}}{n}\right) \circ \exp\left(\frac{v_{n-2}}{n}\right) \dots \circ \exp\left(\frac{v_0}{n}\right)$$

Therefore we decompose here the time varying process into successive stationary ones, in order to apply the same methods derived from the Lie group theory. Thus, the parallel transport methods illustrated in the previous section can be applied on each trajectory v_i to iteratively transport a longitudinal trajectory along the metric geodesic.

Figure 10 shows the resulting vectors transported from the sphere to the ellipse space along the SVF one-parameter subgroup and the LDDMM metric geodesic: transporting along the different paths does not seems to introduce relevant differences. We notice that the right transport Π^R and the Schild's Ladder provide the most consistent and robust results. The left (Π^L) and the symmetric (Π^S) transports appear to be more sensitive to the very high Jacobian matrix variations of the LDDMM trajectory, though they are still consistent with respect to the deformation pattern. The “concentration” of the transported vector field at the high momentum places along the LDDMM trajectory




Fig. 9 The parallel transport of the vector v closely depends on the chosen trajectory, and generally transporting along different curves lead to different parallel vectors.




Fig. 10 Parallel transport of the longitudinal trajectory on the sphere to the ellipse target space along the trajectories defined by the SVF one-parameters subgroup and the LDDMM metric geodesic. The Schild's Ladder and the Right Cartan connection lead to the more consistent results in term of the resulting transported vectors.

could be interesting for statistical studies but might also disperse more easily the information even in case of very small inter-subject matching errors.

This very simple experiment shows that the SVF tools can also be used in different settings such as with LDDMM, where geodesics can be approximated by a sequence of stationary approximations. Moreover, the similarity of the transport on the different paths suggests that the parallel transport is more influenced by the type of parallel transport than by the inter-subject registration method. It would be quite interesting to verify if this still holds for the LDDMM transport. In all cases, the numerics behind the parallel transport scheme seems to be a key issue.

6 Conclusions

This study investigates the theoretical background which underlies diffeomorphic registration based on stationary

velocity fields. We illustrate the use of Lie group concepts to derive effective and efficient solutions for computational anatomy. We showed that the one-parameter subgroups (and their left and right translations) are geodesics of the a family of connections called Cartan connections. This mathematical setting is completely consistent with the Lie group operations (left and right composition, inversion) and leads to a closed form solution for the parallel transport. Moreover, the geodesics of the Cartan connections are "metric-free", and the parallel transport is not related to the preservation of metric properties. This affine geometric mathematical setting differs from invariant Riemannian manifolds used for instance in LDDMM, for which the choice of left or right invariance, as well as the choice of the metric, lead to different geodesics. Among the three canonical (right, left and symmetric) Cartan connections, the transport with the right one was the smoother, due to the simple computational requirements. However, this connection is related to a specific geometry where

the group is flat (no curvature), but has torsion. From a theoretical point of view, it is widely accepted in other domains (e.g. general relativity and gravitation) that working with a symmetric connection is preferable than working in a space with torsion. However, we believe that it would be worth verifying by experiments on real data that the symmetric Cartan connection indeed leads to a better description of the groupwise anatomy than non-symmetric ones. Such test could be performed on disease classification experiments for instance, where the statistical power of the separation would designate the optimal parallel transport method.

The experimental results highlight also the trade-off between the choice of proper mathematical constructions and the related numerical implementation issues. For instance, the left and symmetric Cartan transports could benefit from more robust numerical schemes for the computation of the differential quantities like the Jacobian matrix, which could lead to more stable and accurate results. On the registration side, our experiments showed that the SVF and LDDMM settings performed very similarly in the longitudinal case. This result suggests that, when dealing with small deformations, stationary and time varying parametrization lead to negligible discrepancies. Moreover, even in case of the inter-subject registration, the different parametrization performed similarly for the resulting intensity matching. It seems therefore that the choice of the admissible transformations is decisively circumscribed by the smoothness constraints and by the numerical implementation, which limit the set of possible anatomical deformations that can be retrieved by a registration algorithm.

We should note that the above conclusions come from a precise choice of the registration parameters (fluid/elastic regularization, incompressibility, . . .), and we cannot exclude that a different tuning might lead to a different scenario. However, the choice made here was motivated by the search of an optimal compromise between registration accuracy and smoothness of the deformations and, among the several configurations tested, the proposed was the most suitable one. Moreover, given the lack of a ground truth for the inter-subject registration problem, the reliable comparison of different registration techniques can be assessed only by statistical measures on large datasets, which goes beyond the scope of the presented study.

References

1. URL <http://www.sci.utah.edu/software.html>. AtlasWerks: A set of high-performance tools for diffeomorphic 3D image registration and atlas building. Scientific Computing and Imaging Institute (SCI)
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A Log-Euclidean framework for statistics on diffeomorphisms. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI, vol. 9, pp. 924–931 (2006)
3. Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95 – 113 (2007)
4. Bossa, M., Hernandez, M., Olmos, S.: Contributions to 3D diffeomorphic atlas estimation: Application to brain images. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI, vol. 10, pp. 667–674 (2007)
5. Bossa, M., Zacur, E., Olmos, S.: On changing coordinate systems for longitudinal tensor-based morphometry. In: Proceedings of Spatio Temporal Image Analysis Workshop (STIA) (2010)
6. Cachier, P., Ayache, N.: Isotropic energies, filters and splines for vectorial regularization. *J. of Math. Imaging and Vision* **20**, 251–265 (2004)
7. do Carmo, M.: *Riemannian Geometry*. Mathematics. Birkhäuser, Boston, Basel, Berlin (1992)
8. Cartan, E., Schouten, J.: On the geometry of the group-manifold of simple and semi-simple groups. *Proc. Akad. Wetensch. Amsterdam* **29**, 803–815 (1926)
9. Durrleman, S., Fillard, P., Pennec, X., Trounev, A., Ayache, N.: Registration, atlas estimation and variability analysis of white matter fiber bundles modeled as currents. *NeuroImage* **55**(3), 1073–1090 (2011)
10. Durrleman, S., Pennec, X., Trounev, A., Gerig, G., Ayache, N.: Spatiotemporal atlas estimation for developmental delay detection in longitudinal datasets. In: Medical Image Computing and Computer-Assisted Intervention - MICCAI, vol. 12, pp. 297–304 (2009)
11. Gallot, S., Hulin, D., Lafontaine, J.: *Riemannian Geometry*, 2nd edition edn. Springer Verlag (1993)
12. Galluzzi, S., Testa, C., Boccardi, M., Bresciani, L., Benussi, L., Ghidoni, R., Beltramello, A., Bonetti, M., Bono, G., and G. Magnani and G. Minonzio, A.F., Piovani, E., Binetti, G., Frisoni, G.: The Italian Brain Normative Archive of structural MR scans: norms for medial temporal atrophy and white matter lesions. *Aging Clin Exp Res.* **21**(4-5), 264–5 (2009)
13. Helgason, S.: *Differential Geometry, Lie groups, and Symmetric Spaces*. Academic Press (1978)
14. Helgason, S.: *Differential geometry, Lie groups, and symmetric spaces*. American Mathematical Soc (1978)
15. Hernandez, M., Bossa, M., Olmos, S.: Registration of anatomical images using paths of diffeomorphisms parameterized with stationary vector field flows. *International Journal of Computer Vision* **85**, 291–306 (2009)
16. Joshi, S., Davis, B., Jomier, M., G. Gerig: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* pp. S151–S160 (2004)
17. Joshi, S., Miller, M.: Landmark matching via large deformation diffeomorphisms. *IEEE Transactions on Image Processing* **9**(8), 1357–1370 (2000)
18. Khesin, B.A., Wendt, R.: The Geometry of Infinite Dimensional Lie groups, *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge / A Series of Modern Surveys in Mathematics*, vol. 51. Springer Verlag (2009)
19. Kheifets, A., Miller, W., Newton, G.: Schild's Ladder parallel transport for an arbitrary connection. *International Journal of Theoretical Physics* **39**(12), 41–56 (2000)

20. Kolev, B.: Groupes de lie et mécanique. <http://www.cmi.univ-mrs.fr/~kolev/> (2007). Notes of a Master course in 2006-2007 at Université de Provence
21. KSchild, A.: Tearing geometry to pieces: More on conformal geometry. unpublished lecture at Jan. 19 1970 Princeton Univesity relativity seminar (1970)
22. Lorenzi, M., Ayache, N., Pennec, X.: Schild's Ladder for the parallel transport of deformations in time series of images. In: *Information Processing in Medical Imaging - IPMI*, vol. 22, pp. 463–474 (2011)
23. Lorenzi, M., Frisoni, G., Ayache, N., Pennec, X.: Mapping the effects of $A\beta_{1-42}$ levels on the longitudinal changes in healthy aging: hierarchical modeling based on stationary velocity fields. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, vol. 14, pp. 663–670 (2011)
24. Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N.: LogDemons revisited: Consistent regularisation and incompressibility constraint for soft tissue tracking in medical images. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, vol. 13, pp. 652–659 (2010)
25. Mansi, T., Pennec, X., Sermesant, M., Delingette, H., Ayache, N.: iLogDemons: A Demons-based registration algorithm for tracking incompressible elastic biological tissues. *International Journal of Computer Vision* **92**(1), 92–111 (2011)
26. Mansi, T., Voigt, I., Leonardi, B., Pennec, X., Durrleman, S., Sermesant, M., Delingette, H., Taylor, A., Boudjemline, Y., Pongiglione, G., Ayache, N.: A statistical model for quantification and prediction of cardiac remodelling: Application to tetralogy of fallot. *IEEE Transactions on Medical Images* **30**(9), 1605–1616 (2011)
27. Miller, M., Trounev, A., Younes, L.: On the metrics and Euler-Lagrange equations of computational anatomy. *Annu Rev Biomed Eng* **4**(1), 375–405 (2002)
28. Milnor, J.: Remarks on infinite-dimensional Lie groups. In: *Relativity, Groups and Topology*, pp. 1009–1057. Les Houches (1984)
29. Misner, C.W., Thorne, K.S., Wheeler, J.: *Gravitation*. W.H. Freeman and Compagny (1973)
30. Modat, M., Ridgway, G., Daga, P., Cardoso, M., Hawkes, D., Ashburner, J., Ourselin, S.: Log-euclidean free-form deformation. In: *Proc. of SPIE Medical Imaging 2011*. SPIE (2011)
31. Pennec, X., Arsigny, V.: Exponential Barycenters of the Canonical Cartan Connection and Invariant Means on Lie Groups. In: F. Barbaresco, A. Mishra, F. Nielsen (eds.) *Matrix Information Geometry*. Springer (2012)
32. Postnikov, M.M.: *Geometry VI: Riemannian Geometry*. Encyclopedia of mathematical science. Springer (2001)
33. Rao, A., Chandrashekhara, R., Sanchez-Hortiz, G., Mohiaddin, R., aljabar, P., Hajnal, J., Puri, B., Rueckert, D.: Spatial trasformation of motion and deformation fields using nonrigid registration. *IEEE Transactions on Medical Imaging* **23**(9), 1065–1076 (2004)
34. Schmid, R.: Infinite dimensional lie groups with applications to mathematical physics. *J. Geometry and Symmetry in Physics* **1**, 167 (2004)
35. Schmid, R.: Infinite-Dimensional Lie Groups and Algebras in Mathematical Physics. *Advances in Mathematical Physics* **2010**, 1–36 (2010). DOI 10.1155/2010/280362
36. Seiler, C., Pennec, X., Reyes, M.: Geometry-Aware Multi-scale Image Registration Via OBBTree-Based Polyaffine Log-Demons. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, vol. 14, pp. 631–638 (2011)
37. Shattuck, D., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K., Poldrack, R., Bilder, R., Toga, A.: Construction of a 3D probabilistic atlas of human cortical structures. *NeuroImage* **39**(3), 1064–1080 (2008)
38. Thompson, P., Ayashi, K., Zubicaray, G., Janke, A., Rose, S., Semple, J., Herman, D., Hong, M., Dittmer, S., Dodrell, D., Toga, A.: Dynamics of gray matter loss in Alzheimer's disease. *The Journal of Neuroscience* **23**(3), 994–1005 (2003)
39. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric Log-domain diffeomorphic registration: A Demons-based approach. In: *Medical Image Computing and Computer-Assisted Intervention - MICCAI*, vol. 11, pp. 754–761 (2008)
40. Younes, L.: Jacobi fields in groups of diffeomorphisms and applications. *Q. Appl. Math* pp. 113–134 (2007)
41. Younes, L.: *Shapes and diffeomorphisms*. No. 171 in *Applied Mathematical Sciences*. Springer (2010)
42. Younes, L., Qiu, A., Winslow, R., Miller, M.: Transport of relational structures in groups of diffeomorphisms. *J Math Imaging Vis* **32**(1), 41–56 (2008)