

Efficient Parallel Transport in the Group of Diffeomorphisms via Reduction to the Lie Algebra

Kristen M. Campbell and P. Thomas Fletcher

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

Abstract. This paper presents an efficient, numerically stable algorithm for parallel transport of tangent vectors in the group of diffeomorphisms. Previous approaches to parallel transport in large deformation diffeomorphic metric mapping (LDDMM) of images represent a momenta field, the dual of a tangent vector to the diffeomorphism group, as a scalar field times the image gradient. This “scalar momenta” constraint couples tangent vectors with the images being deformed and leads to computationally costly horizontal lifts in parallel transport. This paper uses the vector momenta formulation of LDDMM, which decouples the diffeomorphisms from the structures being transformed, e.g., images, point sets, etc. This decoupling leads to parallel transport expressed as a linear ODE in the Lie algebra. Solving this ODE directly is numerically stable and significantly faster than other LDDMM parallel transport methods. Results on 2D synthetic data and 3D brain MRI demonstrate that our algorithm is fast and conserves the inner products of the transported tangent vectors.

1 Introduction

Analysis of anatomical shape changes from longitudinal medical imaging requires comparing the changes over time of subjects in disparate groups. For instance, imaging studies have shown that the hippocampi of subjects with Alzheimer’s disease atrophy significantly more over time than those of healthy aging subjects. Trajectories of anatomical shape change can be estimated from sequences of images using regression methods in the space of diffeomorphisms. When trajectories are modeled as geodesics, they can be represented by their initial velocity. However, these velocities are defined with respect to different coordinate systems associated with the baseline image of each subject. As such, they are not directly comparable. In order to perform statistical analysis of these trajectories, multiple researchers have proposed using parallel transport to bring these subject-specific trajectories into a common coordinate system for comparison.

One of the preferred techniques for analyzing images in this context is large deformation diffeomorphic metric mapping (LDDMM) [2], which is a mathematical framework for finding smooth diffeomorphic transformations between images. The main benefits of LDDMM are that the deformations between images are smooth and invertible and that a metric allows distances between diffeomorphisms to be computed in a meaningful way. Existing methods for parallel

translation within the LDDMM setting work well in practice and have been used for longitudinal shape analysis [11, 12]. However, these approaches require computing horizontal lifts at each time step, an expensive computation involving solving a linear system for the scalar momenta using an iterative conjugate gradient method [16]. Additionally, these methods approximate each time step of parallel transport with a short time step evolution of Jacobi fields. Instead of that approximation, this paper works directly with the parallel translation equation.

Another method [4, 5] uses a sparse parameterization of the diffeomorphism by using control points. Other approaches involve using stationary velocity fields (SVFs) to generate the diffeomorphisms. These methods use Schild’s ladder to approximate parallel transport along a curve by taking small steps in the associated tangent space at each time [6]. However, as each rung of Schild’s ladder requires two imperfect, computationally-intensive image registrations, the Schild’s ladder steps are then further approximated using the Baker-Campbell-Hausdorff (BCH) formula. This results in fast parallel transport, however, much like the Jacobi field approximation for LDDMM [16], each time step is an approximation to the direct parallel transport equation. Also, the SVF formulation is different from LDDMM in that it does not result in a distance metric on the space of diffeomorphisms.

This paper uses the vector momenta [13] formulation of LDDMM to decouple the diffeomorphisms from the image data. This decoupling allows us to work in the full Lie algebra of the space of diffeomorphisms, and we can then directly implement the parallel translation equations in terms of right-invariant tangent vectors to this space of diffeomorphisms. This results in a linear ordinary differential equation (ODE) that can be solved with a standard, numerically stable scheme that avoids the need to perform computationally-expensive horizontal lifts to the constraint of scalar momenta, as is done in [16]. Additionally, we use the Fourier approximations of vector fields from [17] to gain more numeric stability as well as a more efficient algorithm.

We perform experiments with 2D synthetic data and 3D brain MRIs in order to show the effectiveness of our approach and demonstrate that we can transport realistic vector fields even for quite large deformations. Our results show that our approach is quite fast, indeed it is two orders of magnitude faster than the LDDMM image matching that is also performed using the efficient Fourier-approximated vector fields. Additionally, we demonstrate conservation of the inner product of the tangent vectors being transported. This approaches nearly exact conservation as we increase the number of time steps in the numerical integration scheme.

2 Background on Diffeomorphisms and LDDMM

We provide a brief review of diffeomorphisms, associated Lie group operators and the LDDMM formulation, highlighting the math relevant to parallel transport of diffeomorphisms and the links to diffeomorphic image registration.

2.1 Diffeomorphisms

Let $\Omega = \mathbb{R}^d/\mathbb{Z}^d$ be a d -dimensional toroidal image domain. A toroidal domain is the natural setting for defining the Fourier transform and assuming cyclical boundary conditions. A diffeomorphism of Ω is a bijective, C^∞ mapping $\phi : \Omega \rightarrow \Omega$ whose inverse, ϕ^{-1} , is also C^∞ . We will denote the space of all such diffeomorphisms as $\text{Diff}(\Omega)$. We are particularly interested in time-varying diffeomorphisms, $\phi(t, x) : [0, 1] \times \Omega \rightarrow \Omega$, which can be generated as flows of time-varying velocity fields $v(t, x) : [0, 1] \times \Omega \rightarrow \mathbb{R}^d$. These will be referred to in this paper as $\phi_t(x)$ and $v_t(x)$, where $t \in [0, 1]$ and $x \in \Omega$. Note that $\phi_t(x)$ is generated by the flow $t \mapsto \phi_t \in \text{Diff}(\Omega)$ by integrating the ODE

$$\frac{d\phi_t}{dt} = v_t \circ \phi_t. \quad (1)$$

We know that $\text{Diff}(\Omega)$ is an infinite-dimensional Lie group, whose associated Lie algebra, $V = \mathfrak{X}(\Omega)$, consists of all C^∞ vector fields on Ω . For two vector fields $v, w \in V$, the Lie bracket is defined as $[v, w] = Dv \cdot w - Dw \cdot v$. Here D is the first derivative operator and \cdot is element-wise matrix-vector multiplication.

In order to define distances on the manifold $\text{Diff}(\Omega)$, we need an appropriate Riemannian metric. Here we use a weak metric

$$\langle v, w \rangle_V = \int_{\Omega} \langle Lv(x), w(x) \rangle dx, \quad (2)$$

where $L : V \rightarrow V$ is a positive-definite, self-adjoint differential operator. In this paper, L is chosen to be a Laplacian operator of the form $L = (-\alpha\Delta + I)^c$ where $\alpha > 0$, $c > 0$ and I is the $d \times d$ identity matrix. In order to compute the inner product of vector fields v and w that belong to the tangent space of any other element $\phi \in \text{Diff}(\Omega)$, we need to pull back the velocities to the tangent space at identity by using a right-invariant metric such as

$$\langle v, w \rangle_{T_\phi \text{Diff}(\Omega)} = \langle v \circ \phi^{-1}, w \circ \phi^{-1} \rangle_V. \quad (3)$$

Then the distance between ϕ and id becomes

$$\text{dist}(\text{id}, \phi) = \int_0^1 \|v_t\|_V dt. \quad (4)$$

2.2 LDDMM Image Registration

For the image registration application, we will be looking at how to find an optimal diffeomorphism that takes us from image I_0 to image I_1 , where optimal will mean that the diffeomorphism is as small as possible and that $I_0 \circ \phi_1^{-1}$ is as close to I_1 as possible. The LDDMM formulation will formulate this problem as an energy-minimization problem. Before we get to the more specific notion of diffeomorphisms acting on images, we look at the geodesic equations in the general Lie group setting.

First, we will need to define some fundamental operators from Lie group theory. We define the adjoint action of $\text{Diff}(\Omega)$ on $\mathfrak{X}(\Omega)$, $\text{Ad}_\psi : V \rightarrow V$, as

$$\text{Ad}_\psi(v) = \frac{d}{dt} (\psi \circ \phi_t \circ \psi^{-1})|_{t=0}, \quad (5)$$

where $\phi_0 = \text{id}$ and $\frac{d\phi}{dt}|_{t=0} = v$. Note that if ϕ_t and ψ commuted, we would simply end up with ϕ_t , thus Ad_ψ is evaluating how well all infinitesimal deformations commute with ψ . Now we let ψ be time-varying and we define the adjoint action, ad , of $\mathfrak{X}(\Omega)$ on itself by

$$\text{ad}_u v = \frac{d}{ds} (\text{Ad}_{\psi_s} v)|_{s=0}, \quad (6)$$

where $\psi_0 = \text{id}$ and $\frac{d\psi}{ds}|_{s=0} = u$. In the case of the Lie group $\text{Diff}(\Omega)$, the adjoint action is given by the formula

$$\text{ad}_u v = [u, v] = Du \cdot v - Dv \cdot u. \quad (7)$$

For the energy optimization, we will use results from Arnold [1] and Miller et al. [9] that show that geodesics are extremal curves that satisfy the Euler-Poincaré equations for diffeomorphisms (EPDiff):

$$\frac{dv_t}{dt} = -\text{ad}_{v_t}^\dagger w_t, \quad (8)$$

where ad^\dagger , the adjoint of the ad operator, is

$$\text{ad}_{v_t}^\dagger w_t = K [(Dv_t)^T Lw_t + D(Lw_t)v_t + Lw_t \text{div } v_t]. \quad (9)$$

Here div denotes the divergence operator. The process of finding the unique geodesic path, ϕ_t by integrating an initial velocity, $v_0 \in V$ at $t = 0$ forward in time according to (8) is known as geodesic shooting.

Now let's look more specifically at diffeomorphisms acting on images $I \in L^2(\Omega, \mathbb{R})$, meaning that images are square-integrable functions defined on Ω . Diffeomorphic image registration is looking for a v_t that minimizes an energy function, $E(v_t)$, that measures how well $I_0 \circ \phi_1^{-1}$ matches I_1 while preferring small diffeomorphisms by adding a regularization term.

$$E(v_t) = \frac{1}{2\sigma^2} \|I_0 \circ \phi_1^{-1} - I_1\|_{L^2}^2 + \int_0^1 \|v_t\|_V^2 dt, \quad (10)$$

where σ^2 represents image noise variance.

Vialard et al. [14] and Younes et al. [15] showed that it is only necessary to estimate the initial velocity, v_0 . Therefore, we can rewrite (10) as

$$E(v_0) = \frac{1}{2\sigma^2} \|I_0 \circ \phi_1^{-1} - I_1\|_{L^2}^2 + \|v_0\|_V^2, \quad \text{s.t. EPDiff (8) holds.} \quad (11)$$

2.3 Decoupling Diffeomorphisms from Images

In the original LDDMM formulation, Beg et al. [2] showed that the initial vector fields that are minimizers of the diffeomorphic image registration energy (11) are of the form $\hat{v}_0 = K(s\nabla I_0)$, where $s : \Omega \rightarrow \mathbb{R}$ is a scalar field. In other words, the initial momenta $m_0 = Lv_0 = s\nabla I_0$ is constrained to be a scalar field times the image gradient. The scalar momenta constraint was also used in the derivation of geodesic shooting by Vialard [14]. This constraint has the practical benefit that it reduces the size needed to represent the initial conditions, i.e., we can discretize the scalar field s , rather than the vector field v_0 . However, Singh et al. [13] showed that removing the scalar momenta constraint, that is, optimizing over initial momenta m_0 that are vector fields, was more numerically stable and converged to better local optima of the target energy.

Removing the scalar momenta constraint also has the effect of *decoupling* the diffeomorphisms from the images that they are acting on. This decoupling has advantages in Bayesian formulations of diffeomorphic image registration and atlas building, as developed by Zhang et al. [19]. In this approach, the decoupling enables formulation of diffeomorphisms as latent random variable with a prior that does not depend on the images (data) in any way. Furthermore, elements of the Lie algebra V are spatially smooth vector fields, and as such, are easier to deal with numerically than non-smooth momenta fields. Zhang and Fletcher [17] used this fact to show that initial velocities could be efficiently represented in the Fourier domain by low-frequency approximations, resulting in much faster image registration and even better optimization of the LDDMM energy. Similarly, we show in the next section that parallel translation benefits from this same decoupling of the diffeomorphisms from images. By working in the full Lie algebra of the space of diffeomorphisms, we can directly implement the equations for parallel translation in terms of right-invariant tangent vectors to $\text{Diff}(\Omega)$. As such, we avoid the need to perform computationally-expensive horizontal lifts to the constraint of scalar momenta, as used in [16]. We are also able to use the Fourier approximations of vector fields from [17]. The end result is an efficient and numerically stable algorithm for directly computing parallel transport in the space of diffeomorphisms.

3 Parallel Transport

In order to do comparisons of trajectories defined by geodesic segments in the space of diffeomorphisms, we need a way to bring the initial velocities of these geodesics to the same reference point. One mechanism to do so is called parallel transport, a generalization of the Euclidean notion of parallel translation of one vector to the origin of another. When this happens in Euclidean space, the angle between the vectors is preserved and the magnitude of the vector is preserved. We will see below that parallel transport along a geodesic on a Riemannian manifold similarly preserves the inner product of the transported vector to the tangent vector of the geodesic and also preserves the norms of the transported vector and the tangent vector.

In this section, we will start from the definition of parallel transport on general Lie groups and then look more specifically at parallel transport on the manifold of diffeomorphisms, $\text{Diff}(\Omega)$. Then we will do a computational complexity analysis of parallel transport. We'll also talk about details related to implementing parallel transport of diffeomorphisms on a computer, including numerical integration details and using a Fourier-approximated Lie algebra to speed up the discrete computation.

3.1 Parallel Transport Equation

Let's start by looking generally at right-invariant vector fields v and w on a Lie group. We can look at how w varies in the direction v by looking at the covariant derivative $\nabla_v w$. Parallel transport of a tangent vector along a curve is defined by this covariant derivative of the transported vector being zero in the direction of the velocity of the curve. The covariant derivative for right-invariant vector fields (c.f. [3]) is given by the equation

$$\nabla_v w = -\frac{1}{2} \left(\text{ad}_v^\dagger w + \text{ad}_w^\dagger v - \text{ad}_v w \right). \quad (12)$$

For our application, we want to transport along a curve ϕ_t in the space of diffeomorphisms. Remember that (1) says the change in ϕ_t over time is equal to a time-varying velocity field v_t composed with ϕ_t . Let's plug these time-varying right-invariant vector fields into (12) to get the following:

$$\nabla_{v_t} w_t = \frac{dw_t}{dt} - \frac{1}{2} \left(\text{ad}_{v_t}^\dagger w_t + \text{ad}_{w_t}^\dagger v_t - \text{ad}_{v_t} w_t \right), \quad (13)$$

where the dw_t/dt comes from needing to take the total derivative since w_t varies with time. Let's set this covariant derivative to 0 and substitute the definitions for $\text{ad}_{v_t}^\dagger w_t$ and $\text{ad}_{w_t}^\dagger v_t$ from (7) and (9) to get:

$$\begin{aligned} \frac{dw_t}{dt} = & -\frac{1}{2} (K [(Dv_t)^T Lw_t + D(Lw_t)v_t + Lw_t \text{div } v_t] \\ & + K [(Dw_t)^T Lv_t + D(Lv_t)w_t + Lv_t \text{div } w_t] \\ & - Dv_t w_t + Dw_t v_t). \end{aligned} \quad (14)$$

Notice that this becomes the geodesic equation when $w_t = v_t$.

3.2 Computational Complexity Analysis

The computational complexity of solving (14) is $O(NM \log M)$ where M is the number of voxels in the image and N is the number of time steps taken. If we instead solve (14) in the Fourier-approximated space, the complexity improves to $O(Nm \log m)$, where m is number of frequencies used in the reduced space. We use $m = 16^3$ in the real data experiments below. Note that computing Jacobi fields in the Lie algebra would be the same complexity. But our method avoids

computing the horizontal lifts needed in order to enforce the scalar momenta constraint used by other LDDMM parallel transport methods. These horizontal lifts involve solving an M -dimensional system of linear equations using an iterative conjugate gradient method.

3.3 Implementation Details

In order to gain the performance benefits from performing operations in the Fourier-approximated Lie algebra (FLASH) [17], we implemented parallel translation in the Flash C++ environment [18], which is built on top of PyCA [10]. Additionally, we implemented a more accurate numerical integration scheme for both forward integrating v_0 at each time step $t = 1/N, 2/N, \dots, 1$ and for solving the parallel translation ODE (13) numerically using N time steps. We provide both an Euler first-order scheme and a Runge-Kutta fourth-order (RK4) scheme to perform these integrations, with the RK4 integration happening independently for w_t and v_t . Ideally, since w_t and v_t are coupled, the integration would be even more accurate by doing a coupled symplectic integration scheme. Below we run experiments in order to find a reasonably small N that gives good enough stability.

4 Experiments

We ran experiments with both synthetic images and real 3D MR images of human brains to explore the accuracy, stability and speed of our parallel translation approach. In all of our experiments, we follow the same general setup. First we use Flash C++ to do image matching between two time points of the first subject, I_0 and I_n , to find a diffeomorphism between the two images represented by an initial velocity w_{0_n} . Then we do image matching between the first subject, I_0 and the template image, T_0 to find a diffeomorphism between the subject and the template represented by the initial velocity v_0 . At this point, we parallel translate w_{0_n} in the direction of v_0 to get the translated diffeomorphism represented by the initial velocity $\pi(w_{0_n})$ as shown in Figure 1. The values $\langle Lv, v \rangle$, $\langle Lv, w \rangle$, and $\langle Lw, w \rangle$, while different from each other, should each remain constant throughout integration. We measure the percent change of these inner products at each time step of the integration in order to quantify the stability.

4.1 Synthetic Data

We modeled our synthetic data experiments to follow the approach of Lorenzi and Pennec [7] that was also used by [8]. These experiments consist of 2D images of size 256 x 256 pixels. The subject’s initial image, I_0 , at time 0 is composed of centered black and white semi-circles with a 21 pixel radius surrounded by a centered grey circle with a 42 pixel radius. Brain atrophy over 3 time steps is modeled by decreasing the volume of the outer grey circle by 5% of the initial time point at each time increment, while simultaneously increasing the volume

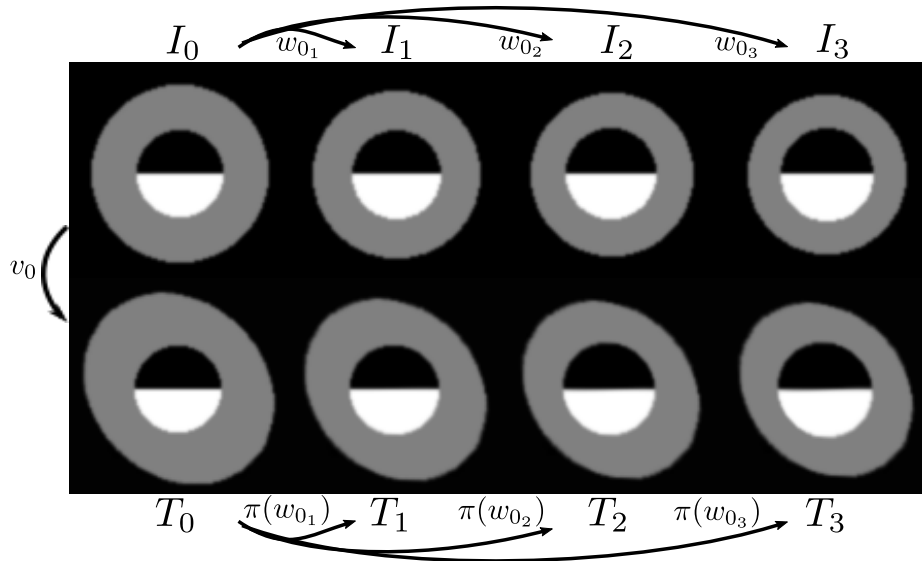


Fig. 1: Results of parallel translating w_{0_n} for each $I_0 \dots I_3$ to the template space T_0 to produce the transformed images T_1, T_2, T_3 .

of the inner semi-circles by 5%. A second image, T_0 , which can represent either a template image, atlas image, or second subject, is composed of the same black and white semi-circles found in I_0 , while the outer grey circle has been deformed into an ellipse by stretching the top and bottom edges an amount equal to 10% of the diameter of the grey circle in I_0 and then rotating the ellipse by 45° .

The initial velocities $v_0, w_{0_1}, w_{0_2}, w_{0_3}$ of the deformations are found with 100 iterations of Flash C++'s image matching using a truncation dimension of 16 and the parameters $\alpha = 3.0, s = 3.0, \sigma = 0.03, \gamma = 0.2$. The results of transporting these velocities between I_0 and each time point $I_n, n \in \{1, 2, 3\}$ can be seen in Figure 1.

Numerical Stability We expect to see that the stability of parallel translation improves as the number of time steps of numerical integration increase for a particular integration scheme, or improves for higher order numerical schemes. Additionally, we want to characterize the stability of translation as the deformations grow larger. Therefore, we compare the percent change relative to the value at $t = 0$ of the norms, $\langle Lv, v \rangle = \|v\|^2, \langle Lw, w \rangle = \|w\|^2$, and the relative percent change of the inner product, $\langle Lv, w \rangle$, all of which should be 0 since these inner products remain constant throughout parallel translation. We look at how the stability changes as we do either 10, 20 or 100 time steps of both an Euler first-order numerical integration scheme and a Runge-Kutta 4 (RK4) fourth-order integration scheme. We then do this same comparison for the 5%, 10%, and 15% volume change to see how the implementation behaves as defor-




Fig. 2: Percent relative change of the inner product over the integration time from 0 to 1. Where relative percent change of $x = \langle \cdot, \cdot \rangle$ is computed by $100*(x_t - x_0)/x_0$ for each time step t . The rows correspond to results for 5%, 10% and 15% simulated atrophy respectively.

mations grow larger. As you can see in Figure 2, the relative change approaches 0 as expected as the number of integration steps increases and as the numerical scheme changes from Euler to RK4 where it becomes effectively 0 for 100 steps of RK4. While it is more expensive to compute this many steps of RK4, it provides excellent preservation of the relationship between v and w throughout the parallel transport. From these results, we see that 20 iterations of RK4 gives reasonable performance without being too computationally intensive. Therefore, we chose to do 20 iterations of RK4 for the real data experiments below. Note also that the percent relative change for the angle between v and w , $\langle Lv, w \rangle$, increases somewhat as the amount of deformation increases from 5% simulated atrophy in the top row to 10% atrophy in the center row to the 15% simulated atrophy in the bottom row. It is expected that the larger deformation leads to somewhat larger errors.

4.2 Real Data

We looked at performance of our parallel translation in the context of 3D brain MRIs from the OASIS database in order to see how well it captures known atrophy associated with the progression of Alzheimer’s disease. We did pairwise comparison of every combination of 11 healthy subjects and 10 subjects with Alzheimers. We start with images of size 128 x 128 x 128 that have had the

skulls stripped out, intensities normalized, and are then rigidly co-registered. For each pair, the subject with Alzheimer’s is I_0 , and I_1 is the same subject’s scan at a later time, between 2 and 5 years later. the healthy subject’s initial scan is T_0 . I_0 is deformably registered to each of I_1 and T_0 using 200 iterations of image matching from Flash C++ with truncated dimension of 16 and parameters $\alpha = 3.0, s = 3.0, \sigma = 0.03, \gamma = 1.0$. The results for a typical pair are shown in Figure 3.




Fig. 3: Axial and coronal views of the results of parallel translating w_0 along v_0 to the template space T_0 to produce the transformed image $\pi(w_0)$. The top row in each group consists of the original images from a subject with Alzheimer’s, I_0 and I_1 as well as the difference image $I_1 - I_0$. The bottom row in each group consists of the original template image from a control subject T_0 , the parallel translated image $\pi(w_0)$, and the difference image $\pi(w_0) - T_0$.

One of the benefits of this parallel transport is that it is significantly faster than the image matching registrations used to produce the diffeomorphisms to be transported. For our experiments with real data, the image registration of one pair of 3D images took on the order of 800 seconds while the parallel transport of 3D vector fields from one subject to another took on the order of 8 seconds. Since parallel transport is consistently 2 orders of magnitude faster than the image matching, it becomes an essentially free operation for an image analysis pipeline.

Numerical Stability In order to evaluate the numerical stability, we look at the maximum percent change of the inner products $\langle Lv, v \rangle, \langle Lv, w \rangle, \langle Lw, w \rangle$ as w is transported along v . A plot of the maximum percent change across all 106 pairs of subjects can be found in Figure 4. The largest maximum relative change across all subjects in $\langle Lv, w \rangle$ is 8.6%. The associated $\langle Lv, v \rangle$ is 0.0009% and the $\langle Lw, w \rangle$ is 0.0000% for this same pair of subjects. In order to understand why this pair had such a large value compared to other pairs, we ran the same pair for 100 steps of RK4. That experiment resulted in much smaller percent change of $\langle Lv, w \rangle = 1.51\%, \langle Lv, v \rangle = 0.00086\%, \langle Lw, w \rangle = 0.0000\%$.




Fig. 4: Maximum amount of change in inner products for each subject. The left plot shows the maximum relative change of each subject $x = \langle \cdot, \cdot \rangle$ is computed by $\max_{0 \leq t \leq 1} ((x_t - x_0)/x_0)$. The right plot shows the maximum absolute change of each subject, computed as $\max_{0 \leq t \leq 1} (x_t - x_0)$.

5 Conclusion

We presented a method to perform parallel translation in the space of diffeomorphisms, allowing us to work with the parallel translation equations directly instead of approximating them. Further, we were able to use FLASH to speed up computations further by performing them in a smaller Fourier-approximated space. We demonstrated that our method is numerically stable and that preservation of the inner products throughout parallel translation can be improved in a predictable manner by increasing the number of integration steps and/or using the RK4 scheme with a modest associated computational cost.

We look forward to applying this method to studying trajectories of anatomical shape change in a variety of medical image analysis contexts. Also, this method is one example of how formulating problems directly in the space of diffeomorphisms and working with that Lie algebra combined with efficiencies gained from FLASH techniques allows us to perform computations efficiently and stably. We expect that this same approach could work well for other analysis such as working with Sasaki metrics.

Acknowledgements This work was supported by NIH grant R01EB022876. The OASIS data was provided by the following grants: P50 AG05681, P01 AG03991, R01 AG021910, P20 MH071616, U24 RR021382.

References

1. Arnold, V.: Sur la géométrie différentielle des groupes de lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits. In: Annales de l'institut Fourier. vol. 16, pp. 319–361 (1966)

2. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* 61(2), 139–157 (2005)
3. Cheeger, J., Ebin, D.G., Ebin, D.G.: *Comparison theorems in Riemannian geometry*, vol. 9. North-Holland Publishing Company Amsterdam (1975)
4. Durrleman, S., Prastawa, M., Gerig, G., Joshi, S.: Optimal data-driven sparse parameterization of diffeomorphisms for population analysis. In: *Information Processing in Medical Imaging*. pp. 123–134. Springer (2011)
5. Durrleman, S., Prastawa, M., Charon, N., Korenberg, J.R., Joshi, S., Gerig, G., Trouvé, A.: Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage* 101, 35–49 (2014)
6. Lorenzi, M., Ayache, N., Pennec, X.: Schilders ladder for the parallel transport of deformations in time series of images. In: *Information Processing in Medical Imaging*. pp. 463–474. Springer (2011)
7. Lorenzi, M., Pennec, X.: Geodesics, parallel transport & one-parameter subgroups for diffeomorphic image registration. *International journal of computer vision* 105(2), 111–127 (2013)
8. Matsui, J.T.: *Development of image processing tools and procedures for analyzing multi-site longitudinal diffusion-weighted imaging studies* (2014)
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. Preston, J.S.: *Python for computational anatomy* (2016), <https://bitbucket.org/scicompanat/pyca>, commit e649151
11. Qiu, A., Albert, M., Younes, L., Miller, M.I.: Time sequence diffeomorphic metric mapping and parallel transport track time-dependent shape changes. *NeuroImage* 45(1), S51–S60 (2009)
12. Qiu, A., Younes, L., Miller, M.I., Csernansky, J.G.: Parallel transport in diffeomorphisms distinguishes the time-dependent pattern of hippocampal surface deformation due to healthy aging and the dementia of the alzheimer’s type. *NeuroImage* 40(1), 68–76 (2008)
13. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.T.: A vector momenta formulation of diffeomorphisms for improved geodesic regression and atlas construction. In: *Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on*. pp. 1219–1222. IEEE (2013)
14. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3d image registration via geodesic shooting using an efficient adjoint calculation. *International Journal of Computer Vision* 97(2), 229–241 (2012)
15. Younes, L., Arrate, F., Miller, M.I.: Evolutions equations in computational anatomy. *NeuroImage* 45(1), S40–S50 (2009)
16. Younes, L., Qiu, A., Winslow, R.L., Miller, M.I.: Transport of relational structures in groups of diffeomorphisms. *Journal of mathematical imaging and vision* 32(1), 41–56 (2008)
17. Zhang, M., Fletcher, P.T.: Finite-dimensional lie algebras for fast diffeomorphic image registration. In: *International Conference on Information Processing in Medical Imaging*. pp. 249–260. Springer (2015)
18. Zhang, M., Fletcher, P.T.: *Flashc++* (2016), <https://bitbucket.org/FlashC/flashc>, commit 33cfd0e
19. Zhang, M., Singh, N., Fletcher, P.T.: Bayesian estimation of regularization and atlas building in diffeomorphic image registration. In: *IPMI*. vol. 23, pp. 37–48 (2013)