Finite-Dimensional Lie Algebras for Fast Diffeomorphic Image Registration

Miaomiao Zhang and P. Thomas Fletcher

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT

Abstract. This paper presents a fast geodesic shooting algorithm for diffeomorphic image registration. We first introduce a novel finite-dimensional Lie algebra structure on the space of bandlimited velocity fields. We then show that this space can effectively represent initial velocities for diffeomorphic image registration at much lower dimensions than typically used, with little to no loss in registration accuracy. We then leverage the fact that the geodesic evolution equations, as well as the adjoint Jacobi field equations needed for gradient descent methods, can be computed entirely in this finite-dimensional Lie algebra. The result is a geodesic shooting method for large deformation metric mapping (LDDMM) that is dramatically faster and less memory intensive than state-of-the-art methods. We demonstrate the effectiveness of our model to register 3D brain images and compare its registration accuracy, runtime, and memory consumption with leading LDDMM methods. We also show how our algorithm breaks through the prohibitive time and memory requirements of diffeomorphic atlas building.

1 Introduction

Deformable image registration is a fundamental tool in medical image analysis that is used for several tasks, including anatomical comparisons across individuals, alignment of functional data to a reference coordinate system, and atlas-based image segmentation. In many applications, it is desirable that image transformations be diffeomorphisms, i.e., differentiable, bijective mappings with differentiable inverses. Such diffeomorphic mappings ensure several properties of the transformed images: (1) topology of objects in the image remain intact; (2) no non-differentiable artifacts, such as creases or sharp corners, are created; and (3) the process can be inverted, for instance, to move back and forth between two individuals, or between an atlas and an individual. An elegant mathematical formulation for diffeomorphic image registration is that of Large Deformation Diffeomorphic Metric Mapping (LDDMM), first proposed by Beg et al. [5]. In this setting, the group of diffeomorphisms is equipped with a Riemannian metric, giving rise to a variational principle that expresses the optimal image registration as a geodesic flow. The result is a distance metric between images that quantifies their geometric similarity. Having a distance metric is a critical component in statistical analysis of anatomical shape, including regression, longitudinal analysis, and group comparisons, as it provides a mathematical foundation for fitting a statistical model to image data via minimization of the sum-of-squared residual distances.

A major barrier to the widespread use of LDDMM, especially in large imaging studies, is its high computational cost and large memory footprint. The original algorithm by Beg et al. computes a geodesic path by gradient descent of a time-varying velocity field. This requires expensive numerical solutions to partial differential equations for the gradient evaluation, as well as time integration of the velocities to compute the diffeomorphic transformation, all of which must be done on dense spatial grids for numerical accuracy. Furthermore, convergence of the algorithm can be slow and prone to getting stuck in local minima, as is often the case with relaxation methods. Addressing some of these weaknesses, Vialard et al. [14] introduced a geodesic shooting algorithm for diffeomorphic image matching. This takes advantage of the fact that a geodesic is determined by its initial momentum at time zero via the geodesic evolution equations. Using a control theory formulation, Vialard et al. then derive the necessary adjoint equations to carry gradients of the image match at the endpoint of the geodesic back to a gradient in the initial momentum. They demonstrated that this led to better convergence and more reliable estimates of the initial momentum. It also avoids having to store the entire time-varying velocity field from one iteration to the next, as just the initial momentum is sufficient to encode the current geodesic. However, despite these advantages, the geodesic shooting and backward adjoint equations must also be solved numerically on a dense grid and are still prohibitively time consuming.

Another class of diffeomorphic registration methods are the “greedy” algorithms, that is, algorithms that iteratively apply gradient updates to a single deformation field, rather than update a full time-dependent flow each iteration. Algorithms in this class include the original diffeomorphic image registration algorithm of Christensen et al. [6] and the diffeomorphic demons algorithm of Vercauteren et al. [13]. Greedy methods are much faster and more memory efficient than LDDMM, but they do not minimize a global variational problem and do not provide a distance metric between images. They also lack the initial velocity parameterization that geodesic methods possess. Such a parameterization is important for statistical analysis because it represents deformations in a linear vector space, which is amenable to statistical models such as principal component analysis [12] and regression [10]. Arsigny et al. [2] introduced the concept of a stationary velocity field representation for diffeomorphisms and demonstrated that one-parameter subgroup flows of diffeomorphisms could be computed efficiently. A similar strategy is used by the DARTEL image registration method of Ashburner [3]. While stationary velocity fields again are more efficient in time and memory than LDDMM, they do not provide distance metrics on the space of diffeomorphisms.

In this paper, we show that it is possible to have a fast diffeomorphic image registration algorithm that retains the metric properties of LDDMM. To do this, we introduce a novel theoretical framework for diffeomorphic image registration based on representing discretized velocity fields as elements in a finite-dimensional Lie algebra. A somewhat surprising feature of our proposed framework is that not only is our algorithm faster than state-of-the-art LDDMM algorithms, it is actually able to converge to better solutions, i.e., lower values of the registration objective function.

2 Background on LDDMM and Geodesic Shooting

In this section we give a brief overview of LDDMM and set up the notation we will use. We consider images defined on a cyclical domain, i.e., a d-dimensional torus, $\Omega = \mathbb{R}^d/\mathbb{Z}^d$. An image will be a square-integrable function on Ω, that is, an element of $L^2(\Omega, \mathbb{R})$. Let $\mathrm{Diff}^\infty(\Omega)$ denote the Lie group of smooth diffeomorphisms on Ω. The Lie algebra for $\mathrm{Diff}^\infty(\Omega)$ is the space $V = \mathfrak{X}^\infty(T\Omega)$ of smooth vector fields on Ω, which is the tangent space at the identity transform. We equip V with the weak Sobolev metric:

$$\langle v, w \rangle_V = \int_\Omega (Lv(x), w(x))dx, \tag{1}$$

where $L = (I - \alpha\Delta)^s$ is a symmetric, positive-definite, differential operator for some scalar, $\alpha > 0$, and integer power, s. The dual to the tangent vector v is a momentum, $m = Lv \in V^*$. The notation (m, v) denotes the pairing of a momentum vector $m \in V^*$ with a tangent vector $v \in V$. The metric L is an invertible operator, and $K = L^{-1}$ maps the momentum vector $m \in V^*$ back to the tangent vector $v = Km \in V$.

LDDMM: In the LDDMM framework, we will consider diffeomorphisms that are generated by flows of time-varying velocity fields from V. More specifically, consider a time-varying velocity field, $v : [0, 1] \to V$, then we may define the flow $t \mapsto \phi_t \in \mathrm{Diff}^s(\Omega)$ as a solution to the equation

$$\frac{d\phi_t}{dt}(x) = v_t \circ \phi_t(x).$$

The registration between two images, $I_0, I_1 \in L^2(\Omega, \mathbb{R})$, is the minimizer of the energy,

$$E(v_t) = \int_0^1 \|v_t\|_V^2 \, dt + \frac{1}{2\sigma^2} \|I_0 \circ \phi_1^{-1} - I_1\|_{L^2}^2. \tag{2}$$

Geodesic Shooting: Given an initial velocity, $v_0 \in V$, at $t = 0$, the geodesic path $t \mapsto \phi_t \in \mathrm{Diff}^\infty(\Omega)$ under the right-invariant Riemannian metric (1) is uniquely determined by the Euler-Poincaré equations (EPDiff) [1, 9],

$$\begin{aligned} \frac{\partial v}{\partial t} = -\mathrm{ad}_v^\dagger v &= -K\mathrm{ad}_v^* m \\ &= -K \left[(Dv)^T m + Dm\, v + m \operatorname{div} v \right], \end{aligned} \tag{3}$$

where D denotes the Jacobian matrix, and the operator ad^* is the dual of the negative Lie bracket of vector fields,

$$\mathrm{ad}_v w = -[v, w] = Dvw - Dwv. \tag{4}$$

By integrating equation (3) forward in time, we generate a time-varying velocity $v_t : [0, 1] \to V$, which itself is subsequently integrated in time by the rule $(d\phi_t(x)/dt) = v_t \circ \phi_t(x)$ to arrive at the geodesic path, $\phi_t(x) \in \mathrm{Diff}^s(\Omega)$. Details are found in [14, 15]. Noting that the geodesic ϕ_t is fully determined by the initial condition v_0, we can rewrite the LDDMM image matching objective function (2) in terms of the initial velocity v_0 as

$$E(v_0) = (Lv_0, v_0) + \frac{1}{2\sigma^2} \| I_0 \circ \phi_1^{-1} - I_1 \|^2. \tag{5}$$

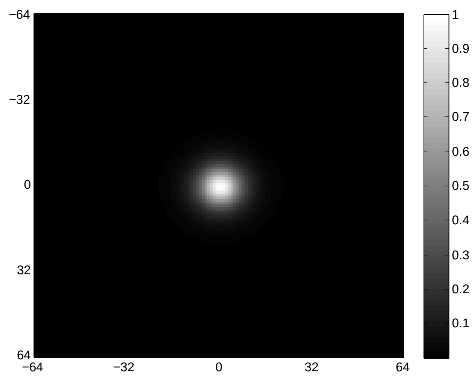

Fig. 1: Fourier coefficients of the discretized K operator on a 128×128 grid, with parameters $\alpha = 3$, $s = 3$.

3 Finite-Dimensional Lie Algebras of Diffeomorphism

While the Lie algebra $V = \mathfrak{X}^\infty(\Omega)$ is an infinite-dimensional vector space, we must of course approximate smooth vector fields in V with finite-dimensional discretizations in order to represent them on a computer. In this section, we show that careful choices for the discretization of vector fields and their corresponding Lie brackets leads to a representation that is itself a finite-dimensional Lie algebra. The key observation is that the K operator is a low-pass filter, and as such, suppresses high frequency components in the velocity fields (see, for example, Figure 1). As the K operator appears as the last operation on the right-hand side of the EPDiff equation (3), we can see that the velocity fields in the geodesic evolution do not develop high frequency components. This suggests that a standard implementation of geodesic shooting, using high-resolution velocity fields, wastes a lot of effort computing the high frequency components, which just end up being forced to zero by K. Instead, we propose to use a low-dimensional discretization of velocity fields as bandlimited signals in the Fourier domain.

More specifically, let $\tilde{V}$ denote the space of bandlimited velocity fields on Ω, with frequency bounds $N_1, N_2, \ldots, N_d$ in each of the dimensions of Ω. It will be convenient to represent an element $\tilde{v} \in \tilde{V}$ in the Fourier domain. That is, let $\tilde{v} \in \tilde{V}$ be a multidimensional array: $\tilde{v}_{k_1,k_2,\ldots,k_d} \in \mathbb{C}^d$, where $k_i \in 0, \ldots, N_i - 1$ is the frequency index along the ith axis. Note that to ensure $\tilde{v}$ represents a real-valued vector field in the spatial domain, we have the constraint that $\tilde{v}_{k_1,\ldots,k_d} = \tilde{v}^*_{N_1-k_1,\ldots,N_d-k_d}$, where $*$ denotes the complex conjugate. There is a natural inclusion mapping, $\iota : \tilde{V} \to V$, of $\tilde{V}$ into the space V of smooth vector fields, given by the Fourier series expansion:

$$\iota(\tilde{v})(x_1, \ldots, x_d) = \sum_{k_1=0}^{N_1} \cdots \sum_{k_d=0}^{N_d} \tilde{v}_{k_1,\ldots,k_d} e^{2\pi k_1 x_1} \cdots e^{2\pi k_d x_d}. \tag{6}$$

Next, we define an operator that is a discrete analog of (4), the Lie bracket on continuous vector fields. First, we will denote by $\tilde{D}\tilde{v}$ the central difference Jacobian matrix of a discrete vector field, $\tilde{v} \in \tilde{V}$. This can be computed in the discrete Fourier domain as a tensor product $\tilde{D}\tilde{v} = \eta \otimes \tilde{v}$, where $\eta \in \tilde{V}$ is given by

$$\eta_{k_1,k_2,\ldots,k_d} = (i \sin(2\pi k_1), \ldots, i \sin(2\pi k_d)).$$

Second, we note that pointwise multiplication of matrix and vector field in the spatial domain corresponds to convolution in the Fourier domain. Because convolution between two bandlimited signals does not preserve the bandlimit, we must follow a convolution operation by truncation back to the bandlimits, N_j, in each dimension. We denote this truncated convolution operator between a matrix and a vector field as $\star$. Now we are ready to define our discrete bracket operator for any two vectors $\tilde{v}, \tilde{w} \in \tilde{V}$ as

$$[\tilde{v}, \tilde{w}] = (\tilde{D}\tilde{v}) \star \tilde{w} - (\tilde{D}\tilde{w}) \star \tilde{v}. \tag{7}$$

The next theorem proves that this operation satisfies the properties to be a Lie bracket on $\tilde{V}$.

Theorem 1. *The vector space $\tilde{V}$, when equipped with the bracket operation* (7)*, is a finite-dimensional Lie algebra. That is to say, $\forall\, \tilde{x}, \tilde{y}, \tilde{z} \in \tilde{V}$ and $a, b \in \mathbb{R}$, the following properties are satisfied:*

(a) Linearity: $[a\tilde{x} + b\tilde{y}, \tilde{z}] = a\,[\tilde{x}, \tilde{z}] + b\,[\tilde{y}, \tilde{z}]$,
(b) Anticommutativity: $[\tilde{x}, \tilde{y}] = -[\tilde{y}, \tilde{x}]$,
(c) Jacobi identity: $[\tilde{x}, [\tilde{y}, \tilde{z}]] + [\tilde{z}, [\tilde{x}, \tilde{y}]] + [\tilde{y}, [\tilde{z}, \tilde{x}]] = 0$.

Proof. Linearity and anticommutativity are immediate. We have *(a) Linearity:*

$$\begin{aligned}[a\tilde{x} + b\tilde{y},\, \tilde{z}] &= \tilde{D}(a\tilde{x} + b\tilde{y}) \star \tilde{z} - \tilde{D}\tilde{z} \star (a\tilde{x} + b\tilde{y}) \\ &= a(\tilde{D}\tilde{x} \star \tilde{z} - \tilde{D}\tilde{z} \star \tilde{x}) + b(\tilde{D}\tilde{y} \star \tilde{z} - \tilde{D}\tilde{z} \star \tilde{y}) \\ &= a[\tilde{x},\, \tilde{z}] + b[\tilde{y},\, \tilde{z}],\end{aligned}$$

(b) Anticommutativity:

$$[\tilde{x},\, \tilde{y}] = \tilde{D}\tilde{x} \star \tilde{y} - \tilde{D}\tilde{y} \star \tilde{x} = -\left(\tilde{D}\tilde{y} \star \tilde{x} - \tilde{D}\tilde{x} \star \tilde{y}\right) = -[\tilde{y},\, \tilde{x}].$$

(c) Jacobi identity: The proof of the Jacobi identity follows closely that of the continuous case. First, note that the iterated central difference operator results in a third-order tensor:

$$\tilde{D}^2\tilde{x} = \tilde{D}\tilde{D}\tilde{x} = \eta \otimes \eta \otimes \tilde{x}.$$

Much like the Hessian tensor of a vector-valued function in the continuous case, the discrete Hessian is also symmetric with respect to contraction with a pair of vectors. That is,

$$(\tilde{D}^2\tilde{x} \star \tilde{y}) \star \tilde{z} = (\tilde{D}^2\tilde{x} \star \tilde{z}) \star \tilde{y},$$

where the convolution between $\tilde{D}^2\tilde{x}$ and $\tilde{y}$ is now analogous to the pointwise contraction of a third-order tensor field with a vector field in spatial domain.

Next, we note that the product rule of differentiation also carries over to the discrete Fourier representation, and we have the identity

$$\tilde{D}(\tilde{D}\tilde{x} \star \tilde{y}) = \tilde{D}^2\tilde{x} \star \tilde{y} + \tilde{D}\tilde{x} \star \tilde{D}\tilde{y},$$

where the second convolution operator is analogous to pointwise matrix field multiplication in the spatial domain. We then have

$$\begin{aligned}[\tilde{x},\, [\tilde{y},\, \tilde{z}]] &= [\tilde{x},\, \tilde{D}\tilde{y} \star \tilde{z} - \tilde{D}\tilde{z} \star \tilde{y}] \\ &= \tilde{D}\tilde{x} \star (\tilde{D}\tilde{y} \star \tilde{z} - \tilde{D}\tilde{z} \star \tilde{y}) - \tilde{D}(\tilde{D}\tilde{y} \star \tilde{z} - \tilde{D}\tilde{z} \star \tilde{y}) \star \tilde{x} \\ &= \tilde{D}\tilde{x} \star \tilde{D}\tilde{y} \star \tilde{z} - \tilde{D}\tilde{x} \star \tilde{D}\tilde{z} \star \tilde{y} - (\tilde{D}^2\tilde{y} \star \tilde{z}) \star \tilde{x} - \tilde{D}\tilde{y} \star \tilde{D}\tilde{z} \star \tilde{x} \\ &\quad + (\tilde{D}^2\tilde{z} \star \tilde{y}) \star \tilde{x} + \tilde{D}\tilde{z} \star \tilde{D}\tilde{y} \star \tilde{x}\end{aligned} \tag{8}$$

Similarly, we rewrite the other two terms as

$$\begin{aligned}[\tilde{z},\, [\tilde{x},\, \tilde{y}]] =& \tilde{D}\tilde{z} \star \tilde{D}\tilde{x} \star \tilde{y} - \tilde{D}\tilde{z} \star \tilde{D}\tilde{y} \star \tilde{x} - (\tilde{D}^2\tilde{x} \star \tilde{y}) \star \tilde{z} - \tilde{D}\tilde{x} \star \tilde{D}\tilde{y} \star \tilde{z} \\ &+ (\tilde{D}^2\tilde{y} \star \tilde{x}) \star \tilde{z} + \tilde{D}\tilde{y} \star \tilde{D}\tilde{x} \star \tilde{z}\end{aligned} \tag{9}$$

$$\begin{aligned}[\tilde{y},\, [\tilde{z},\, \tilde{x}]] =& \tilde{D}\tilde{y} \star \tilde{D}\tilde{z} \star \tilde{x} - \tilde{D}\tilde{y} \star \tilde{D}\tilde{x} \star \tilde{z} - (\tilde{D}^2\tilde{z} \star \tilde{x}) \star \tilde{y} - \tilde{D}\tilde{z} \star \tilde{D}\tilde{x} \star \tilde{y} \\ &+ (\tilde{D}^2\tilde{x} \star \tilde{z}) \star \tilde{y} + \tilde{D}\tilde{x} \star \tilde{D}\tilde{z} \star \tilde{y}\end{aligned} \tag{10}$$

Finally, by combining the equations (8), (9), (10), and using the symmetric rule above, we obtain

$$[\tilde{x},\, [\tilde{y},\, \tilde{z}]] + [\tilde{z},\, [\tilde{x},\, \tilde{y}]] + [\tilde{y},\, [\tilde{z},\, \tilde{x}]] = 0.$$

4 Estimation of Diffeomorphic Image Registration

In this section, we present a geodesic shooting algorithm for diffeomorphic image registration using our finite-dimensional Lie algebra representation. This is a gradient descent algorithm on an initial velocity $\tilde{v}_0 \in \tilde{V}$. Geodesic shooting of $\tilde{v}_0$ proceeds entirely in the reduced finite-dimensional Lie algebra, producing a time-varying velocity, $t \mapsto \tilde{v}_t \in \tilde{V}$. Such a geodesic path in $\tilde{V}$ can consequently generate a flow of diffeomorphisms, $t \mapsto \phi_t \in \mathrm{Diff}^\infty(\Omega)$, in the following way. Using the inclusion mapping $\iota : \tilde{V} \to V$ defined in (6), we can generate the diffeomorphic flow as

$$\frac{d\phi_t(x)}{dt} = \iota(\tilde{v}_t) \circ \phi_t(x), \qquad x \in \Omega.$$

This leads to a modification for the energy function (5) for LDDMM, where we now parameterize diffeomorphisms by the finite-dimensional velocity $\tilde{v}_0$:

$$E(\tilde{v}_0) = \|\tilde{v}_0\|^2_{\tilde{L}} + \frac{1}{2\sigma^2}\|I_0 \circ \phi_1^{-1} - I_1\|^2. \tag{11}$$

Here the metric on $\tilde{V}$ is given by the discretized version of the $\tilde{L}$ operator. The Fourier transformation of $L = (-\alpha\Delta + I)^s$ is a diagonal operator. Discretizing this operator by only keeping the frequencies up to our bandlimits, N_j, we get a diagonal matrix $\tilde{L}$. Analogous to the L operator, this $\tilde{L} : \tilde{V} \to \tilde{V}^*$ maps a tangent vector in Fourier domain to its dual momentum vector $\tilde{m}$. For a 3D grid, the coefficients $\tilde{L}_{ijk}$ of this operator at coordinate (i, j, k) in the Fourier domain are

$$\tilde{L}_{ijk} = \left[-2\alpha\left(\cos\frac{2\pi i}{W} + \cos\frac{2\pi j}{H} + \cos\frac{2\pi k}{D}\right) + 7\right]^s,$$

where W, H, D are the dimension of each direction. Vice versa, the Fourier coefficients of the K operator are $\tilde{K}_{ijk} = \tilde{L}^{-1}_{ijk}$.

Before describing the details of our diffeomorphic image matching algorithm, we first provide an outline of the general steps. Beginning with the initialization $\tilde{v}_0 = 0$, the gradient descent algorithm to minimize the energy (11) proceeds by iterating the following:

1. **Forward shooting of $\tilde{v}_0$:** Forward integrate the geodesic evolution equations on $\tilde{V}$ to generate $\tilde{v}_{t_k}$ at discrete time points $t_1 = 0, t_2, \ldots, t_T = 1$.
2. **Compute the diffeomorphism ϕ_1^{-1}:** Compute the inverse diffeomorphism, ϕ_1^{-1}, by integrating the negative velocity field backward in time.
3. **Compute gradient at $t = 1$:** Compute the gradient, $\nabla_{\tilde{v}_1} E$, of the energy (11) at $t = 1$.
4. **Bring gradient to $t = 0$ by adjoint Jacobi field:** Integrate the reduced adjoint Jacobi field equations in $\tilde{V}$ to get the gradient update $\nabla_{\tilde{v}_0} E$.

We note that Steps 2 and 3 are computed at the full resolution of the input images in the spatial domain. However, Steps 1 and 4 to compute the geodesic and the adjoint Jacobi fields are computed entirely in the finite-dimensional space $\tilde{V}$, resulting in greatly reduced computation time and memory requirements. We now provide details for the computations in each of these steps.

4.1 Geodesic Shooting in the Finite-Dimensional Lie Algebra

In the previous section, we introduced a finite-dimensional Lie algebra. Analogous to the EPDiff equation (3), we define its geodesic evolution equation as

$$\frac{\partial \tilde{v}}{\partial t} = -\mathrm{ad}^{\dagger}_{\tilde{v}} \tilde{v} = -\tilde{K} \mathrm{ad}^{*}_{\tilde{v}} \tilde{m}, \tag{12}$$

where the operator $\mathrm{ad}^* : \tilde{V}^* \to \tilde{V}^*$ is the dual of the negative finite-dimensional Lie bracket of vector fields in the Fourier space, and its discrete formulation is

$$\mathrm{ad}^{*}_{\tilde{v}} \tilde{m} = (\tilde{D}\tilde{v})^T \star \tilde{m} + \tilde{D}\tilde{m} \star \tilde{v} + \tilde{m} \star (\tilde{\Gamma}\tilde{v}), \tag{13}$$

where $\tilde{\Gamma}\tilde{v}$ is divergence of the discrete vector field $\tilde{v}$. It is computed as sum of the point-wise multiplication $\tilde{\Gamma}\tilde{v} = \sum_{k_1=0}^{N_1} \cdots \sum_{k_d=0}^{N_d} \eta_i \tilde{v}_i$, where η_i is the Fourier coefficient of the central differential operator on each dimension.

Plugging (13) back into the geodesic evolution equation (12), we have

$$\frac{\partial \tilde{v}}{\partial t} = -\mathrm{ad}^{\dagger}_{\tilde{v}} \tilde{v} = -\tilde{K} \left[(\tilde{D}\tilde{v})^T \star \tilde{m} + \tilde{D}\tilde{m} \star \tilde{v} + \tilde{m} \star (\tilde{\Gamma}\tilde{v}) \right]. \tag{14}$$

4.2 Adjoint Jacobi Fields

The computation of the gradient term in our model requires the adjoint Jacobi fields, which are used to integrate the gradient term at $t = 1$ backward to the initial point $t = 0$. To derive this, consider a variation of geodesics $\gamma : (-\epsilon, \epsilon) \times [0, 1] \to \mathrm{Diff}^{\infty}(\Omega)$, with $\gamma(0, t) = \phi_t$ and $\gamma(s, 0) = \mathrm{Id}$. Such a variation corresponds to a variation of the initial velocity $(d/dt)\gamma(s, 0) = v_0 + s\delta v_0$. The variation $\gamma(s, t)$ produces a "fan" of geodesics, illustrated in Figure 2. Taking the derivative of this variation results in a Jacobi field: $J_v(t) = d\gamma/ds(0, t)$.

In this paper, we use a simple version of reduced Jacobi field from Bullo [7], which is also used by [8]. Under the right invariant metric of diffeomorphisms, we define a vector $U(t) \in \tilde{V}$ as a right trivialized reduced Jacobi field $U(t) = J_{\tilde{v}}(t)\phi(t)^{-1}$, and a variation of the right trivialized reduced velocity $\tilde{v}$ is $\delta\tilde{v}$.

By introducing adjoint variables $\hat{U}, \delta\hat{v} \in \tilde{V}$, we have the adjoint Jacobi equations as

$$\begin{aligned} \frac{d\hat{U}}{dt} &= -\mathrm{ad}^{\dagger}_{\tilde{v}} \hat{U} \\ \frac{d\delta\hat{v}}{dt} &= -\hat{U} - \mathrm{sym}^{\dagger}_{\tilde{v}} \delta\hat{v}\,, \end{aligned} \tag{15}$$

where $\mathrm{sym}^{\dagger}_{v} \delta\hat{v} = -\mathrm{ad}_v \delta\hat{v} + \mathrm{ad}^{\dagger}_{\delta\hat{v}} v$. For more details on the derivation of the adjoint Jacobi field equations, see [7].

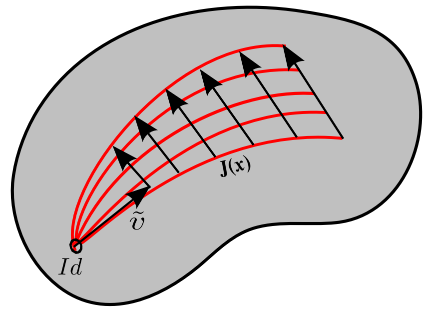

Fig. 2: Jacobi fields

5 Results

We demonstrate the effectiveness of our proposed model using real 3D OASIS MRI brain data. The MRI have resolution $128 \times 128 \times 128$ and are skull-stripped, intensity normalized, and co-registered with rigid transforms. We use $\alpha = 3.0, s = 3.0, \sigma = 0.03$ with $T = 10$ time-steps in geodesic shooting for all the experiments in this paper.

Image registration. We tested our algorithm for geodesic shooting for pairwise image registration at different levels of truncated dimension $N = 4, 8, 16, 32, 64$. We compared the total energy formulated in (5), time consumption, and memory requirement of our model versus the open source implementation of vector momenta LDDMM [11] (`https://bitbucket.org/scicompanat/vectormomentum`). For comparison, we use the same integration method and (α, s, σ, T) parameters for both models.

Figure 3 displays the comparison of convergence, time, and memory at different levels of truncated dimensions. It indicates that our model gains better registration accuracy but with much less time and memory. We see that our method actually achieves a lower overall energy than vector momenta LDDMM for truncated dimension $N = 16$ and higher. Note that increasing the dimension beyond $N = 16$ does not improve the image registration energy, indicating that $N = 16$ is sufficient to capture the transformations between images. We emphasize that we used the same full-dimensional registration energy from (2) for all runs so that they would be comparable. In addition, our model arrives at the optimal solution for $N = 16$ in 1.59s per iteration, and 168.4 MB memory. In comparison, vector momenta LDDMM requires 46s per iteration, and 1708.1 MB memory.

Atlas building We also used our algorithm to build an atlas from a set of 3D brain MRIs from the OASIS database, consisting of 60 healthy subjects between the age of 60 to 95. We initialize the template I as the average of image intensities, and set the truncated dimension as $N = 16$ that was shown to be optimal

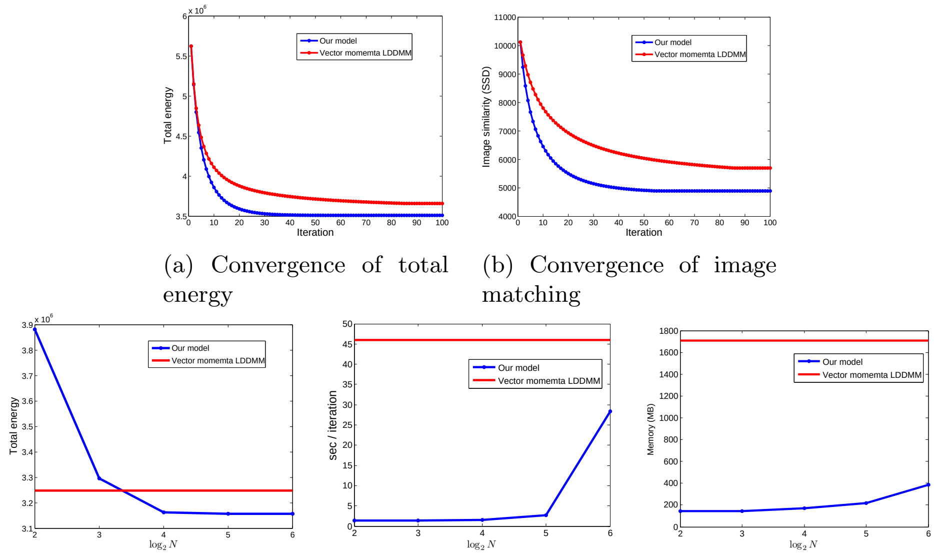

(a) Convergence of total energy

(b) Convergence of image matching

(c) Total energy

(d) Time per iteration

(e) Memory requirement

Fig. 3: Comparison between our model at different scale of truncated dimension and vector momenta LDDMM for (a) convergence of total energy with $N = 16$ truncated dimension; (b) convergence of image matching with $N = 16$ truncated dimension; (c) total energy; (d) time consumption per iteration; (e) memory requirement.

in the previous section. We used a message passing interface (MPI) parallel programming implementation for both our model and vector momenta LDDMM, and scattered each image onto an individual processor. With 100 iterations for gradient descent, our model builds the atlas in 7.5 minutes, while the vector momenta LDDMM in [11] requires 2 hours.

The left side of Figure 4 shows the axial and coronal slices from 8 of the selected 3D MRI dataset. The right side shows the atlas image estimated by our algorithm, followed by the atlas estimated by vector momenta LDDMM. We see from the difference image between the two atlas results that our algorithm generated a very similar atlas to vector momenta LDDMM, but at a fraction of the time and memory cost.

6 Conclusion

We presented a fast geodesic shooting algorithm for diffeomorphic image registration. Our method is the first to introduce a finite-dimensional Lie algebra that can represent discretized velocity fields of diffeomorphisms in a lower-dimensional space. Another key contribution of this finite-dimensional Lie algebra is that we can compute the geodesic evolution equations, as well as the

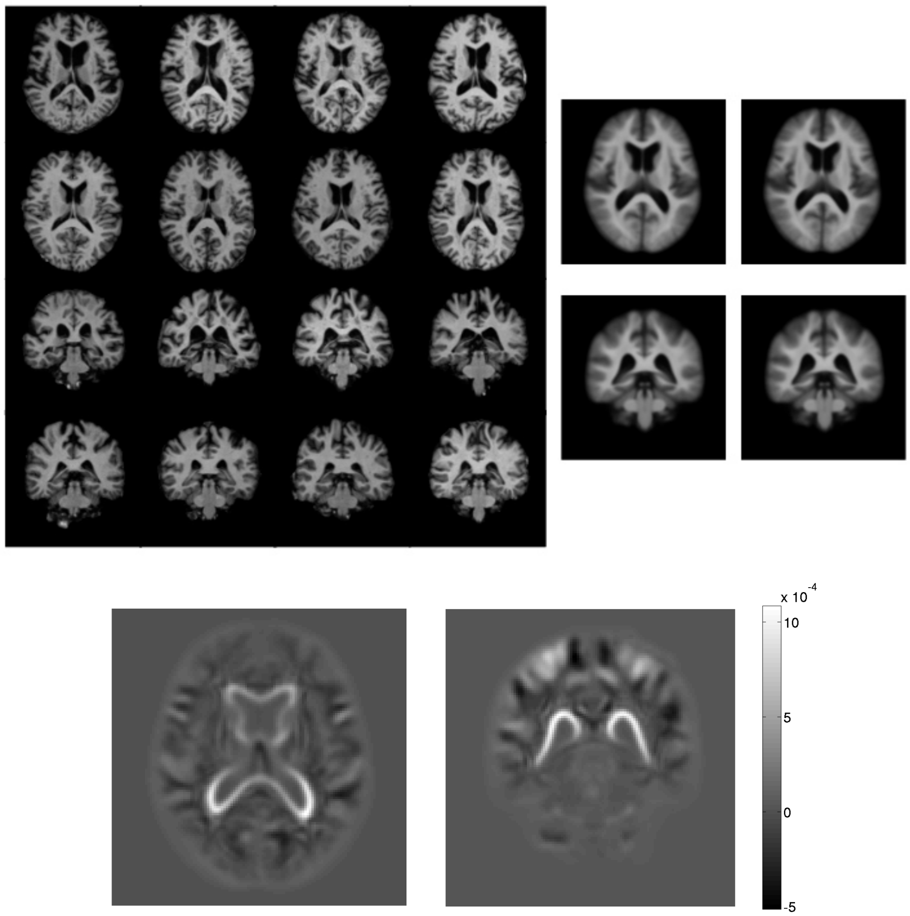

Fig. 4: Top left: axial and coronal slices from 8 of the input 3D MRIs. Middle to right: atlas estimated by our model and vector momenta LDDMM. Bottom: axial and coronal view of atlas intensity difference.

adjoint Jacobi field equations required for gradient descent methods entirely in a low-dimensional vector space. This leads to a dramatically fast diffeomorphic image registration algorithm without loss in accuracy. This work paves the way for efficient computations in large statistical studies using LDDMM. Other speed up strategies, for instance, a second-order Gauss-Newton step, similar to the one proposed in [4], or a multi-resolution optimization scheme could be added on top of our algorithm for further speed improvement. Another interesting possibility is that our algorithm could make inference by Monte Carlo sampling of diffeomorphisms [16] more feasible.

Acknowledgments This work was supported by NIH Grant 5R01EB007688 and NSF CAREER Grant 1054057.

References

1. Arnol'd, V.I.: Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. Ann. Inst. Fourier 16, 319–361 (1966)
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 2006, pp. 924–931. Springer (2006)
3. Ashburner, J.: A fast diffeomorphic image registration algorithm. Neuroimage 38(1), 95–113 (2007)
4. Ashburner, J., Friston, K.J.: Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. NeuroImage 55(3), 954–967 (2011)
5. Beg, M., Miller, M., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision 61(2), 139–157 (2005)
6. Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. IEEE Transactions on Image Processing 5(10), 1435–1447 (1996)
7. Francesco, B.: Invariant affine connections and controllability on lie groups. Tech. rep., technical Report for Geometric Mechanics, California Institute of Technology (1995)
8. Hinkle, J., Fletcher, P.T., Joshi, S.: Intrinsic polynomials for regression on riemannian manifolds. Journal of Mathematical Imaging and Vision (2013)
9. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. Journal of Mathematical Imaging and Vision 24(2), 209–228 (2006)
10. Singh, N., Fletcher, P.T., Preston, J.S., Ha, L., King, R., Marron, J.S., Wiener, M., Joshi, S.: Multivariate statistical analysis of deformation momenta relating anatomical shape to neuropsychological measures. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2010, pp. 529–537. Springer (2010)
11. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.T.: A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: International Symposium on Biomedial Imaging (ISBI) (April 2013)
12. Vaillant, M., Miller, M.I., Younes, L., Trouvé, A.: Statistics on diffeomorphisms via tangent space representations. NeuroImage 23, S161–S169 (2004)
13. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. NeuroImage 45(1), S61–S72 (2009)
14. Vialard, F.X., Risser, L., Rueckert, D., J.Cotter, C.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. In: International Journal of Computer Vision. pp. 229–241 (2012)
15. Younes, L., Arrate, F., Miller, M.: Evolutions equations in computational anatomy. NeuroImage 45(1S1), 40–50 (2009)
16. Zhang, M., Singh, N., Fletcher, P.T.: Bayesian estimation of regularization and atlas building in diffeomorphic image registration. In: Information Processing in Medical Imaging. pp. 37–48. Springer (2013)