



# NeurEPDiff: Neural Operators to Predict Geodesics in Deformation Spaces

Nian Wu<sup>1,2(✉)</sup> and Miaomiao Zhang<sup>2,3</sup>

<sup>1</sup> Computer Science, East China Normal University, Shanghai, China  
nellie\_nw@stu.ecnu.edu.cn

<sup>2</sup> Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, USA

<sup>3</sup> Computer Science, University of Virginia, Charlottesville, VA, USA

**Abstract.** This paper presents NeurEPDiff, a novel network to fast predict the geodesics in deformation spaces generated by a well known Euler-Poincaré differential equation (EPDiff). To achieve this, we develop a neural operator that for the first time learns the evolving trajectory of geodesic deformations parameterized in the tangent space of diffeomorphisms (a.k.a velocity fields). In contrast to previous methods that purely fit the training images, our proposed NeurEPDiff learns a nonlinear mapping function between the time-dependent velocity fields. A composition of integral operators and smooth activation functions is formulated in each layer of NeurEPDiff to effectively approximate such mappings. The fact that NeurEPDiff is able to rapidly provide the numerical solution of EPDiff (given any initial condition) results in a significantly reduced computational cost of geodesic shooting of diffeomorphisms in a high-dimensional image space. Additionally, the properties of discretization/resolution-invariant of NeurEPDiff make its performance generalizable to multiple image resolutions after being trained offline. We demonstrate the effectiveness of NeurEPDiff in registering two image datasets: 2D synthetic data and 3D brain resonance imaging (MRI). The registration accuracy and computational efficiency are compared with the state-of-the-art diffeomorphic registration algorithms with geodesic shooting.

## 1 Introduction

Deformable image registration is a fundamental tool in medical image analysis, for example, computational anatomy and shape analysis [20, 37], statistical analysis of groupwise images [21, 31], and template-based image segmentation [4, 9]. Among many applications, it is ideal that the transformation is diffeomorphic (i.e., bijective, smooth, and invertible smooth) to maintain an intact topology of objects presented in images. A bountiful literature has studied various parameterizations of diffeomorphisms, including stationary velocity fields [2, 3], B-spline free-form deformations [26, 27], and the large deformation diffeomorphic metric mapping (LDDMM) [7, 29]. In this paper, we focus on LDDMM, which provides a rigorous mathematical definition of distance metrics on the space of diffeomorphisms. Such distance metrics play an important role in deformation-based shape analysis of images, such as geometric shape regression [13, 20], longitudinal shape analysis [28], and group shape comparisons [18, 22].

Despite the advantages of LDDMM [7], the substantial time and computational cost involved in searching for the optimal diffeomorphism on high-dimensional image grids limits its applicability in real-world applications. To alleviate this problem, a geodesic shooting algorithm [29] was developed to reformulate the optimization of LDDMM (originally over a sequence of time-dependent diffeomorphisms) equivalently in the tangent space of diffeomorphisms at time point zero (a.k.a., initial velocity fields). The geodesic is then uniquely determined by integrating a given initial velocity field via the Euler-Poincaré differential equation (EPDiff) [1, 19]. This geodesic shooting algorithm demonstrated a faster convergence and avoided having to store the entire time-varying velocity fields from one iterative search to the next. A recent work FLASH has further reduced the computational cost of geodesic shooting by reparameterizing the velocity fields in a low-dimensional bandlimited space [35, 36]. The numerical solution to EPDiff was rapidly computed with much fewer parameters in a frequency domain. In spite of a significant speedup, it still takes FLASH minutes to compute when the dimension of the bandlimited space grows higher.

In this paper, we present a novel method, NeurEPDiff, that can fast predict the geodesics in deformation spaces generated by EPDiff. Inspired by the recent achievements on learning-based partial/ordinary differential equation solvers [16, 17], we develop a neural operator to learn the evolving functions/mappings of geodesics in the space of diffeomorphisms parameterized by bandlimited initial velocity fields. In contrast to current deep learning-based approaches that are tied to the resolution and discretization of training images [6, 8, 24, 30, 33], our proposed method shows consistent performances on different scales of image resolutions without the need of re-training. We also define a composition of integral operators and smooth activation functions in each layer of the NeurEPDiff network to effectively approximate the mapping functions of EPDiff over time. To summarize, our main contributions are threefold:

- (i) We are the first to develop a neural operator that fast predicts a numerical solution to the EPDiff [1, 19], which is a critical step to generate geodesics of diffeomorphisms.
- (ii) A composition of integral operators and smooth activation functions is defined to approximate the mapping functions of EPDiff characterized in the low-dimensional bandlimited space. This makes the training and inference of our proposed NeurEPDiff much more efficient in high-dimensional image spaces.
- (iii) The performance of NeurEPDiff is invariant to image resolutions. That is to say, once it is trained, this operator can be tested on multiple image resolutions without the need of being re-trained.

To demonstrate the effectiveness of our proposed model in diffeomorphic image registration, we run experiments on both 2D synthetic data and 3D real brain MRIs. We first show that a trained NeurEPDiff is able to consistently predict the diffeomorphic transformations at multiple different image resolutions. We then compare the registration accuracy (via registration-based image segmentation) and computational efficiency of NeurEPDiff vs. the state-of-the-art LDDMM with geodesic shooting [29, 35, 36]. Experimental results show that our method outperforms the baselines in term of time consumption, while with a comparable registration accuracy.

## 2 Background: Geodesics of Diffeomorphisms via EPDiff

In this section, we first briefly review the basic concepts of geodesic shooting, which generates geodesics of diffeomorphisms via EPDiff [29, 34]. We then show how to optimize the problem of diffeomorphic image registration in the setting of LDDMM [7] with geodesic shooting [29, 34].

**Geodesic Shooting via EPDiff.** Let  $\text{Diff}^\infty(\Omega)$  denote the space of smooth diffeomorphisms on an image domain  $\Omega$ . The tangent space of diffeomorphisms is the space  $V = \mathfrak{X}^\infty(T\Omega)$  of smooth vector fields on  $\Omega$ . Consider a time-varying velocity field,  $\{v_t\} : [0, \tau] \rightarrow V$ , we can generate diffeomorphisms  $\{\phi_t\}$  as a solution to the equation

$$\frac{d\phi_t}{dt} = v_t(\phi_t), \quad t \in [0, \tau]. \quad (1)$$

The geodesic shooting algorithm [23] states that a geodesic path of the diffeomorphisms (in Eq. (1)) is uniquely determined by integrating a well-studied EPDiff [1, 19] with an initial condition. That is, given an initial velocity,  $v_0 \in V$ , at  $t = 0$ , a geodesic path  $t \mapsto \phi_t \in \text{Diff}^\infty(\Omega)$  in the space of diffeomorphisms can be computed by forward shooting the EPDiff equation

$$\frac{\partial v_t}{\partial t} = -K [(Dv)^T m_t + Dm_t v_t + m_t \text{div } v_t], \quad (2)$$

where  $D$  denotes the Jacobian matrix. Here  $K$  is an inverse operator of  $L : V \rightarrow V^*$ , which is a positive-definite differential operator that maps a tangent vector  $v \in V$  into the dual space  $m \in V^*$ . This paper employs a commonly used Laplacian operator  $L = (-\alpha\Delta + e)^c$ , where  $\alpha$  is a positive weight parameter,  $e$  denotes an identity matrix, and  $c$  is a smoothness parameters. A larger value of  $\alpha$  and  $c$  indicates more smoothness.

A recent model FLASH has later developed a Fourier variant of the EPDiff and demonstrated that the numerical solution to the original EPDiff can be efficiently computed in a low-dimensional bandlimited space with dramatically reduced computational cost [35, 36]. Let  $\widehat{\text{Diff}}(\Omega)$  and  $\widehat{V}$  denote the space of Fourier representations of diffeomorphisms and velocity fields respectively. The equation of EPDiff is reformulated in a bandlimited space as

$$\frac{\partial \tilde{v}_t}{\partial t} = -\tilde{K} \left[ (\tilde{D}\tilde{v}_t)^T \star \tilde{m}_t + \tilde{\nabla} \cdot (\tilde{m}_t \otimes \tilde{v}_t) \right], \quad (3)$$

where  $\tilde{v}_t$  and  $\tilde{m}_t$  are the Fourier transforms of velocity fields  $v_t$  and momentum  $m_t$  respectively. The  $\tilde{D}$  represents Fourier frequencies of a Jacobian matrix  $D$  with central difference approximation,  $\star$  is the truncated matrix-vector field auto-correlation with zero-padding, and  $\otimes$  denotes a tensor product. The Fourier coefficients of a laplacian operator  $L$  is  $\tilde{L}(\xi_1, \dots, \xi_d) = \left( -2\alpha \sum_{j=1}^d (\cos(2\pi\xi_j) - 1) + 1 \right)^c$ , where  $(\xi_1, \dots, \xi_d)$  is a  $d$ -dimensional frequency vector.

**LDDMM with Geodesic Shooting.** Given a source image  $S$  and a target image  $T$  defined on a  $d$ -dimensional torus domain  $\Omega = \mathbb{R}^d / \mathbb{Z}^d$  ( $S(x), T(x) : \Omega \rightarrow \mathbb{R}$ ). The space of diffeomorphisms is denoted by  $\text{Diff}(\Omega)$ . The problem of diffeomorphic image

registration is to find the shortest path, i.e., geodesic, to generate time-varying diffeomorphisms  $\{\phi_t\} : t \in [0, \tau]$ , such that a deformed source image by the smooth mapping  $\phi_\tau$ , noted as  $S(\phi_\tau)$ , is similar to  $T$ . By parameterizing the transformation  $\phi_\tau$  with an initial velocity field  $\tilde{v}_0$ , we write the optimization of diffeomorphic registration in the setting of LDDMM with geodesic shooting in a low-dimensional bandlimited space as

$$E(\tilde{v}_0) = \lambda \text{Dist}(S(\phi_\tau), T) + \frac{1}{2} \|\tilde{v}_0\|_{\mathcal{V}}^2 \text{ s.t. Eq. (1) and (3)} \quad (4)$$

Here,  $\text{Dist}(\cdot, \cdot)$  is a distance function that measures the dissimilarity between images and  $\lambda$  is a positive weighting parameter. The commonly used distance functions include the sum-of-squared intensity differences ( $L_2$ -norm) [7], normalized cross correlation (NCC) [5], and mutual information (MI) [32]. In this paper, we will use the sum-of-squared intensity differences. The  $\|\cdot\|_{\mathcal{V}}$  represents a Sobolev space that enforces smoothness of the velocity fields.

Despite a dramatic speed up of LDDMM with the reduced cost of Fourier EPDiff for shooting [35, 36], the numerical solution operator in Eq. (3) does not scale well due to the introduced correlation and tensor convolutions.

### 3 Our Method: NeurEPDiff

We introduce a novel network, NeurEPDiff, that learns the solution operator to EPDiff in a low-dimensional bandlimited space. Inspired by the recent work on neural operators [15, 16], we propose to develop NeurEPDiff,  $\mathcal{G}_\theta$  with parameters  $\theta$ , as a surrogate model to approximate the solution operator to EPDiff. Once it is trained, our NeurEPDiff can be used to fast predict the original EPDiff with a given initial condition  $\tilde{v}_0$  on various resolutions of image grids.

Given a set of  $N$  observations,  $\{\tilde{v}_0^n, \tilde{v}^n\}_{n=1}^N$ , where  $\tilde{v}^n \triangleq \{\tilde{v}_1^n, \dots, \tilde{v}_\tau^n\}$  is a time-sequence of numerical solutions to EPDiff. The solution operator can be learned by minimizing the empirical data loss defined as

$$E(\theta) = \frac{1}{N} \sum_{n=1}^N \|\tilde{v}^n - \mathcal{G}_\theta(v_0^n)\|_{\mathcal{V}}^2 + \beta \cdot \text{Reg}(\theta), \quad (5)$$

where  $\beta$  is a weighting parameter and  $\text{Reg}(\cdot)$  regularizes the network parameter  $\theta$ .

Following the similar principles in [16], we formulate  $\mathcal{G}_\theta$  as an iterative architecture  $\tilde{v}_0 \mapsto \tilde{v}_1 \mapsto \dots \mapsto \tilde{v}_\tau$ , which is a sequence of functions carefully parameterized in the hidden network layers. The input  $\tilde{v}_0$  is first lifted to a higher dimensional representation by a complex-valued local transformation  $\tilde{P}$ , which is usually parameterized by a shallow fully connected neural network. We then develop a key component, called iterative evolution layer, with operations designed in the bandlimited space to learn the nonlinear function of EPDiff.

**Iterative Evolution Layer.** Let  $\tilde{z} \in \mathbb{C}^d$  be a  $d$ -dimensional complex-valued latent feature vector encoded from the initial velocity  $\tilde{v}_0$ . An iterative evolution layer (with the number of  $J$  hidden layers) to update  $\tilde{z}_j \mapsto \tilde{z}_{j+1}$ , where  $j \in \{1, \dots, J\}$ , is defined

as the composition of a nonlinear smooth activation function  $\sigma$  of a complex-valued linear transform  $\tilde{W}_j$  with a global convolutional kernel  $\tilde{\mathcal{H}}_j$ :

$$\tilde{z}_{j+1} := \sigma(\tilde{W}_j \tilde{z}_j + \tilde{\mathcal{H}}_j * \tilde{z}_j). \tag{6}$$

The  $*$  represents a complex convolution, which is efficiently computed by multiplying two signals in a real space. Similar to [25, 30], we employ a complex-valued Gaussian error linear unit (CGeLU) that applies real-valued activation functions separately to the real and imaginary parts. With a smoothing operator  $\tilde{K}$  to ensure the smoothness of the output signal, we have  $\sigma := \tilde{K}(\text{GeLU}(\mathcal{R}(\cdot)) + i\text{GeLU}(\mathcal{I}(\cdot)))$ , where  $\mathcal{R}(\cdot)$  denotes the real part of a complex-valued vector, and  $\mathcal{I}(\cdot)$  denotes an imaginary part.

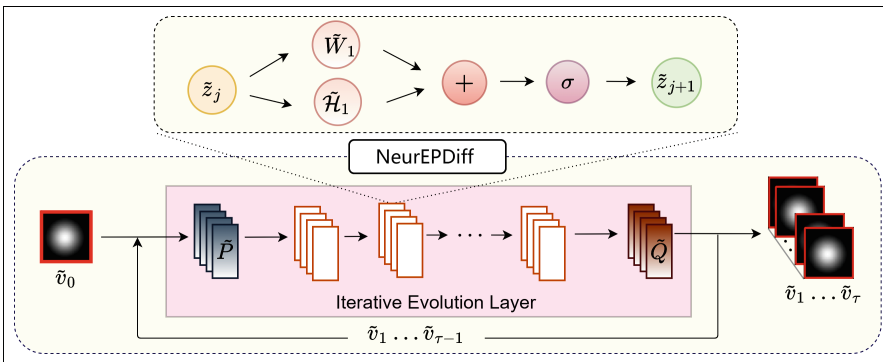
We are now ready to introduce a formal definition of the NeurEPDiff  $\mathcal{G}_\theta$  developed in a bandlimited space as follows.

**Definition 1 (Fourier NeurEPDiff).** *The neural operator  $\mathcal{G}_\theta$  is defined as a composition of a complex-valued encoder  $\tilde{P}$  (a decoder  $\tilde{Q}$ ) (a.k.a. local linear transformations) to project the lower dimension function into higher dimensional space and vice versa, and a nonlinear smooth activation function  $\sigma$  of a local linear transform  $\tilde{W}$  with a non-local convolution kernel  $\tilde{\mathcal{H}}$  parameterized by network parameters, i.e.,*

$$\mathcal{G}_\theta := \tilde{Q} \circ \sigma(\tilde{W}_J, \tilde{\mathcal{H}}_J) \circ \dots \circ \sigma(\tilde{W}_1, \tilde{\mathcal{H}}_1) \circ \tilde{P}. \tag{7}$$

The  $\circ$  denotes function composition, and the parameter  $\theta$  includes all parameters from  $\{\tilde{P}, \tilde{Q}, \tilde{W}, \tilde{\mathcal{H}}\}$ .

The network architecture of our proposed NeurEPDiff is shown in Fig. 1.



**Fig. 1.** An overview of our proposed NeurEPDiff. Given an initial input  $\tilde{v}_0$ , the iterative evolution layer updates the time sequence of  $\{\tilde{v}_1, \dots, \tilde{v}_\tau\}$ .

### 3.1 NeurEPDiff for Diffeomorphic Image Registration

This section introduces a new diffeomorphic image registration algorithm with our developed NeurEPDiff. We will first train a neural operator  $\mathcal{G}_\theta$  offline and then use

it as predictive model in each iteration of the optimization of LDDMM (Eq. (4)), where the numerical solution of EPDiff is required. In this paper, we use a sum-of-squares-distance (SSD) metric for image dissimilarity. The optimization

$$E(\tilde{v}_0) = \frac{1}{2\lambda^2} \|(S(\phi_{\tau=1}) - T\|_2^2 + \frac{1}{2} \|\tilde{v}_0\|_V^2, \text{ s.t. Eq. (1) \& } \mathcal{G}_\theta(\tilde{v}_0). \quad (8)$$

The final transformation  $\phi_{\tau=1}$  is computed by Eq. (1) after predicting  $\{\tilde{v}_t\}$  from a well-trained network  $\mathcal{G}_\theta$ .

**Inference.** Similar to [29, 36], we employ a gradient decent algorithm to minimize the problem (8). The three main steps of our inference are described below:


- (i) Forward geodesic shooting by NeurEPDiff. Given an initialized velocity field  $\tilde{v}_0$ , we compute the transformation  $\phi_\tau$  by first predicting  $\{\tilde{v}_t\}$  and then solving Eq. (1).
- (ii) Compute the gradient  $\nabla_{\tilde{v}_\tau} E$  at the ending time point  $t = \tau$  as

$$\nabla_{\tilde{v}_\tau} E = \tilde{K} \mathcal{F} \left( \frac{1}{\lambda^2} (S(\phi_\tau) - T) \cdot \nabla (S(\phi_\tau)) \right). \quad (9)$$

Recall that  $\mathcal{F}$  denotes a Fourier transformation.

- (iii) Backward integration to bring the gradient  $\nabla_{\tilde{v}_\tau} E$  back to  $t = 0$  by integrating the reduced adjoint Jacobi field equations defined in [35].

Figure 2 visualize the process to optimize image registration with NeurEPDiff.



**Fig. 2.** An illustration of our optimization-based diffeomorphic registration with NeurEPDiff.

## 4 Experimental Evaluation

To evaluate the effectiveness of our proposed model, NeurEPDiff, for diffeomorphic image registration, we compare its performance with the original LDDMM with geodesic shooting [29] and a fast variant of LDDMM (FLASH) [36] on both 2D synthetic data and 3D real brain MR images with different levels of resolution.

The training of NeurEPDiff in all our experiments is implemented on Nvidia GeForce GTX 1080Ti GPUs. We use ADAM optimizer [14] with the batch size as 20, learning rate of  $1e^{-3}$ , and weight decay as  $1e^{-4}$ , and the number of epochs as 500.

For a fair comparison with the baseline algorithms that are optimized by gradient decent on CPU, we also test our model on CPU once it is trained.

We evaluate the registration accuracy by performing registration-based segmentation and examine the resulting segmentation accuracy and runtime. 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 [10].

## 4.1 Data

**2D Synthetic Data.** We first simulate 2200 “bull-eye” synthetic data (as shown in Fig 3) with the resolution  $192^2$  by manipulating the width  $a$  and height  $b$  of an ellipse, formulated as  $\frac{(x-96)^2}{a^2} + \frac{(y-96)^2}{b^2} = 1$ . We draw the parameters  $a, b$  randomly from a Gaussian distribution  $\mathcal{N}(40, 4.2^2)$  for the outer ellipse, and  $a, b \sim \mathcal{N}(17, 1.6^2)$  for the inner ellipse. We then carefully down-sample the images to different resolutions of  $64^3$  and  $128^3$ . The dataset is split into training (2000 volumes of resolution  $64^2$ ) and testing part (200 volumes of resolution  $64^2$ ,  $128^2$ , and  $192^2$  respectively).

**3D Brain MRI.** We include 2200 T1-weighted 3D brain MRI scans from Open Access Series of Imaging Studies (OASIS) [12] in our experiments. We randomly select 2000 of them and down-sampled to the resolution of  $64^3$  for training. The remaining 200 are re-sampled to the resolution of  $64^3$ ,  $128^2$ , and  $192^2$  for testing separately. The 200 testing volumes include manually delineated anatomical structures, i.e., segmentation labels, which also have been carefully down-sampled to the resolution of  $64^3$ ,  $128^3$ , and  $192^3$ . All of the MRI scans have undergone skull-stripping, intensity normalization, bias field correction, and affine alignment.

**Ground Truth Velocity Fields.** We use numerical solutions to Fourier EPDiff in Eq. (3) (i.e., computed by a Euler integrator) as the ground truth data for network training. More advanced integration methods, such as Runge-Kutta RK4, can be easily applied. In all experiments, we first run pairwise image registration using FLASH algorithm [35, 36] on randomly selected images with lower resolution of  $64^2$  or  $64^3$ . We then collect all numerical solutions to Fourier EPDiff as ground truth data. We set all integration steps as 10, the smoothing parameters for the operator  $\tilde{K}$  as  $\alpha = 3.0, c = 3.0$ , the positive weighting parameter of registration as  $\lambda = 0.03$ , and the dimension of the bandlimited velocity field as 16.

## 4.2 Experiments


We first test our NeurEPDiff-based registration methods on 2D synthetic data and compare the final estimated transformation fields with the baseline algorithms. To validate whether our model is resolution-invariant, we train NeurEPDiff on images with lower resolution of  $64^2$ , and then test its prediction accuracy on different levels of higher resolution (i.e., images with  $128^2$  and  $192^2$ ). We examine the optimization energy and compare it with all baseline algorithms.

We further analyze the performance of NeurEPDiff on real 3D brain MRI scans. The convergence graphs of NeurEPDiff-based registration are reported at different image resolutions. To investigate the predictive efficiency of NeurEPDiff in terms of both time-consumption and accuracy, we compare the actual runtime of solving the original EPDiff by our method and all baselines.

The registration accuracy is validated through a registration-based segmentation. Given a source image, we first run pairwise image registration with all algorithms and propagate segmentation labels of the template image to the number of 100 target images. We then examine the resulting segmentation accuracy by computing dice scores of the propagated segmentations and manually delineated labels. We report averaged dice scores of all images on four brain structures, including cortex, subcortical-gray-matter (SGM), white-matter (WM), and cerebrospinal fluid (CSF).

### 4.3 Results

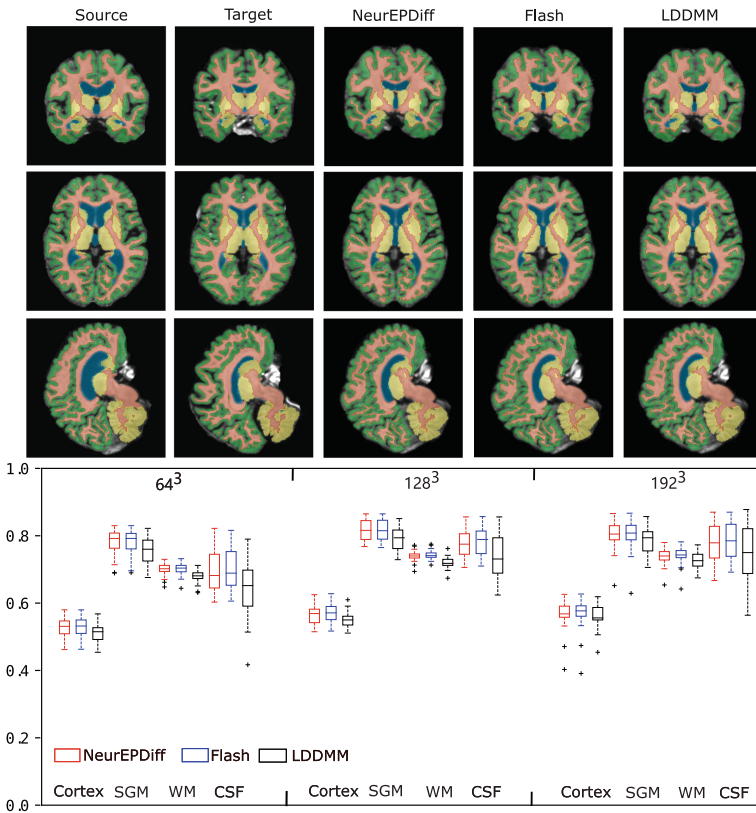
Figure 3 visualizes examples of registration results on 2D synthetic images with the size of  $192^2$ . The comparison is on deformed source images and transformation fields estimated by our NeurEPDiff, FLASH [36], and the original LDDMM with geodesic shooting [29]. Note that even though our method is trained on a low image resolution of  $64^2$ , the predicted results are fairly close to the estimates from registration algorithms of both FLASH and LDDMM on a higher resolution of  $192^2$ .



**Fig. 3.** Examples of registration results on 2D synthetic data. Left to right: 2D synthetic source and target images, deformed images, and transformation grids estimated from all algorithms.

The top panel of Fig. 4 displays an example of the comparison between the manually labeled segmentations and propagated segmentations on 3D brain MRI scans (at the resolution of  $192^3$ ) generated by our algorithm NeurEPDiff and the baseline models. The bottom panel of Fig. 4 reports the statistics of average dice scores over four brain structures from 100 registration pairs at different scales of image resolution. It shows that our method NeurEPDiff produces comparable dice scores. More importantly, the





**Fig. 4.** Quantitative result of average dice for all methods on 3D brain MRIs at different scales of resolution:  $64^3$ ,  $128^3$ , and  $192^3$ .


pre-trained NeurEPDiff at a lower resolution of  $64^3$  also achieves comparable results when tested on higher resolutions.

The left panel of Table 1 compares the runtime of solving the EPDiff equation with all methods on a CPU machine. It indicates that our method gains orders of magnitude speed compared to other baselines, while maintaining a comparable registration accuracy evaluated through the averaged dice (shown on the the right panel of Table 1). Overall, NeurEPDiff is approximately more than 10 times faster than FLASH and 100 times faster than LDDMM with geodesic shooting.

Figure 5 demonstrates the convergence graphs of optimized total energy by our developed NeurEPDiff on different levels of resolution. It shows that our method consistently converges to a better solution than LDDMM at multiple different resolutions. However, it achieves a slightly worse but fairly close solution to FLASH [36]. Please refer to Sect. 5 for more detailed discussions.

**Table 1.** Left to right: a runtime on images with different scales of resolutions and comparison of average dice on all brain structures.

Methods / Time	$64^3$	$128^3$	$192^3$	Methods/Dice	$64^3$	$128^3$	$192^3$
NeurEPDiff	<b>0.045 s</b>	<b>0.12 s</b>	<b>0.33 s</b>	NeurEPDiff	0.866	0.883	0.878
Flash	0.37 s	1.24 s	21.8 s	Flash	0.867	0.884	0.879
LDDMM	3.85 s	30.76 s	103.09 s	LDDMM	0.858	0.874	0.876

**Fig. 5.** A comparison of convergence graphs of all methods on 3D brain MRIs.

## 5 Discussion and Conclusion

We presented a novel neural operator, NeurEPDiff, to predict the geodesics of diffeomorphisms in a low-dimensional bandlimited space. Our method is the first to learn a solution operator to the geodesic evolution equation (EPDiff) with significantly improved computational efficiency. More importantly, the prediction capability of a pre-trained NeurEPDiff can be generalized to various resolutions of image grids because of its special property of resolution-invariance. To achieve this, we developed a brand-new network architecture, inspired by recent works on neural operators [15, 16], to approximate the nonlinear EPDiff equations as time-series functions carefully parameterized in the network hidden layers. A composition of network operators were defined in a complex-valued space to process the bandlimited signals of the velocity fields.

We validated NeurEPDiff on both simulated 2D and real 3D images. Experimental results show that our method achieves significantly faster speed at different scales of image resolution. Interestingly, we also find out that the NeurEPDiff-based registration algorithm converges to an optimal solution with slightly higher energy compared to FLASH [36]. This may be due to a relaxed regularization on the smoothness of the velocity fields parameterized in the network, which often happens in learning-based approaches. A possible solution is to enforce a smooth convolutional kernel  $\mathcal{H}$  [11] for each iterative evolution layer of our developed network architecture. This will be further investigated in our future work. Another potential future direction could be applying the network architectures developed in NeurEPDiff to learn another solution operator to predict transformations (Eq. (1)), which will further speed up optimization-based and learning-based diffeomorphic registration algorithms [6, 29, 30].

**Acknowledgement.** This work was supported by NSF CAREER Grant 2239977.

## 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. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Larsen, R., Nielsen, M., Sparring, J. (eds.) *MICCAI 2006*. LNCS, vol. 4190, pp. 924–931. Springer, Heidelberg (2006). <https://doi.org/10.1007/11866565-113>
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. *Med. Image Anal.* **12**(1), 26–41 (2008)
6. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V.: VoxelMorph: a learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* **38**, 1788–1800 (2019)
7. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vision* **61**(2), 139–157 (2005)
8. Cao, T., Singh, N., Jovic, V., Niethammer, M.: Semi-coupled dictionary learning for deformation prediction. In: *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pp. 691–694. IEEE (2015)
9. Christensen, G.E., Rabbitt, R.D., Miller, M.I.: Deformable templates using large deformation kinematics. *IEEE Trans. Image Process.* **5**(10), 1435–1447 (1996)
10. Dice, L.R.: Measures of the amount of ecologic association between species. *Ecology* **26**(3), 297–302 (1945)
11. Feinman, R., Lake, B.M.: Learning a smooth kernel regularizer for convolutional neural networks. arXiv preprint [arXiv:1903.01882](https://arxiv.org/abs/1903.01882) (2019)
12. Fotenos, A.F., Snyder, A., Girton, L., Morris, J., Buckner, R.: Normative estimates of cross-sectional and longitudinal brain volume decline in aging and ad. *Neurology* **64**(6), 1032–1039 (2005)
13. Hong, Y., Golland, P., Zhang, M.: Fast geodesic regression for population-based image analysis. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017*. LNCS, vol. 10433, pp. 317–325. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66182-7\\_37](https://doi.org/10.1007/978-3-319-66182-7_37)
14. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
15. Kovachki, N., et al.: Neural operator: learning maps between function spaces. arXiv preprint [arXiv:2108.08481](https://arxiv.org/abs/2108.08481) (2021)
16. Li, Z., Kovachki, N., Azzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., Anandkumar, A.: Fourier neural operator for parametric partial differential equations. arXiv preprint [arXiv:2010.08895](https://arxiv.org/abs/2010.08895) (2020)
17. Lu, L., Jin, P., Karniadakis, G.E.: DeepoNet: learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. arXiv preprint [arXiv:1910.03193](https://arxiv.org/abs/1910.03193) (2019)
18. Miller, M.I.: Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms. *Neuroimage* **23**, S19–S33 (2004)

19. Miller, M.I., Trounev, A., Younes, L.: Geodesic shooting for computational anatomy. *J. Math. Imaging Vis.* **24**(2), 209–228 (2006)
20. Niethammer, M., Huang, Y., Vialard, F.-X.: Geodesic regression for image time-series. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011. LNCS*, vol. 6892, pp. 655–662. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-23629-7\\_80](https://doi.org/10.1007/978-3-642-23629-7_80)
21. O’Donnell, L.J., Wells, W.M., Golby, A.J., Westin, C.-F.: Unbiased groupwise registration of white matter tractography. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012. LNCS*, vol. 7512, pp. 123–130. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-33454-2\\_16](https://doi.org/10.1007/978-3-642-33454-2_16)
22. 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)
23. Risser, L., Holm, D., Rueckert, D., Vialard, F.X.: Diffeomorphic atlas estimation using karcher mean and geodesic shooting on volumetric images. In: *MIUA* (2011)
24. Rohé, M.-M., Datar, M., Heimann, T., Sermesant, M., Pennec, X.: SVF-Net: learning deformable image registration using shape matching. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *MICCAI 2017. LNCS*, vol. 10433, pp. 266–274. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66182-7\\_31](https://doi.org/10.1007/978-3-319-66182-7_31)
25. Scardapane, S., Van Vaerenbergh, S., Hussain, A., Uncini, A.: Complex-valued neural networks with nonparametric activation functions. *IEEE Trans. Emerg. Top. Comput. Intell.* **4**(2), 140–150 (2018)
26. Schnabel, J.A., et al.: A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In: Niessen, W.J., Viergever, M.A. (eds.) *MICCAI 2001. LNCS*, vol. 2208, pp. 573–581. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-45468-3\\_69](https://doi.org/10.1007/3-540-45468-3_69)
27. Shi, W., et al.: Registration using sparse free-form deformations. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012. LNCS*, vol. 7511, pp. 659–666. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-33418-4\\_81](https://doi.org/10.1007/978-3-642-33418-4_81)
28. Singh, N., Hinkle, J., Joshi, S., Fletcher, P.T.: A hierarchical geodesic model for diffeomorphic longitudinal shape analysis. In: Gee, J.C., Joshi, S., Pohl, K.M., Wells, W.M., Zöllei, L. (eds.) *IPMI 2013. LNCS*, vol. 7917, pp. 560–571. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-38868-2\\_47](https://doi.org/10.1007/978-3-642-38868-2_47)
29. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *Int. J. Comput. Vision* **97**(2), 229–241 (2012)
30. Wang, J., Zhang, M.: DeepFlash: an efficient network for learning-based medical image registration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
31. Wang, J., Zhang, M.: Geo-sic: learning deformable geometric shapes in deep image classifiers. In: *The Conference on Neural Information Processing Systems* (2022)
32. Wells, W., Viola, P., Atsumi, H., Nakajima, S., Kikinis, R.: Multi-modal volume registration by maximization of mutual information. *Med. Image Anal.* **1**, 35–51 (1996)
33. Yang, X., Kwitt, R., Styner, M., Niethammer, M.: Quicksilver: fast predictive image registration—a deep learning approach. *Neuroimage* **158**, 378–396 (2017)
34. Younes, L., Arrate, F., Miller, M.I.: Evolutions equations in computational anatomy. *Neuroimage* **45**(1), S40–S50 (2009)
35. Zhang, M., Fletcher, P.T.: Finite-dimensional lie algebras for fast diffeomorphic image registration. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) *IPMI 2015. LNCS*, vol. 9123, pp. 249–260. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-19992-4\\_19](https://doi.org/10.1007/978-3-319-19992-4_19)

36. Zhang, M., et al.: Frequency diffeomorphisms for efficient image registration. In: Niethammer, M., et al. (eds.) IPMI 2017. LNCS, vol. 10265, pp. 559–570. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-59050-9\\_44](https://doi.org/10.1007/978-3-319-59050-9_44)
37. Zhang, M., Wells, W.M., Golland, P.: Low-dimensional statistics of anatomical variability via compact representation of image deformations. In: Ourselin, S., Joskowicz, L., Sabuncu, M.R., Unal, G., Wells, W. (eds.) MICCAI 2016. LNCS, vol. 9902, pp. 166–173. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-46726-9\\_20](https://doi.org/10.1007/978-3-319-46726-9_20)