

Frequency Diffeomorphisms for Efficient Image Registration

Miaomiao Zhang¹, Ruizhi Liao¹, Adrian V. Dalca¹, Esra A. Turk²,
Jie Luo², P. Ellen Grant², and Polina Golland¹

¹ Computer Science and Artificial Intelligence Laboratory, MIT

² Boston Children's Hospital, Harvard Medical School

Abstract. This paper presents an efficient algorithm for large deformation diffeomorphic metric mapping (LDDMM) with geodesic shooting for image registration. We introduce a novel finite dimensional Fourier representation of diffeomorphic deformations based on the key fact that the high frequency components of a diffeomorphism remain stationary throughout the integration process when computing the deformation associated with smooth velocity fields. We show that manipulating high dimensional diffeomorphisms can be carried out entirely in the bandlimited space by integrating the nonstationary low frequency components of the displacement field. This insight substantially reduces the computational cost of the registration problem. Experimental results show that our method is significantly faster than the state-of-the-art diffeomorphic image registration methods while producing equally accurate alignment. We demonstrate our algorithm in two different applications of image registration: neuroimaging and in-utero imaging.

1 Introduction

Diffeomorphisms have been widely used in the field of image registration [7, 6], atlas-based image segmentation [4, 10], and anatomical shape analysis [14, 22]. In this paper, we focus on a time-varying velocity field representation for diffeomorphisms as it supplies a distance metric that is critical to statistical analysis of anatomical shapes, for instance, by least squares, geodesic regression, or principal modes detection [13, 16, 22].

In spite of the advantages of supporting Riemannian metrics in LDDMM, the extremely high computational cost and large memory footprint of the current implementations has limited the use of time-varying velocity representations in important applications that require computational efficiency. The original LDDMM optimization performs gradient descent on the entire time-varying velocity field that is defined on a dense image grid. Since a geodesic is uniquely determined by its initial conditions on the velocity field, the geodesic shooting algorithm has been shown to reduce the computational complexity and improve optimization landscape by manipulating the initial velocity via the geodesic evolution equations [20]. FLASH (finite dimensional Lie algebras for shooting) [23]

is a recent variant of LDDMM with geodesic shooting that employs a low dimensional bandlimited representation of the initial velocity field to further improve the convergence and efficiency of the optimization. The algorithm still maps the velocity fields from the low dimensional Fourier space back to the full image domain to perform forward integration at each iteration [23]. The computational complexity of this step thus dominates the entire procedure of diffeomorphic image registration.

Previous works that aimed to improve diffeomorphic representations have reduced the high degrees of freedom available to represent the velocity fields, but not the diffeomorphisms themselves. In this paper, we adopt the low dimensional representation of the tangent space of diffeomorphisms [23] and propose an efficient way to compute diffeomorphisms in the bandlimited space, thus further reducing the computational complexity of image registration. Our approach is based on the important insight that only the low frequency components of the diffeomorphisms vary over time when integrating a bandlimited velocity field to obtain the deformation. Since the optimization of image registration can be directly solved by advecting the inverse of diffeomorphisms [17], we propose a novel Fourier representation of the deformation in the inverse coordinate system that is computed entirely in the low dimensional bandlimited space. The theoretical tools developed in this paper are broadly applicable to other parametrization of diffeomorphic transformations, such as stationary velocity fields [2, 3, 19]. To evaluate the proposed algorithm, we perform image registration of real 3D MR images and show that the accuracy of the propagated segmentations is comparable to that obtained via the state-of-the-art diffeomorphic image registration methods, while the runtime and the memory demands are dramatically lower for our method. We demonstrate the method in the context of atlas-based segmentation of brain images and of temporal alignment of in-utero MRI scans.

2 Background

Before introducing our development, we provide a brief overview of continuous diffeomorphisms endowed with metrics on vector fields [6, 21] and of the finite dimensional Fourier representation of the tangent space of diffeomorphisms [23].

Given an open and bounded d -dimensional domain $\Omega \subset \mathbb{R}^d$, we use $\text{Diff}(\Omega)$ to denote a space of continuous differentiable and inverse differentiable mappings of Ω onto itself. The distance metric between the identity element e and any diffeomorphism ϕ

$$\text{dist}(e, \phi) = \int_0^1 (\mathcal{L}v_t, v_t) dt \quad (1)$$

depends on the time-varying Eulerian velocity field v_t ($t \in [0, 1]$) in the tangent space of diffeomorphisms $V = T\text{Diff}(\Omega)$. Here $\mathcal{L} : V \rightarrow V^*$ is a symmetric, positive-definite differential operator, e.g., discrete Laplacian, with its inverse $\mathcal{K} : V^* \rightarrow V$, and $m_t \in V^*$ is a momentum vector that lies in the dual space V^* such that $m_t = \mathcal{L}v_t$ and $v_t = \mathcal{K}m_t$.

The path of deformation fields ϕ_t parametrized by $t \in [0, 1]$ is generated by

$$\frac{d\phi_t}{dt} = v_t \circ \phi_t, \quad (2)$$

where \circ is a composition operator. The inverse mapping of ϕ_t is defined via

$$\frac{d\phi_t^{-1}}{dt} = -D\phi_t^{-1} \cdot v_t, \quad (3)$$

where D is the $d \times d$ Jacobian matrix at each voxel and \cdot is an element-wise multiplication.

The geodesic shooting algorithm estimates the initial velocity at $t = 0$ and relies on the fact that a geodesic path of transformations ϕ_t and its inverse ϕ_t^{-1} with a given initial condition v_0 can be uniquely determined through integrating the Euler-Poincaré differential equation (EPDiff) [1, 12] as

$$\frac{\partial v_t}{\partial t} = -\mathcal{K} [(Dv_t)^T m_t + Dm_t v_t + m_t \operatorname{div}(v_t)], \quad (4)$$

where div is the divergence operator and \mathcal{K} is a smoothing operator that guarantees the smoothness of the velocity fields.

2.1 Fourier Representation of Velocity Fields

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a real-valued function. The Fourier transform \mathcal{F} of f is given by

$$\mathcal{F}[f](\xi_1, \dots, \xi_d) = \int_{\mathbb{R}^d} f(x_1, \dots, x_d) e^{-2\pi i(x_1 \xi_1 + \dots + x_d \xi_d)} dx_1, \dots, dx_d, \quad (5)$$

where (ξ_1, \dots, ξ_d) is a d -dimensional frequency vector. The inverse Fourier transform \mathcal{F}^{-1} of a discretized Fourier signal \tilde{f}

$$\mathcal{F}^{-1}[\tilde{f}](x_1, \dots, x_d) = \sum_{\xi_1, \dots, \xi_d} \tilde{f}(\xi_1, \dots, \xi_d) e^{2\pi i(\xi_1 x_1 + \dots + \xi_d x_d)} \quad (6)$$

is an approximation of the original signal f . To ensure that \tilde{f} represents a real-valued vector field in the spatial domain, we require $\tilde{f}(\xi_1, \dots, \xi_d) = \tilde{f}^*(-\xi_1, \dots, -\xi_d)$, where $*$ denotes the complex conjugate. When working with vector-valued functions of diffeomorphisms ϕ and velocity fields v , we apply the Fourier transform to each vector component separately.

Since \mathcal{K} is a smoothing operator that suppresses high frequencies in the Fourier domain, the geodesic evolution equation (4) suggests that the velocity field v_t can be represented efficiently as a bandlimited signal in Fourier space. Let \tilde{V} denote the discrete Fourier space of velocity fields. As shown in [23], for any two elements $\tilde{v}, \tilde{w} \in \tilde{V}$, the distance metric at identity \tilde{e} is defined as

$$\langle \tilde{u}, \tilde{v} \rangle_{\tilde{V}} = \int (\tilde{\mathcal{L}}\tilde{v}(\xi), \tilde{w}(\xi)) d\xi,$$

where $\tilde{\mathcal{L}} : \tilde{V} \mapsto \tilde{V}^*$ is the Fourier transform of a commonly used Laplacian operator $(-\alpha\Delta + e)^c$, with a positive weight parameter α and a smoothness parameter c , i.e.,

$$\tilde{\mathcal{L}}(\xi_1, \dots, \xi_d) = \left[-2\alpha \sum_{j=1}^d (\cos(2\pi\xi_j) - 1) + 1 \right]^c.$$

The Fourier representation of the inverse operator $\mathcal{K} : \tilde{V}^* \mapsto \tilde{V}$ is equal to $\tilde{\mathcal{K}}(\xi_1, \dots, \xi_d) = \tilde{\mathcal{L}}^{-1}(\xi_1, \dots, \xi_d)$.

Therefore, the geodesic shooting equation (4) can be efficiently computed in a low dimensional bandlimited velocity space:

$$\frac{\partial \tilde{v}_t}{\partial t} = -\tilde{\mathcal{K}} \left[(\tilde{D}\tilde{v}_t)^T \star \tilde{m}_t + \tilde{\nabla} \cdot (\tilde{m}_t \otimes \tilde{v}_t) \right], \quad (7)$$

where \star is the truncated matrix-vector field auto-correlation and $\tilde{D}\tilde{v}_t$ is a tensor product $\tilde{D} \otimes \tilde{v}_t$ with $\tilde{D}(\xi_1, \dots, \xi_d) = (i \sin(2\pi\xi_1), \dots, i \sin(2\pi\xi_d))$ representing the Fourier frequencies of a central difference Jacobian matrix D . Operator $\tilde{\nabla} \cdot$ is the discrete divergence of a vector field \tilde{v} , which is computed as the sum of the Fourier coefficients of the central difference operator \tilde{D} along each dimension, i.e., $\tilde{\nabla} \cdot (\xi_1, \dots, \xi_d) = \sum_{j=1}^d i \sin(2\pi\xi_j)$.

3 Frequency Diffeomorphisms and Signal Decomposition

While geodesic shooting in the Fourier space (7) is efficient, integrating the ordinary differential equation (2) to compute the corresponding diffeomorphism from a velocity field remains computationally intensive. To address this problem, we introduce a Fourier representation of diffeomorphisms that is simple and easy to manipulate in the bandlimited space. The proposed representation promises improved computational efficiency of any algorithm that requires generation of diffeomorphisms from velocity fields.

Analogous to the continuous inverse flow in (3), we define a sequence of time-dependent inverse diffeomorphisms ϕ_t^{-1} in the Fourier domain that consequently generates a path of geodesic flow. Because that the pointwise multiplication of two vector fields in the spatial domain corresponds to convolution in the Fourier domain, we can easily compute the multiplication of a square matrix and a vector field in the Fourier domain as a single convolution for each row of the matrix.

Let $\widetilde{\text{Diff}}(\Omega)$ denote the space of Fourier representations of diffeomorphisms. To simplify the notation, we use $\psi \triangleq \phi^{-1}$ in the remainder of this section. Given time-dependent velocity field $\tilde{v}_t \in \tilde{V}$, the diffeomorphism $\tilde{\psi}_t \in \widetilde{\text{Diff}}(\Omega)$ in the finite-dimensional Fourier domain can be computed as

$$\frac{d\tilde{\psi}_t}{dt} = -\tilde{D}\tilde{\psi}_t \star \tilde{v}_t, \quad (8)$$

where $*$ is a circular convolution with zero padding to avoid aliasing. To prevent the domain from growing infinity, we truncate the output of the convolution in each dimension to a suitable finite set.

Representing diffeomorphisms entirely in the Fourier space eliminates the effort of converting their associated velocity fields back and forth from the Fourier domain to the spatial domain. While current implementations of computing a diffeomorphism in (3) have complexity $O(N^d)$ on a full image grid of linear size N , the complexity of naively integrating (8) is $O(N^d \log N)$ if the convolution operator is implemented via fast Fourier transform (FFT) [18]. Here we show how to reduce the complexity of (8) via frequency decomposition of the deformations.

In particular, we consider a representation $\tilde{\psi} = \tilde{e} + \tilde{u}$, where \tilde{e} is the Fourier transform of the identity transformation, and \tilde{u} is the Fourier transform of the displacement field. We can now isolate the frequency response of the identity transformation in (8) as follows:

$$\frac{d\tilde{\psi}_t}{dt} = -\tilde{D}(\tilde{e} + \tilde{u}_t) * \tilde{v}_t = -\tilde{D}\tilde{e} * \tilde{v}_t - \tilde{D}\tilde{u}_t * \tilde{v}_t. \quad (9)$$

Since $De \cdot v_t = v_t$ in the spatial domain, we have

$$\tilde{D}\tilde{e} * \tilde{v}_t = \mathcal{F}(De \cdot v_t) = \mathcal{F}(v_t) = \tilde{v}_t. \quad (10)$$

Substituting (10) into (9), we arrive at

$$\begin{aligned} \frac{d\tilde{\psi}_t}{dt} &= -\tilde{v}_t - \tilde{D}\tilde{u}_t * \tilde{v}_t \quad \text{with} \quad \tilde{\psi}_0 = \tilde{e}, \quad \text{or} \\ \tilde{u}_t &= -\int_0^t \tilde{v}_\tau + \tilde{D}\tilde{u}_\tau * \tilde{v}_\tau d\tau \quad \text{with} \quad \tilde{u}_0 = 0. \end{aligned} \quad (11)$$

Importantly, we observe that the high frequency components of the diffeomorphisms that come from the initial condition $\tilde{\psi}_0 = \tilde{e}$ remain unchanged, and only low frequency components vary throughout the integration. Moreover, the evolution scheme for geodesic shooting in the Fourier domain (7) indeed maintains \tilde{v}_t as a bandlimited signal. The truncated convolution operation does not introduce high frequencies if the displacement \tilde{u}_t is also bandlimited. Fig. 1 illustrates a 1D example of the integration (11).

The initial condition $\tilde{\psi}_0 = \tilde{e}$ corresponds to $\tilde{u}_0 = 0$. We first integrate (11) for bandlimited low frequency components of the signal and then add the high frequency components back. Therefore we have

$$\begin{aligned} \text{for } \xi \leq \eta, \quad \tilde{u}_t(\xi) &= -\int_0^t \tilde{v}_\tau(\xi) + \tilde{D}\tilde{u}_\tau(\xi) * \tilde{v}_\tau(\xi) d\tau, \\ \text{for } \xi > \eta, \quad \tilde{u}_t(\xi) &= 0, \end{aligned} \quad (12)$$

where $\xi = (\xi_1, \dots, \xi_d)$ and $\eta = (\eta_1, \dots, \eta_d)$ is the vector of upper bounds on the frequency in the bandlimited representation of ψ .

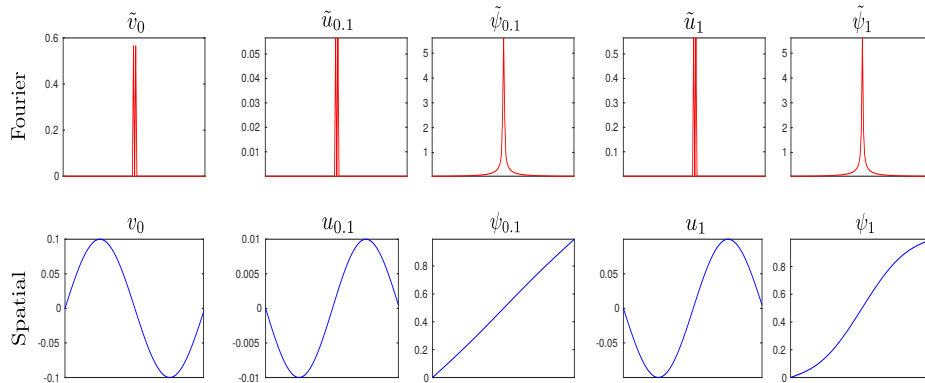


Fig. 1: A 1D example of an initial velocity v_0 as a sinusoid function and the resulting displacement field u_t and diffeomorphism ψ_t at $t = 0.1$ and $t = 1$. Both Fourier (red) and spatial (blue) representations are shown.

3.1 Frequency Diffeomorphisms for Image Registration

In this section, we present a diffeomorphic image registration algorithm based on geodesic shooting that is carried out entirely in the Fourier space.

Let S be the source image and T be the target image defined on a torus domain $\Gamma = \mathbb{R}^d/\mathbb{Z}^d$ ($S(x), T(x) : \Gamma \rightarrow \mathbb{R}$). The problem of diffeomorphic image registration is to find the shortest path of diffeomorphisms $\psi_t \in \text{Diff}(\Gamma) : \Gamma \rightarrow \Gamma, t \in [0, 1]$ such that $S \circ \psi_1$ is similar to T , where \circ is a composition operator that resamples S by the smooth mapping ψ_1 . LDDMM with geodesic shooting [20] leads to a gradient decent optimization of an explicit energy function

$$E(v_0) = \frac{\lambda}{2} \text{dist}(S \circ \psi_1, T) + \frac{1}{2} \|v_0\|_V^2 \quad (13)$$

under the constraints (2) and (4). The distance function $\text{dist}(\cdot, \cdot)$ measures the dissimilarity between images. Commonly used distance metrics include sum-of-squared difference (L^2 -norm) of image intensities, normalized cross correlation (NCC), and mutual information (MI). Here $\lambda > 0$ is a weight parameter.

The energy function in the finite-dimensional Fourier space can be equivalently formulated as

$$E(\tilde{v}_0) = \frac{\lambda}{2} \text{dist}(S \circ \psi_1, T) + \frac{1}{2} \|\tilde{v}_0\|_{\tilde{V}}^2 \quad (14)$$

with the new constraints (7) and (8).

We use a gradient decent algorithm on the initial velocity \tilde{v}_0 to estimate the path of diffeomorphic flow $\{\tilde{\psi}_t\}$ entirely in the bandlimited space. Beginning with the initialization $\tilde{v}_0 = 0$, the gradient of the energy (14) is computed via two steps below:

1. Compute the gradient $\nabla_{\tilde{v}_1} E$ of the energy (14) at $t = 1$. This requires integrating both the diffeomorphism $\tilde{\psi}_t$ and the velocity field \tilde{v}_t forward in time, and then mapping $\tilde{\psi}_1$ to the spatial domain to obtain ψ_1 . Formally,

$$\nabla_{\tilde{v}_1} E = \mathcal{K} \mathcal{F} \left(\frac{\partial \text{dist}(S \circ \psi_1, T)}{\partial (S \circ \psi_1)} \cdot \nabla (S \circ \psi_1) \right). \quad (15)$$

2. Bring the gradient $\nabla_{\tilde{v}_1} E$ in (15) back to $t = 0$ by reduced adjoint Jacobi field equations [8, 23]. Integrate the reduced adjoint Jacobi field equations

$$\frac{d\hat{v}}{dt} = -\text{ad}_{\hat{v}}^\dagger \hat{h}, \quad \frac{d\hat{h}}{dt} = -\hat{v} - \text{ad}_{\hat{v}} \hat{h} + \text{ad}_{\hat{h}}^\dagger \tilde{v} \quad (16)$$

in \tilde{V} to get the gradient update $\nabla_{\tilde{v}_0} E$, where ad^\dagger is an adjoint operator and $\hat{v}, \hat{h} \in \tilde{V}$ are introduced adjoint variables.

The algorithm is summarized below.

Algorithm 1: Frequency Diffeomorphisms for Image Registration

Input: source image S , target image T .

Initialize $\tilde{v}_0 = 0$.

repeat

(a) Integrate (7) to compute \tilde{v}_t at discrete time points $t = [0, \dots, 1]$.

(b) Integrate (8) to generate $\tilde{\psi}_t$.

(c_i) Convert $\tilde{\psi}_1$ back to the spatial domain to transport the source image S .

(c_{ii}) Compute the gradient $\nabla_{\tilde{v}_1} E$ (15) at time point $t = 1$.

(d) Integrate $\nabla_{\tilde{v}_1} E$ backward in time via (16) to obtain $\nabla_{\tilde{v}_0} E$.

(e) Update $\tilde{v}_0 \leftarrow \tilde{v}_0 - \delta \nabla_{\tilde{v}_0} E$, where δ is the step size.

until convergence

Computational Complexity It has previously been shown that the complexity of steps (a) and (d) is $O(Tn^d)$, where T is the number of time steps in the integration and n is the truncated dimension in the bandlimited space. The complexity of the current methods for computing diffeomorphisms in the high-dimensional image space via (2) is $O(TN^d)$. In contrast, our algorithm reduces the complexity of this step to $O(Tn^d)$ (step (b)). To transport the images and measure the image dissimilarity at $t = 1$, we convert the transformation $\tilde{\psi}_1$ into the spatial domain via FFT ($O(N^d \log N)$) (step (c_i)).

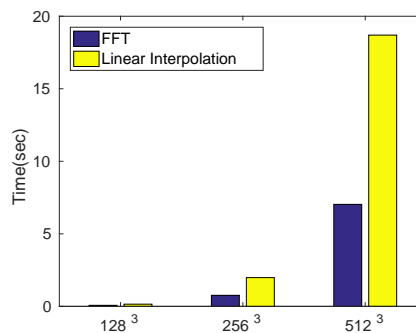


Fig. 2: Exact run-time of FFT (blue) and linear interpolation (yellow).

While the theoretical complexity of FFT is higher than the complexity of computing $S \circ \psi_1^{-1}$, which is $O(N^d)$, its empirical runtime for N a power of 2 is quite low by comparison. Fig. 2 reports the empirical runtime of FFT and of linear interpolation for different image grid sizes ($N = 2^7, 2^8, 2^9$) as used by the current methods and ours to transport S . We note that the linear interpolation requires more than twice the amount of time than that of FFT.

4 Experimental Evaluation

To evaluate the proposed approach, we perform registration-based segmentation and examine the resulting segmentation accuracy, runtime and memory consumption of the algorithm. We compare the proposed method with the diffeomorphic demons implementation in ANTS software package [5] and the state-of-the-art fast geodesic shooting for LDDMM method FLASH [23] (downloaded from: <https://bitbucket.org/FlashC/flashc>). In all experiments, we set $\alpha = 1.5, c = 3.0, \lambda = 1.0e4$ and $T = 10$ for the time integration. A normalized cross correlation (NCC) metric for image dissimilarity and a multi-resolution optimization scheme with three levels are used in all three methods. To evaluate volume overlap between the propagated segmentation A and the manual segmentation B for each structure, we compute the Dice Similarity Coefficient $DSC(A, B) = 2(|A \cap B|)/(|A| + |B|)$ where \cap denotes an intersection of two regions.

4.1 Data

We evaluate the method on 3D brain MRI scans [9] and 4D in-utero MRI time series [11].

3D brain MRI Thirty six brain MRI scans (T1-weighted MP-RAGE) were acquired in normal subjects and patients with Alzheimer’s disease across a broad age range. The MRI images are of dimension 256^3 , 1mm isotropic voxels and were computed by averaging three or four scans. All scans underwent skull stripping, intensity normalization, bias field correction, and co-registration with affine transformation. An atlas was built from 20 images as a reference. Manual segmentations are available for all scans in the set. We perform image registration of the 3D brain atlas to the remaining 16 subjects. We evaluate registration by examining the accuracy of atlas-based delineations for white matter (WM), cortex (Cor), ventricles (Vent), hippocampus (Hipp), and caudate (Caud).

4D in-utero time-series volumetric MRI Ten pregnant women (three singleton pregnancies, six twin pregnancies, and one triplet pregnancy) were recruited and consented. Single-shot GRE-EPI image series were acquired for each woman on a 3T MR scanner (Skyra Siemens, 18-channel body and 12-channel spine receive arrays, $3 \times 3\text{mm}^2$ in-plane resolution, 3mm slice thickness, interleaved slice acquisition, $TR = 5.8 - 8\text{s}$, $TE = 32 - 36\text{ms}$, $FA = 90^\circ$). Odd and even slices of each volume were resampled onto an isotropic 3mm^3 image grid to reduce the effects of interleaved acquisition. Each series includes around 300 3D volumes. The placentae (total of 10) and fetal brains (total of 18) were




Fig. 3: Left: average final energy for different values of the truncated dimension $n = 8, 16, \dots, 128$. Right: example propagated segmentation with 35 structures obtained by our method. 2D slices are shown for visualization only, all computations are carried out fully in 3D.

manually delineated in the reference template and in five additional volumes in each series. When applying our method to in-utero MRI, we use it as part of sequential registration of the consecutive frames in the scan series. Each series requires about 300 consecutive image registration steps over time.

4.2 Results

3D brain MRI Fig. 3 reports the total energy (14) averaged over 16 test images for different values of truncated dimension $n = 8, 16, \dots, 128$. Our method arrives at the same solution at $n = 32$ and higher. For the remainder of this section, we use $n = 32$ to illustrate the results. An example segmentation obtained by our algorithm is also illustrated in Fig. 3. Fig. 4 reports the volume overlap of segmentations for our method and the two baseline algorithms. All three algorithms produce comparable segmentation accuracy. The difference is not statistically significant in a paired t-test for each labelled structure ($p = 0.369$). Our algorithm has substantially lower computational cost than ANTS and FLASH. Fig. 4 also provides runtime and memory consumption across all methods.

4D in-utero time-series volumetric MRI Similar to the previous experiment, we cross-validated the optimal truncated dimension n at different scales and set $n = 16$ for the 4D in-utero time series. Once all the volumes in the series are aligned, we transform the manual segmentations in the first volume to other volumes in each series by using the estimated deformations. Fig. 5 illustrates results for an example case from the study. We observe that the delineations achieved by transferring manual segmentations from the reference frame to the coordinate system of the target frame 25 align fairly well with the manual segmentations. Fig. 6 reports segmentation volume overlap for the fetal brains and placenta, as well as the time and memory consumption for our method and the two baseline algorithms. Again, our algorithm achieves comparable results

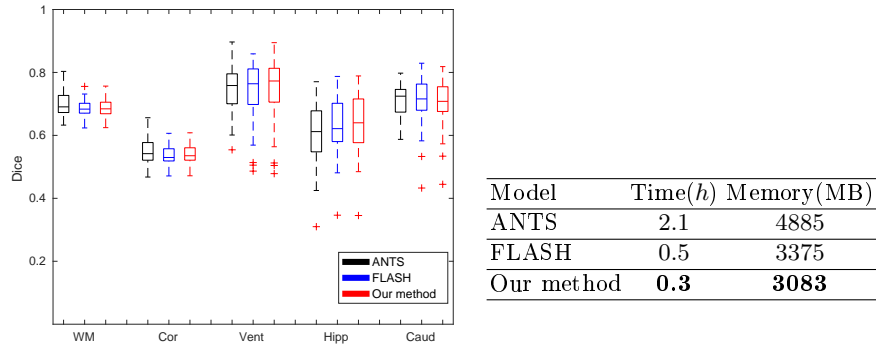


Fig. 4: Left: volume overlap between atlas-based and manual segmentations for five important regions (white matter, cortex, ventricles, hippocampus, and caudate) estimated via ANTS (black), FLASH (blue), and our method (red). Right: Runtime and memory consumption per image for all three methods.

(paired t-test $p = 0.4671$) while offering significant improvements in computational efficiency.




Fig. 5: An example case from the in-utero MRI study. Left to right: source with manual segmentation, deformed source, target with manual segmentation, and target with propagated segmentations for the fetal brains (green) and placenta (pink). 2D slices of axial view are shown for visualization only, all computations are carried out fully in 3D.

5 Conclusion

We presented an efficient way to compute diffeomorphisms in the setting of LD-DMM with geodesic shooting for image registration. Our method is the first to represent diffeomorphisms in the Fourier space, which provides a way to compute transformations from the associated velocity fields entirely in the low di-

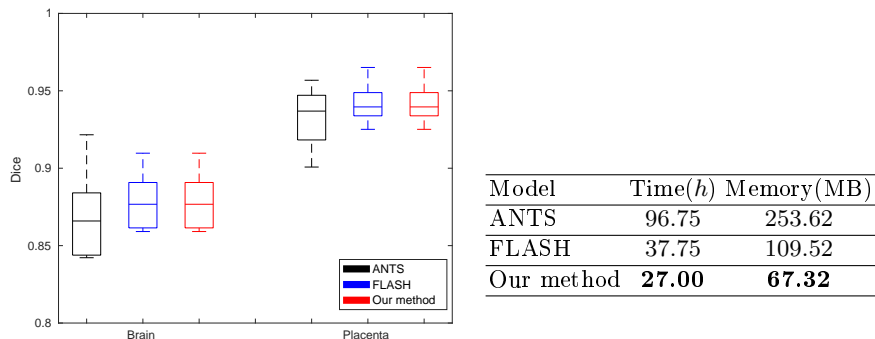


Fig. 6: Left: volume overlap between transferred and manual segmentations of fetal brains and placenta averaged over all subjects. Statistics are reported for ANTS (black), FLASH (blue), and our method (red). Right: Runtime and memory consumption of 300 registrations for all three algorithms.

mensional bandlimited space. This approach reduces the computational cost of the algorithms without loss in accuracy. The theoretical tools employed in this work are not only broadly applicable to other representations of diffeomorphic transformations such as stationary velocity fields, but also to the standard path optimization strategy in the original LDDMM. Our method can be naturally interpreted as operating with the left-invariant metric on the space of diffeomorphisms [15], which accepts spatially-varying smoothing kernels in special image registration scenarios where different regularizations are required at different locations. Future research will explore more of numerical analysis, as well as its theoretical connection and implications for diffeomorphic image registration.

Acknowledgments This work was supported by NIH NIBIB NAC P41EB015902, NIH NINDS R01NS086905, NIH NICHD U01HD087211, and Wistron Corporation.

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*, 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.: Unified segmentation. *Neuroimage* 26(3), 839–851 (2005)
5. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* 12(1), 26–41 (2008)

6. 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)
7. 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)
8. Francesco, B.: Invariant affine connections and controllability on lie groups. Tech. rep., technical Report for Geometric Mechanics, California Institute of Technology (1995)
9. Johnson, K.A., Jones, K., Holman, B.L., Becker, J.A., Spiers, P.A., Satlin, A., Albert, M.S.: Preclinical prediction of alzheimer’s disease using spect. *Neurology* 50(6), 1563–1571 (1998)
10. Joshi, S., Davis, B., Jomier, M., Gerig, G.: Unbiased diffeomorphic atlas construction for computational anatomy. *NeuroImage* 23, Supplement1, 151–160 (2004)
11. Liao, R., Turk, E.A., Zhang, M., Luo, J., Grant, P.E., Adalsteinsson, E., Golland, P.: Temporal registration in in-utero volumetric mri time series. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 54–62. Springer (2016)
12. Miller, M.I., Trouvé, A., Younes, L.: Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision* 24(2), 209–228 (2006)
13. Niethammer, M., Huang, Y., Vialard, F.X.: Geodesic regression for image time-series. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 655–662. Springer (2011)
14. Qiu, A., Younes, L., Miller, M.I.: Principal component based diffeomorphic surface mapping. *Medical Imaging, IEEE Transactions on* 31(2), 302–311 (2012)
15. Schmah, T., Risser, L., Vialard, F.X.: Diffeomorphic image matching with left-invariant metrics. In: *Geometry, Mechanics, and Dynamics*, pp. 373–392. Springer (2015)
16. 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*, pp. 529–537. Springer (2010)
17. Singh, N., Vialard, F.X., Niethammer, M.: Splines for diffeomorphisms. *Medical image analysis* 25(1), 56–71 (2015)
18. Van Loan, C.: *Computational frameworks for the fast Fourier transform*, vol. 10. Siam (1992)
19. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* 45(1), S61–S72 (2009)
20. 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)
21. Younes, L., Arrate, F., Miller, M.: Evolutions equations in computational anatomy. *NeuroImage* 45(1S1), 40–50 (2009)
22. Zhang, M., Fletcher, P.T.: Bayesian principal geodesic analysis in diffeomorphic image registration. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 121–128. Springer (2014)
23. 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)