

Two-Parameter Flow Map Learning for Continuous-Time Diffeomorphic Image Registration

Mohammadjavad Matinkia[✉] and Nilanjan Ray[✉]

University of Alberta, Canada
matinkia@ualberta.ca, nray1@ualberta.ca

Abstract. Diffeomorphic image registration is central to medical image analysis, enabling anatomically consistent alignment across subjects. Most learning-based diffeomorphic methods model autonomous ODEs (ordinary differential equations) by parameterizing a stationary velocity field and recovering deformations via scaling-and-squaring. While non-autonomous ODEs with time-dependent velocities increase expressiveness, existing approaches rely on numerical integration to implicitly enforce flow structure that entangles model expressiveness with discretization accuracy. We propose a framework to directly learn the continuous-time solution of a non-autonomous ODE formulated as a two-parameter flow map. By enforcing cocycle consistency, a fundamental structural property of time-varying flows, we learn the flow maps without time discretization and velocity integration during training. The framework recovers diffeomorphic mappings at inference using a small number of compositions. Our proposed framework seamlessly incorporates standard registration backbones and improves alignment accuracy consistently across nine datasets while preserving diffeomorphic structure. Notably, the proposed method achieves an average Dice improvement of 2.1% on brain MRI benchmarks, a 12% TRE reduction on lung CT, and a 2.6% Dice gain on cardiac MRI and ultrasound datasets (<https://mattkia.github.io/TPFMDIR/>).

Keywords: Diffeomorphic Image Registration · Medical Image Registration · Flow-Based Models · Non-autonomous ODEs

1 Introduction

Diffeomorphic image registration (DIR) aims to estimate a smooth, bijective transformation that warps a moving image onto a fixed image while preserving topology. Unlike general deformable registration methods, which prioritize image similarity under smoothness constraints, diffeomorphic approaches explicitly enforce smooth invertibility to prevent foldings and ensure anatomically plausible alignment. Although this constraint provides strong geometric guarantees,




Fig. 1: Overview of TPFM-DIR. The left panel illustrates time-dependent similarity supervision, while the right panel depicts cocycle consistency regularization. Top: forward-time training; bottom: backward-time training.

diffeomorphic methods have historically faced a trade-off between anatomical consistency and alignment accuracy [14].

Traditional DIR methods rely on variational optimization to compute a deformation for each image pair, resulting in high computational cost [2, 6, 29]. Neural approaches address this limitation by either optimizing network parameters per image pair or learning deformation priors across datasets for rapid inference [4, 13, 46]. A particularly influential class of neural DIR methods models deformations as flows generated by ordinary differential equations (ODEs) [5, 8, 30]. In this formulation, the transformation is obtained by integrating a velocity field over time. Autonomous ODEs, being based on stationary velocity fields, allow efficient integration via scaling-and-squaring [1]. Non-autonomous formulations generalize this setting by allowing time-dependent velocity fields. However, they require costly numerical integration during both training and inference to induce flow structure. As a result, deformation consistency becomes coupled to discretization schemes and solver accuracy [43].

We introduce **TPFM-DIR**, a flow-based framework, which directly models the continuous-time solution operator of a non-autonomous ODE. By doing so, the method eliminates explicit velocity parameterization and numerical integration during training. Our approach is grounded in the theoretical properties of non-autonomous ODE solutions. In particular, we exploit the **cocycle property**, which characterizes valid flow maps of non-autonomous systems [23]. We propose a cocycle-based regularization for flow maps. This regularization enforces temporal consistency, and promotes invertibility and regularity of the deformation. As a result, TPFM-DIR does not require explicit integration and




Fig. 2: Architecture of TPFM-DIR (left), training scheme (top right), and inference scheme (bottom right). The training diagram illustrates the forward direction; the backward direction follows the formulation shown in Fig. 1. In all experiments, we use $N = 4$ compositions during inference.

multiple handcrafted smoothness penalties that are commonly used in diffeomorphic learning frameworks.

By modeling the flow map directly, the framework captures complex time-dependent deformations. At inference time, global transformations are constructed through a small number of compositions of predicted short-time maps. This strategy improves deformation regularity while maintaining computational efficiency.

Empirically, TPFM-DIR narrows the performance gap between diffeomorphic and deformable registration, matching or surpassing strong deformable baselines while preserving topology. We further demonstrate that TPFM-DIR integrates with diverse backbone architectures, consistently improving both alignment accuracy and topology preservation. The contributions of this work are summarized as follows:

1. We introduce TPFM-DIR, a flow-map-based formulation of diffeomorphic image registration, that directly models the two-parameter solution of a non-autonomous ODE. Unlike velocity-based parameterizations that learn the generating field, our formulation operates at the level of the flow map itself, enabling structural modeling of the time-dependent deformations.
2. We propose a cocycle-based regularization with formal guarantees that characterize valid solutions of non-autonomous ODEs. This removes the need for explicit numerical integration during training while promoting temporal consistency and invertibility without relying on multiple handcrafted smoothness penalties.
3. We validate TPFM-DIR on nine benchmark datasets spanning 2D and 3D MRI, CT, and ultrasound, demonstrating consistent improvements over strong deformable and diffeomorphic methods.




Fig. 3: Continuous registration trajectory on CANDI. Top: warped image evolution with overlaid velocity norms computed by Eq. (9) indicating time-varying velocity. Bottom: incremental image differences illustrating temporal smoothness.

2 Related Works

Deformable vs. Diffeomorphic Registration. Deformable registration methods typically predict dense displacement fields and prioritize alignment accuracy. Recent performance gains are largely driven by increasingly expressive architectures, including vision transformers [8, 9, 17] and correlation or structure aware modules [10, 27]. However, such models generally lack explicit structural constraints, which may result in folding under large deformations. To ensure anatomical plausibility, Diffeomorphic Image Registration (DIR) seeks smooth, invertible transformations [38, 42]. Classical geodesic formulations (e.g., LDDMM [6], SyN [2]) provide strong theoretical guarantees but at high computational cost. Modern learning-based DIR approaches amortize this optimization across datasets [4, 13, 32, 36, 46], typically enforcing regularity through architectural symmetry [30] or surrogate penalties such as cycle, inverse, and gradient consistency [11, 12, 21, 39, 40]. These constraints may trade registration accuracy for regularity [42] and are usually defined only between image pairs.

Flow-Based Diffeomorphic Models. Flow-based DIR represents transformations as ODE-generated flows [3, 15, 28]. Autonomous formulations assume stationary velocity fields and enable efficient scaling-and-squaring integration to recover the deformations [5, 7, 30, 31, 37]. However, the stationarity assumption restricts the admissible deformation dynamics, whereas non-autonomous models allow time-varying velocities and provide a more general deformation framework. Classical LDDMM [6] and neural extensions such as NODEO [43] and R2Net [18] fall into this category. The increased expressiveness of non-autonomous flows comes at both computational and structural complexity. Existing approaches rely on explicit numerical integration [43, 44], coupling diffeomorphic guarantees to discretization accuracy and solver design. In contrast, our approach directly models the two-parameter flow family, eliminating explicit velocity integration while preserving the structural foundations of flow-based diffeomorphisms.




Fig. 4: Visual comparisons on OASIS and LungCT. Dice/TRE are shown at top-left corners and $|J|_{<0}\%$ is shown at top-right corners.

3 Proposed Method

3.1 Preliminaries

Let $I_m, I_f : \Omega \rightarrow \mathbb{R}$ denote the moving and fixed images defined over a d -dimensional spatial domain $\Omega \subset \mathbb{R}^d$. Diffeomorphic registration seeks a transformation $\phi : \Omega \rightarrow \Omega$ that balances image similarity and transformation regularity. To obtain diffeomorphic transformations, many methods model ϕ as the flow of an ODE under suitable regularity conditions. A common approach utilizes an *autonomous* ODE [8, 30]

$$\frac{d\phi_t}{dt} = v(\phi_t), \quad \phi_0 = Id, \quad (1)$$

where v is a stationary velocity field, Id denotes the identity mapping, and ϕ_t is the deformation at time t . This formulation is favored for its computational efficiency via the scaling-and-squaring integration scheme [1]. On the other hand, non-autonomous ODEs generalize the autonomous case by adopting time-dependent velocity fields [6, 43]:

$$\frac{d\phi_t}{dt} = v(\phi_t, t), \quad \phi_0 = Id. \quad (2)$$

While more expressive, non-autonomous ODEs rely on numerical integration during training, coupling deformation structure to time discretization and solver accuracy.

In this work, we shift the focus from the velocity field v to the direct modeling of the **two-parameter flow map** $\phi_{s,t} : \Omega \rightarrow \Omega$, which represents the transformation from time s to time t and encodes the solution of a non-autonomous ODE, eliminating the need for discretized velocity integration to construct the flow. For a differentiable mapping $\phi_{s,t}$ to be a valid solution of a non-autonomous ODE, it must satisfy the identity condition and the **cocycle property** [23]:

$$\phi_{s,s} = Id, \quad \phi_{r,t} \circ \phi_{s,r} = \phi_{s,t}, \quad \forall s \leq r \leq t \in [0, 1]. \quad (3)$$

The cocycle property dictates that the mapping from time s to t is equivalent to the composition of intermediate mappings through any intermediate time r . Leveraging this structural constraint, we learn $\phi_{s,t}$ directly without explicit velocity parameterization, decoupling flow modeling from velocity integration. In the rest of the paper, the moving image is associated with time 0, and the fixed image with time 1.

3.2 Two-Parameter Flow Maps For Diffeomorphic Image Registration

The core objective of TPFM-DIR is to directly model the flow map associated with the non-autonomous ODE in Eq. (2), represented as a two-parameter family of transformations. We parameterize the flow map $\phi_{s,t}$ as

$$\phi_{s,t} = Id + (t - s)f_\theta(s, t; I_m, I_f), \quad (4)$$

where $f_\theta(\cdot)$ is a time-dependent neural network with parameters θ , conditioned on the moving and fixed images and time parameters s and t . This construction enforces $\phi_{s,s} = Id$ and defines all transformations in a shared spatial frame, enabling valid compositions. This parameterization alone does not guarantee a valid flow map; the cocycle regularization imposes the necessary structural constraints to promote consistency with a continuous-time flow. We implement f_θ using a doubly time-embedded U-Net (Fig. 2) with sinusoidal embeddings for time samples s and t [16, 25]. Smoothness and invertibility are further promoted through the shared-coordinate formulation and cocycle regularization.

To train the model without numerical integration, we propose a temporal similarity loss and an *anchored* cocycle regularization with fixed endpoints. Following the conventional approaches of time-dependent architectures, we sample $s, t \sim \text{Uni}[0, 1]$ with $s \leq t$ [16]. To ensure unbiased learning and memory efficiency, we employ a symmetric training scheme that alternates which branch carries the composed transformation at each epoch. This avoids doubled memory usage while preserving symmetry. We define the similarity loss as

$$\mathcal{L}_{\text{sim}}^{s,t} = \begin{cases} -NCC(I_m(\phi_{s,t} \circ \phi_{0,s}), I_f(\phi_{1,t})) & \text{even epochs} \\ -NCC(I_m(\phi_{0,s}), I_f(\phi_{t,s} \circ \phi_{1,t})) & \text{odd epochs} \end{cases} \quad (5)$$

This aligns the forward trajectory of the moving image with the backward trajectory of the fixed image at time t . To enforce the structural consistency of the flow, we introduce the *anchored* cocycle regularization:

$$\mathcal{L}_{\text{cocycle}}^{s,t} = \begin{cases} \|\phi_{s,t} \circ \phi_{0,s} - \phi_{0,t}\|_2 & \text{even epochs} \\ \|\phi_{t,s} \circ \phi_{1,t} - \phi_{1,s}\|_2 & \text{odd epochs} \end{cases} \quad (6)$$

Figure 1 illustrates the training process. In practice, compositions are implemented via spatial warping with linear interpolation. Minimizing Eq. (6) enforces the anchored cocycle constraints:

$$\phi_{s,t} \circ \phi_{0,s} = \phi_{0,t}, \quad \phi_{t,s} \circ \phi_{1,t} = \phi_{1,s} \quad (7)$$

Hence, the total loss becomes

$$\mathcal{L}^{s,t} = \mathcal{L}_{\text{sim}}^{s,t} + \lambda \mathcal{L}_{\text{cocycle}}^{s,t} \quad (8)$$

where λ balances the trajectory consistency. In our experiments we show that a value $\lambda = 10$ leads to the most promising results. Note that enforcing the full cocycle constraint for arbitrary triplets (r, s, t) would require evaluating multiple nested compositions per iteration and is substantially more expensive. We therefore enforce only the anchored cocycle constraints in Eq. (7), which are computationally more efficient while remaining sufficient in practice:

Theorem 1. *Under standard regularity assumptions, a deformation $\phi_{s,t}$ satisfying the identity condition and the anchored cocycle properties of Eq. (7) is a two-parameter diffeomorphism solving a non-autonomous ODE of Eq. (2) on a bounded open domain Ω .*

Proof. Appendix A.

An immediate consequence of Eq. (4) is that the instantaneous velocity admits a closed-form expression without numerical differentiation:

Corollary 1. *For the flow map $\phi_{s,t}$ parameterized as Eq. (4), the instantaneous velocity can be analytically obtained via*

$$v_t = f_\theta(t, t; I_m, I_f). \quad (9)$$

Proof. Appendix A.

This allows direct extraction of the velocity fields for further investigations (see Fig. 3). The complete training and inference schema can be viewed in Figs. 1 and 2.

Uniform sampling of s and t biases the time-difference distribution toward small increments, as shorter intervals occur more frequently under uniform sampling. Consequently, large transitions are relatively underrepresented during optimization, which may lead to mild deviations from ideal flow structure when directly predicting $\phi_{0,1}$ (see Appendix B for a complete discussion). Thus, although $\phi_{0,1}$ can be predicted in a single pass at inference, we partition the interval into N subintervals and compose locally predicted maps to better approximate a continuous-time evolution at inference:

$$\phi_{0,1} = \phi_{t_N,1} \circ \cdots \circ \phi_{0,t_1}, \quad (10)$$

Each segment corresponds to a well-trained local deformation, improving stability and reducing the risk of foldings. In our experiments, we find that a $N = 4$ compositions suffices for constructing an accurate diffeomorphic deformation.

Table 1: Quantitative comparison on the OASIS and IXI datasets. The best result per metric is shown in **bold**. The second-best diffeomorphic and deformable results are highlighted in **blue** and **red**, respectively.

Method	OASIS				IXI			
	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	SDLogJ $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	SDLogJ $\downarrow$
LDDMM	76.59 $\pm$ 2.42	0.0064 $\pm$ 0.0051	3.89 $\pm$ 0.93	0.012 $\pm$ 0.009	67.90 $\pm$ 1.15	0.0034 $\pm$ 0.0021	4.11 $\pm$ 1.02	0.058 $\pm$ 0.020
CycleMorph	81.93 $\pm$ 2.14	0.0211 $\pm$ 0.0091	2.36 $\pm$ 0.81	0.040 $\pm$ 0.026	74.67 $\pm$ 1.33	0.0224 $\pm$ 0.0110	3.34 $\pm$ 0.87	0.068 $\pm$ 0.032
GradICON	83.74 $\pm$ 1.42	0.0039 $\pm$ 0.0012	2.09 $\pm$ 0.36	0.011 $\pm$ 0.006	76.43 $\pm$ 1.28	0.0018 $\pm$ 0.0022	3.19 $\pm$ 0.55	0.015 $\pm$ 0.006
TransMorph-diff	83.51 $\pm$ 1.52	0.0066 $\pm$ 0.0073	2.35 $\pm$ 0.76	0.071 $\pm$ 0.041	75.98 $\pm$ 2.01	0.0142 $\pm$ 0.0098	3.81 $\pm$ 0.89	0.102 $\pm$ 0.048
R2Net	80.11 $\pm$ 1.53	0.0093 $\pm$ 0.0054	2.15 $\pm$ 0.49	0.081 $\pm$ 0.055	73.39 $\pm$ 1.77	0.0031 $\pm$ 0.0016	3.69 $\pm$ 0.66	0.042 $\pm$ 0.026
NODEO	81.97 $\pm$ 1.39	0.0024 $\pm$ 0.0010	2.18 $\pm$ 0.71	0.008 $\pm$ 0.005	75.68 $\pm$ 1.82	0.0031 $\pm$ 0.0015	3.46 $\pm$ 0.84	0.017 $\pm$ 0.008
NePhi	80.09 $\pm$ 1.17	0.0056 $\pm$ 0.0023	3.28 $\pm$ 0.44	0.022 $\pm$ 0.016	75.34 $\pm$ 1.78	0.0029 $\pm$ 0.0016	3.55 $\pm$ 0.55	0.018 $\pm$ 0.011
PULPo	80.51 $\pm$ 0.99	0.0527 $\pm$ 0.0240	3.45 $\pm$ 0.98	0.077 $\pm$ 0.031	73.91 $\pm$ 2.11	0.0066 $\pm$ 0.0043	3.30 $\pm$ 0.39	0.028 $\pm$ 0.016
TransMorph	84.11 $\pm$ 1.30	1.0665 $\pm$ 0.5631	2.26 $\pm$ 0.68	0.821 $\pm$ 0.252	77.72 $\pm$ 1.76	1.2574 $\pm$ 0.7574	3.71 $\pm$ 0.80	1.083 $\pm$ 0.742
TransMatch	83.48 $\pm$ 1.80	0.1493 $\pm$ 0.0151	2.23 $\pm$ 0.66	0.741 $\pm$ 0.344	76.80 $\pm$ 1.85	0.0406 $\pm$ 0.0120	3.15 $\pm$ 0.41	0.580 $\pm$ 0.411
DiffuseMorph	79.13 $\pm$ 2.03	1.1563 $\pm$ 0.2512	3.51 $\pm$ 0.77	0.336 $\pm$ 0.189	71.35 $\pm$ 2.05	1.6042 $\pm$ 0.7212	5.48 $\pm$ 1.45	0.938 $\pm$ 0.508
DiffuseReg	80.12 $\pm$ 1.22	1.0877 $\pm$ 0.4312	3.73 $\pm$ 1.01	0.377 $\pm$ 0.130	72.13 $\pm$ 1.88	0.5279 $\pm$ 0.2834	5.04 $\pm$ 1.12	0.488 $\pm$ 0.276
HVT	85.07 $\pm$ 1.05	0.4812 $\pm$ 0.0114	1.92 $\pm$ 0.46	0.631 $\pm$ 0.412	80.67 $\pm$ 1.67	0.5933 $\pm$ 0.1028	2.98 $\pm$ 0.46	0.679 $\pm$ 0.290
CorrMLP	84.66 $\pm$ 1.24	0.4640 $\pm$ 0.1776	2.23 $\pm$ 0.62	0.608 $\pm$ 0.278	77.54 $\pm$ 1.68	0.3675 $\pm$ 0.2168	3.15 $\pm$ 0.39	0.547 $\pm$ 0.231
SACB-Net	83.17 $\pm$ 1.18	0.0184 $\pm$ 0.0093	2.33 $\pm$ 0.48	0.182 $\pm$ 0.087	78.15 $\pm$ 1.87	0.0201 $\pm$ 0.0099	3.19 $\pm$ 0.51	0.144 $\pm$ 0.056
DGIR	82.37 $\pm$ 1.33	0.0325 $\pm$ 0.0076	3.29 $\pm$ 0.91	0.392 $\pm$ 0.182	75.34 $\pm$ 1.55	0.0766 $\pm$ 0.0101	3.22 $\pm$ 0.72	0.193 $\pm$ 0.101
TPFM-DIR	87.08 $\pm$ 1.07	0.0019 $\pm$ 0.0007	1.78 $\pm$ 0.44	0.006 $\pm$ 0.002	82.52 $\pm$ 0.98	0.0015 $\pm$ 0.0006	2.95 $\pm$ 0.68	0.007 $\pm$ 0.003

4 Experimental Results

Unless otherwise specified, all experiments use $\lambda = 10$ (Eq. (8)) and inference uses 4 compositions (Eq. (10)). TPFM-DIR is optimized with AdamW with learning rate 10^{-4} and batch size of 1. Training is performed for 200 epochs on all datasets except LungCT (100 epochs). We set a windows size of 7 for the NCC loss across all experiments. All experiments are conducted on a single NVIDIA RTX 3090 GPU (24GB VRAM). Additional implementation details are provided in Appendix D.

4.1 Datasets and Metrics

We evaluate TPFM-DIR on nine datasets spanning multiple modalities, anatomical regions, and spatial dimensions. Brain MRI benchmarks include OASIS [26], IXI¹, Mindboggle101 [22], LPBA40 [35], and CANDI [19]. For thoracic and abdominal CT, we use LungCT and AbdomenCT from the Learn2Reg challenge², representing large respiratory motion and multi-organ abdominal alignment, respectively. To assess 2D registration, we use ACDC [34] (cardiac MRI) and CAMUS [24] (cardiac ultrasound) datasets.

For datasets with segmentation masks, we report the Dice similarity coefficient as the primary overlap metric. Surface alignment is evaluated using the 95th percentile Hausdorff Distance (HD95) and the Average Symmetric Surface Distance (ASSD). For LungCT, where expert landmarks are available, we compute Target Registration Error (TRE), and additionally report Structural Similarity Index (SSIM) to assess intensity consistency. Diffeomorphic and smoothness

¹ <https://brain-development.org/ixi-dataset>

² <https://learn2reg.grand-challenge.org/>




Fig. 5: Qualitative comparisons on ACDC and CAMUS. Warped images with overlaid segmentations are shown with Dice reported at the top-left corners; $|J|_{<0}\%$ and SDLogJ are shown at the bottom-left and bottom-right of the deformation grids, respectively.

properties are evaluated via percentage of voxels with negative Jacobian determinant ($|J|_{<0}\%$) and the standard deviation of the logarithm of the Jacobian determinant (SDLogJ). To enhance readability of the tables, we only report the main metrics for each dataset, and present the tables with complete metrics in Appendix F.

Our evaluations include both inter-subject (brain MRI and AbdomenCT) and intra-subject (LungCT and cardiac sequences) registration. For inter-subject experiments, subject pairs are randomly sampled following a fixed protocol. Complete dataset splits and preprocessing are detailed in Appendix C.

4.2 Results

Across all datasets, we compare TPFM-DIR against a broad range of state-of-the-art methods. The diffeomorphic baselines include CycleMorph, GradICON, TransMorph-diff, R2Net, NODEO, NePhi, and PULPo [8, 18, 21, 37, 39, 40, 43]. Deformable baselines include TransMorph, TransMatch, DiffuseMorph, DiffuseReg, HViT, CorrMLP, SACB-Net, and DGIR [8–10, 17, 20, 27, 41, 47]. Among the diffeomorphic methods, R2Net and NODEO adopt non-autonomous ODE formulations. Moreover, we include comparisons with LDDMM [6] as a classical diffeomorphic method based on non-autonomous ODEs. For 2D datasets, we evaluate only methods with official 2D implementations and additionally include AdaCS as a strong baseline for ACDC and CAMUS [45].

All quantitative results are reported as mean $\pm$ standard deviation. Tables 1 and 2 present the results on the neuroimaging 3D MRI datasets. TPFM-DIR consistently achieves the highest Dice scores while maintaining near-zero folding

Table 2: Quantitative comparison on the LPBA40, Mindboggle101, and CANDI datasets. The best result per metric is shown in **bold**. The second-best diffeomorphic and deformable results are highlighted in **blue** and **red**, respectively.

Method	LPBA40			Mindboggle101			CANDI		
	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$
LDDMM	69.23 $\pm$ 1.78	0.0011 $\pm$ 0.0004	7.38 $\pm$ 1.11	64.57 $\pm$ 2.08	0.0061 $\pm$ 0.0039	7.79 $\pm$ 0.84	78.52 $\pm$ 1.66	0.0017 $\pm$ 0.0007	2.44 $\pm$ 0.31
CycleMorph	71.90 $\pm$ 1.41	0.0064 $\pm$ 0.0032	6.76 $\pm$ 0.94	70.64 $\pm$ 1.48	0.1282 $\pm$ 0.0873	7.75 $\pm$ 0.98	81.37 $\pm$ 1.25	0.0027 $\pm$ 0.0013	2.23 $\pm$ 0.29
GradICON	74.21 $\pm$ 1.26	0.0011 $\pm$ 0.0006	6.89 $\pm$ 0.89	70.78 $\pm$ 1.52	0.0117 $\pm$ 0.0043	7.59 $\pm$ 1.12	83.39 $\pm$ 1.09	0.0015 $\pm$ 0.0008	2.14 $\pm$ 0.27
TransMorph-diff	70.98 $\pm$ 1.52	0.0086 $\pm$ 0.0044	7.22 $\pm$ 1.03	69.45 $\pm$ 1.66	0.0364 $\pm$ 0.0157	7.89 $\pm$ 1.05	83.20 $\pm$ 1.14	0.0026 $\pm$ 0.0012	2.06 $\pm$ 0.26
R2Net	70.71 $\pm$ 1.49	0.0017 $\pm$ 0.0009	7.13 $\pm$ 1.01	68.34 $\pm$ 1.69	0.0086 $\pm$ 0.0031	8.97 $\pm$ 1.44	78.79 $\pm$ 1.49	0.0032 $\pm$ 0.0016	2.78 $\pm$ 0.32
NODEO	72.73 $\pm$ 1.25	0.0008 $\pm$ 0.0004	7.25 $\pm$ 0.97	70.18 $\pm$ 1.58	0.0023 $\pm$ 0.0011	7.70 $\pm$ 0.96	80.60 $\pm$ 1.26	0.0012 $\pm$ 0.0007	2.46 $\pm$ 0.29
NePhi	72.39 $\pm$ 1.33	0.0024 $\pm$ 0.0016	6.75 $\pm$ 0.91	67.97 $\pm$ 1.91	0.0019 $\pm$ 0.0005	8.20 $\pm$ 1.23	79.15 $\pm$ 1.46	0.0009 $\pm$ 0.0005	2.36 $\pm$ 0.30
PULPo	74.82 $\pm$ 1.11	0.0009 $\pm$ 0.0005	6.22 $\pm$ 0.78	68.69 $\pm$ 1.74	0.0287 $\pm$ 0.0104	7.92 $\pm$ 1.07	78.38 $\pm$ 1.57	0.0136 $\pm$ 0.0047	2.89 $\pm$ 0.33
TransMorph	74.42 $\pm$ 1.33	0.5275 $\pm$ 0.1835	6.88 $\pm$ 0.85	71.46 $\pm$ 1.63	1.4883 $\pm$ 0.5328	7.27 $\pm$ 0.92	83.48 $\pm$ 1.10	0.1059 $\pm$ 0.0388	2.13 $\pm$ 0.22
TransMatch	74.13 $\pm$ 1.31	0.0251 $\pm$ 0.0116	6.53 $\pm$ 0.83	71.78 $\pm$ 1.49	0.2483 $\pm$ 0.1020	7.63 $\pm$ 1.08	82.60 $\pm$ 1.18	0.0160 $\pm$ 0.0061	2.28 $\pm$ 0.25
DiffuseMorph	73.11 $\pm$ 1.42	0.3610 $\pm$ 0.1429	7.21 $\pm$ 1.06	68.22 $\pm$ 1.77	1.7413 $\pm$ 0.6521	8.43 $\pm$ 1.30	77.22 $\pm$ 1.51	0.1277 $\pm$ 0.0412	2.53 $\pm$ 0.30
DiffuseReg	72.98 $\pm$ 1.46	0.0723 $\pm$ 0.0289	7.27 $\pm$ 1.04	71.88 $\pm$ 1.54	1.0823 $\pm$ 0.4882	7.84 $\pm$ 1.03	79.31 $\pm$ 1.38	0.0322 $\pm$ 0.0125	2.31 $\pm$ 0.26
HViT	76.57 $\pm$ 1.18	0.4428 $\pm$ 0.1655	6.74 $\pm$ 0.89	71.80 $\pm$ 1.72	1.4051 $\pm$ 0.6153	7.56 $\pm$ 1.11	83.67 $\pm$ 1.07	0.1368 $\pm$ 0.0435	2.09 $\pm$ 0.21
CorrMLP	75.72 $\pm$ 1.27	0.0145 $\pm$ 0.0068	6.55 $\pm$ 0.87	71.92 $\pm$ 1.68	0.1895 $\pm$ 0.0782	7.65 $\pm$ 1.06	81.96 $\pm$ 1.25	0.0098 $\pm$ 0.0043	2.15 $\pm$ 0.26
SACB-Net	73.91 $\pm$ 1.35	0.0099 $\pm$ 0.0041	6.70 $\pm$ 0.93	72.75 $\pm$ 1.57	0.3542 $\pm$ 0.1433	7.46 $\pm$ 0.97	82.75 $\pm$ 1.22	0.0047 $\pm$ 0.0021	2.19 $\pm$ 0.27
DGIR	71.99 $\pm$ 1.44	0.0189 $\pm$ 0.0073	7.01 $\pm$ 0.89	68.12 $\pm$ 1.93	0.0674 $\pm$ 0.0213	7.94 $\pm$ 1.08	80.13 $\pm$ 1.39	0.0077 $\pm$ 0.0028	2.74 $\pm$ 0.41
TPFM-DIR	77.87 $\pm$ 1.09	0.0005 $\pm$ 0.0002	5.85 $\pm$ 1.63	74.79 $\pm$ 0.98	0.0012 $\pm$ 0.0017	7.53 $\pm$ 1.86	84.77 $\pm$ 1.37	0.0001 $\pm$ 0.0001	1.95 $\pm$ 0.49




Fig. 6: Qualitative comparisons on the AbdomenCT dataset. The segmentation masks are overlaid on all images. Dice score and $|J|_{<0}\%$ of each result are reported at the lower-left and lower-right corners of each sample.

ratios ($|J|_{<0}\%$). On average, Dice improves by 2.1% over the next best method. Visual comparisons on OASIS are shown in Fig. 4 (top).

Table 3 reports results on LungCT and AbdomenCT, where TPFM-DIR achieves zero folding on LungCT and reduces TRE by 17% and 12% relative to GradICON and CorrMLP, respectively. Figure 4 (bottom) and Fig. 6 provide qualitative comparisons for LungCT and AbdomenCT, respectively. Finally, Table 4 summarizes results on ACDC and CAMUS. On CAMUS dataset TPFM-DIR improves Dice by 3% while preserving topology with zero folding. Visual comparisons and deformation grids are shown in Fig. 5. More visual results can be found in Appendix G.

Table 3: Quantitative comparison on the LungCT and AbdomenCT datasets. The best result per metric is shown in **bold**. The second-best diffeomorphic and deformable results are highlighted in **blue** and **red**, respectively.

Method	LungCT			AbdomenCT			
	TRE ↓	$ J _{<0}\%$ ↓	SSIM ↑	Dice ↑	$ J _{<0}\%$ ↓	HD95 ↓	SSIM ↑
LDDMM	3.09 ± 0.26	0.0033 ± 0.0015	61.14 ± 3.11	37.36 ± 2.71	0.0112 ± 0.0072	12.61 ± 1.46	53.83 ± 3.17
CycleMorph	3.04 ± 0.19	0.7180 ± 0.2117	56.99 ± 3.44	42.56 ± 2.65	0.0053 ± 0.0025	13.04 ± 1.28	52.68 ± 3.02
GradICON	2.64 ± 0.17	0.0009 ± 0.0004	54.06 ± 2.86	44.02 ± 2.81	0.0009 ± 0.0004	12.21 ± 1.33	54.86 ± 3.27
TransMorph-diff	2.89 ± 0.18	0.0042 ± 0.0015	66.85 ± 2.95	41.41 ± 2.72	0.0034 ± 0.0017	12.39 ± 1.25	62.61 ± 2.98
R2Net	2.99 ± 0.19	0.0211 ± 0.0068	68.06 ± 2.81	41.99 ± 2.89	0.0010 ± 0.0005	13.05 ± 1.32	51.52 ± 3.36
NODEO	2.87 ± 0.18	0.0038 ± 0.0013	70.80 ± 2.74	40.37 ± 2.96	0.0007 ± 0.0004	13.16 ± 1.38	52.45 ± 3.42
NePhi	3.12 ± 0.21	0.0093 ± 0.0027	56.66 ± 3.33	45.32 ± 2.24	0.0008 ± 0.0003	12.48 ± 1.19	66.34 ± 2.54
TransMorph	2.85 ± 0.17	0.4767 ± 0.1546	68.82 ± 2.65	46.66 ± 2.48	3.1310 ± 0.8113	12.61 ± 1.28	71.84 ± 2.54
TransMatch	3.02 ± 0.19	0.1343 ± 0.0482	67.47 ± 2.89	43.64 ± 2.35	2.6378 ± 0.6922	12.50 ± 1.22	71.47 ± 2.33
DiffuseMorph	3.38 ± 0.21	0.0794 ± 0.0195	62.98 ± 2.92	39.50 ± 2.81	0.0292 ± 0.0113	12.88 ± 1.34	53.02 ± 3.09
DiffuseReg	3.26 ± 0.20	0.7765 ± 0.2337	65.75 ± 2.71	45.91 ± 2.56	1.8106 ± 0.7348	12.55 ± 1.25	56.87 ± 2.88
HViT	3.22 ± 0.20	0.1791 ± 0.0554	66.72 ± 3.41	47.66 ± 2.41	0.9601 ± 0.4175	10.11 ± 1.10	66.90 ± 2.74
CorrMLP	2.48 ± 0.16	0.2673 ± 0.0711	65.57 ± 3.08	50.28 ± 2.18	0.1656 ± 0.0643	11.47 ± 1.16	62.66 ± 2.59
SACB-Net	3.01 ± 0.19	0.9824 ± 0.2761	70.06 ± 2.77	53.38 ± 2.63	0.9348 ± 0.1135	13.09 ± 1.41	66.73 ± 2.96
DGIR	3.24 ± 0.42	0.0711 ± 0.0428	67.72 ± 2.41	42.08 ± 2.23	0.0277 ± 0.0114	12.24 ± 1.80	64.48 ± 2.71
TPFM-DIR	2.19 ± 0.14	0.0 ± 0.0	71.18 ± 2.35	54.54 ± 1.86	0.0002 ± 0.0001	9.81 ± 1.13	68.94 ± 1.77

5 Ablation Studies

We systematically analyze key design components of TPFM-DIR. Specifically, we study: (i) the influence of the cocycle regularization through the weight λ in Eq. (8) on accuracy and topology preservation, (ii) the effect of progressive deformation compositions at inference, (iii) the architectural flexibility of the proposed framework, and (iv) the computational analysis of TPFM-DIR.

5.1 Effect of Regularization Factor

The regularization weight λ controls the trade-off between alignment accuracy and topology preservation. The value $\lambda = 10$ is fixed across all datasets and used in all main experiments since it consistently provided a favorable trade-off between alignment accuracy and topology preservation across datasets. We evaluate the impact of varying λ on three representative datasets. Table 5 shows that increasing λ progressively eliminates voxels with negative Jacobian determinant, at the cost of a moderate decrease in the performance metric. Conversely, reducing λ does not improve alignment accuracy, while significantly increasing folding. In the extreme case of $\lambda = 0$, both the performance metric and topology deteriorate substantially. This confirms that cocycle regularization is essential for stabilizing the learned deformation field and promoting physically plausible solutions of the underlying non-autonomous ODE.

5.2 Effect of Inference Compositions

Progressive compositions in TPFM-DIR are performed only at the inference. We observe that as few as four compositions are sufficient to obtain near-zero


Table 4: Quantitative comparison on the ACDC and CAMUS datasets. The best result per metric is shown in **bold**. The second-best result is highlighted in **red**.

Method	ACDC					CAMUS				
	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	ASSD $\downarrow$	SDLogJ $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	HD95 $\downarrow$	ASSD $\downarrow$	SDLogJ $\downarrow$
LDDMM	78.17 $\pm$ 1.52	0.0062 $\pm$ 0.0044	5.43 $\pm$ 0.63	2.02 $\pm$ 0.96	0.197 $\pm$ 0.131	79.30 $\pm$ 5.21	0.0019 $\pm$ 0.0013	6.25 $\pm$ 1.35	2.61 $\pm$ 0.92	0.195 $\pm$ 0.031
AdaCS	85.46 $\pm$ 1.12	0.1493 $\pm$ 0.0185	5.56 $\pm$ 0.34	2.12 $\pm$ 1.20	0.326 $\pm$ 0.131	82.67 $\pm$ 5.14	0.0029 $\pm$ 0.0130	5.23 $\pm$ 1.95	2.11 $\pm$ 0.76	0.193 $\pm$ 0.033
TransMorph-diff	82.13 $\pm$ 1.42	0.0012 $\pm$ 0.0005	6.55 $\pm$ 0.43	2.67 $\pm$ 1.36	0.162 $\pm$ 0.069	78.63 $\pm$ 5.54	0.0011 $\pm$ 0.0016	10.15 $\pm$ 3.16	4.04 $\pm$ 1.12	0.166 $\pm$ 0.028
TransMorph	84.47 $\pm$ 1.18	0.0418 $\pm$ 0.0087	6.24 $\pm$ 0.36	2.54 $\pm$ 1.22	0.575 $\pm$ 0.094	80.53 $\pm$ 5.93	0.1017 $\pm$ 0.1376	9.10 $\pm$ 3.16	3.68 $\pm$ 1.21	0.312 $\pm$ 0.075
DiffuseMorph	81.90 $\pm$ 1.35	0.0383 $\pm$ 0.0091	7.34 $\pm$ 0.47	2.81 $\pm$ 1.48	0.251 $\pm$ 0.092	76.43 $\pm$ 6.72	0.0932 $\pm$ 0.1329	9.77 $\pm$ 3.18	4.28 $\pm$ 1.27	0.393 $\pm$ 0.087
DiffuseReg	83.31 $\pm$ 1.21	0.2345 $\pm$ 0.0228	6.17 $\pm$ 0.39	2.35 $\pm$ 1.31	0.398 $\pm$ 0.107	77.29 $\pm$ 6.19	0.4026 $\pm$ 0.1402	10.60 $\pm$ 3.04	4.22 $\pm$ 1.23	0.395 $\pm$ 0.062
DGIR	82.39 $\pm$ 2.97	0.0059 $\pm$ 0.0125	7.43 $\pm$ 1.08	2.26 $\pm$ 0.79	0.241 $\pm$ 0.021	81.83 $\pm$ 6.27	0.0315 $\pm$ 0.0282	6.71 $\pm$ 2.02	2.78 $\pm$ 0.82	0.301 $\pm$ 0.066
TPFM-DIR	87.42 $\pm$ 0.99	0.0002 $\pm$ 0.0001	5.36 $\pm$ 0.35	1.91 $\pm$ 1.07	0.182 $\pm$ 0.023	85.12 $\pm$ 2.12	0.0 $\pm$ 0.0	4.72 $\pm$ 2.19	2.06 $\pm$ 0.94	0.162 $\pm$ 0.020

Table 5: Effect of the regularization weight λ on performance and topology preservation of TPFM-DIR across OASIS, LungCT, and ACDC.

λ	OASIS			LungCT			ACDC		
	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	TRE $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	TRE $\downarrow$	Dice $\uparrow$	$ J _{<0}\%$ $\downarrow$	TRE $\downarrow$
0	79.48	3.1263	4.02	2.1821	80.19	3.0671	80.19	3.0671	
5	86.90	0.0415	3.36	0.0019	84.52	0.0039	84.52	0.0039	
10	87.08	0.0019	2.19	0.0	87.42	0.0002	87.42	0.0002	
15	85.84	0.0001	2.41	0.0	87.08	0.0001	87.08	0.0001	
20	85.02	0.0	3.11	0.0	85.74	0.0	85.74	0.0	

$|J|_{<0}\%$. Figure 7 illustrates the effect of increasing the number of steps on Dice and $|J|_{<0}\%$. Although a single-step prediction is feasible, progressive compositions improve smoothness and reduce folding, with only a minor degradation in alignment performance. Beyond four steps, improvements in topology preservation saturate, making this choice a practical trade-off between computational cost and deformation regularity.

**Fig. 7:** The effect of number of compositions at inference on the Dice score and $|J|_{<0}\%$.

5.3 The Choice of Architecture

TPFM-DIR is a diffeomorphic registration framework rather than a specific network architecture. Although the backbone must support temporal conditioning




Fig. 8: Quantitative (table) and visual (figure) analysis of improvements obtained by training CorrMLP and TransMorph in the TPFM-DIR setting. Dice scores increase while $|J|_{<0\%}$ is significantly reduced across OASIS, IXI, Mindboggle, and AbdomenCT.

(e.g., via context addition or FiLM-based modulation [33]), the framework is not restricted to a particular design. To validate architectural flexibility, we adapt two strong deformable baselines, CorrMLP and TransMorph, by modifying only their decoders to incorporate time context via context addition (see Appendix E for more details). Figure 8 reports the results.

For each backbone, we compare three settings: the original deformable model, a conventional diffeomorphic variant trained with scaling and squaring (*-diff*), and the proposed TPFM-DIR formulation. In both architectures, TPFM-DIR improves Dice while drastically reducing folding percentage relative to the original deformable models. Moreover, compared with the scaling-and-squaring variants, TPFM-DIR consistently achieves higher Dice scores while maintaining comparable or better topology preservation. These findings demonstrate that TPFM-DIR is architecture-agnostic and can enhance topology preservation without sacrificing alignment accuracy.

5.4 Computational Analysis

To assess the computational characteristics of TPFM-DIR, we report the number of parameters, training and inference runtime per image pair, and GPU memory consumption in Table 6. We compare against a strong deformable model (HViT), a scaling-and-squaring-based diffeomorphic method (TransMorph-diff), a learning-based non-autonomous approach (R2Net), and an instance-optimization non-autonomous method (NODEO). Since NODEO is a pair-wise optimization method, its training statistics are reported with “-”. The performance is measured on OASIS as a representative large 3D dataset. TPFM-DIR achieves the fastest

Table 6: Computational analysis across methods for per-pair training/inference iterations. “Auto.” and “Non-Auto.” denote autonomous and non-autonomous ODE settings.

Method	# Params (M)	Train Time (s)	Inference Time (s)	Train Mem. (GB)	Inference Mem. (GB)	ODE Type
TransMorph-diff	46.8	0.98	0.47	12.3	4.3	Auto.
NODEO	4.2	-	214	-	3.5	Non-Auto.
R2Net	11.22	1.53	0.96	11.8	3.3	Non-Auto.
HViT	21.2	1.57	0.41	24.6	5.76	-
TPFM-DIR	30.6	1.31	0.38	18.1	2.7	Non-Auto.

inference time while maintaining a low inference memory footprint. Comparing to models incorporating non-autonomous ODEs, TPFM-DIR exhibits faster training and inference times due to avoiding numerical integration during optimization. All measurements are obtained on an NVIDIA RTX3090 GPU.

6 Conclusion

We introduced TPFM-DIR, a diffeomorphic image registration framework that directly parameterizes the flow map of a non-autonomous differential equation. By enforcing the anchored cocycle property during training, the method learns temporally consistent diffeomorphic trajectories without requiring numerical integration during optimization. Extensive experiments across nine datasets spanning multiple modalities, anatomical regions, and spatial dimensions demonstrate that TPFM-DIR substantially narrows the long-standing performance gap between diffeomorphic models and strong unconstrained deformable methods. In many settings, it achieves superior registration accuracy while preserving topology. Importantly, TPFM-DIR is architecture-agnostic and can be integrated into existing backbones with minimal modifications, resulting in consistent improvements of deformation regularity and empirical performance. Beyond medical image registration, the proposed formulation provides a principled framework for learning solution operators of non-autonomous ODEs in a data-driven manner.

Acknowledgements

The authors acknowledge funding support from NSERC Discovery Grants and the Department of Computing Science, University of Alberta. The project received compute support from Digital Research Alliance of Canada.

References

1. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2006: 9th International Conference, Copenhagen, Denmark, October 1-6, 2006. Proceedings, Part I 9. pp. 924–931. Springer (2006)
2. 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)
3. Avants, B.B., Tustison, N.J., Song, G., Cook, P.A., Klein, A., Gee, J.C.: A reproducible evaluation of ants similarity metric performance in brain image registration. *Neuroimage* **54**(3), 2033–2044 (2011)
4. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9252–9260 (2018)
5. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
6. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* **61**, 139–157 (2005)
7. Bostelmann, J., Lellmann, J.: Liemorph: Transformer-based image registration using flows on lie groups (2025)
8. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: Transmorph: Transformer for unsupervised medical image registration. *Medical image analysis* **82**, 102615 (2022)
9. Chen, Z., Zheng, Y., Gee, J.C.: Transmatch: a transformer-based multilevel dual-stream feature matching network for unsupervised deformable image registration. *IEEE transactions on medical imaging* **43**(1), 15–27 (2023)
10. Cheng, X., Zhang, T., Lu, W., Meng, Q., Frangi, A.F., Duan, J.: Sacb-net: Spatial-awareness convolutions for medical image registration. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 5227–5237 (2025)
11. Christensen, G.E., Johnson, H.J.: Consistent image registration. *IEEE transactions on medical imaging* **20**(7), 568–582 (2001)
12. Christensen, G.E., Song, J.H., Lu, W., El Naqa, I., Low, D.A.: Tracking lung tissue motion and expansion/compression with inverse consistent image registration and spirometry. *Medical physics* **34**(6Part1), 2155–2163 (2007)
13. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I. pp. 729–738. Springer (2018)

14. Duan, B., Zhong, M., Yan, Y.: Towards saner deep image registration. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12459–12468 (2023)
15. Dupuis, P., Grenander, U., Miller, M.I.: Variational problems on flows of diffeomorphisms for image matching. *Quarterly of applied mathematics* pp. 587–600 (1998)
16. Geng, Z., Deng, M., Bai, X., Kolter, J.Z., He, K.: Mean flows for one-step generative modeling. *arXiv preprint arXiv:2505.13447* (2025)
17. Ghahremani, M., Khateri, M., Jian, B., Wiestler, B., Adeli, E., Wachinger, C.: H-vit: A hierarchical vision transformer for deformable image registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11513–11523 (2024)
18. Joshi, A., Hong, Y.: R2net: Efficient and flexible diffeomorphic image registration using lipschitz continuous residual networks. *Medical Image Analysis* **89**, 102917 (2023)
19. Kennedy, D.N., Haselgrove, C., Hodge, S.M., Rane, P.S., Makris, N., Frazier, J.A.: Candishare: a resource for pediatric neuroimaging data. *Neuroinformatics* **10**, 319–322 (2012)
20. Kim, B., Han, I., Ye, J.C.: Diffusemorph: Unsupervised deformable image registration using diffusion model. In: European conference on computer vision. pp. 347–364. Springer (2022)
21. Kim, B., Kim, D.H., Park, S.H., Kim, J., Lee, J.G., Ye, J.C.: Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis* **71**, 102036 (2021)
22. Klein, A., Tourville, J.: 101 labeled brain images and a consistent human cortical labeling protocol. *Frontiers in neuroscience* **6**, 171 (2012)
23. Kloeden, P.E., Rasmussen, M.: Nonautonomous dynamical systems. No. 176, American Mathematical Soc. (2011)
24. Leclerc, S., Smistad, E., Pedrosa, J., Østvik, A., Cervenansky, F., Espinosa, F., Espeland, T., Berg, E.A.R., Jodoin, P.M., Grenier, T., et al.: Deep learning for segmentation using an open large-scale dataset in 2d echocardiography. *IEEE transactions on medical imaging* **38**(9), 2198–2210 (2019)
25. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747* (2022)
26. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of cognitive neuroscience* **19**(9), 1498–1507 (2007)
27. Meng, M., Feng, D., Bi, L., Kim, J.: Correlation-aware coarse-to-fine mlps for deformable medical image registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9645–9654 (2024)
28. Miller, M.I., Trounev, A., Younes, L.: On the metrics and euler-lagrange equations of computational anatomy. *Annual review of biomedical engineering* **4**(1), 375–405 (2002)
29. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer methods and programs in biomedicine* **98**(3), 278–284 (2010)
30. Mok, T.C., Chung, A.: Fast symmetric diffeomorphic image registration with convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4644–4653 (2020)

31. Mok, T.C., Chung, A.C.: Large deformation diffeomorphic image registration with laplacian pyramid networks. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23. pp. 211–221. Springer (2020)
32. Mok, T.C., Chung, A.C.: Conditional deformable image registration with convolutional neural network. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 35–45. Springer (2021)
33. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
34. Sakaridis, C., Dai, D., Van Gool, L.: Acdc: The adverse conditions dataset with correspondences for semantic driving scene understanding. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10765–10775 (2021)
35. Shattuck, D.W., Mirza, M., Adisetiyo, V., Hojatkashani, C., Salamon, G., Narr, K.L., Poldrack, R.A., Bilder, R.M., Toga, A.W.: Construction of a 3d probabilistic atlas of human cortical structures. *Neuroimage* **39**(3), 1064–1080 (2008)
36. Sheikhsafari, A., Noga, M., Punithakumar, K., Ray, N.: Unsupervised deformable image registration with fully connected generative neural network. In: Medical imaging with deep learning (2022)
37. Siegert, L., Fischer, P., Heinrich, M.P., Baumgartner, C.F.: Pulpo: Probabilistic unsupervised laplacian pyramid registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 717–727. Springer (2024)
38. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: A survey. *IEEE transactions on medical imaging* **32**(7), 1153–1190 (2013)
39. Tian, L., Greer, H., Vialard, F.X., Kwitt, R., Estépar, R.S.J., Rushmore, R.J., Makris, N., Bouix, S., Niethammer, M.: Gradicon: Approximate diffeomorphisms via gradient inverse consistency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18084–18094 (2023)
40. Tian, L., Sengupta, S., Greer, H., Estépar, R.S.J., Niethammer, M.: Nephi: Neural deformation fields for approximately diffeomorphic medical image registration. *arXiv preprint arXiv:2309.07322* (2023)
41. Tursynbek, N., Greer, H., Demir, B., Niethammer, M.: Guiding registration with emergent similarity from pre-trained diffusion models. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 240–251. Springer (2025)
42. Tustison, N.J., Avants, B.B., Gee, J.C.: Learning image-based spatial transformations via convolutional neural networks: A review. *Magnetic resonance imaging* **64**, 142–153 (2019)
43. Wu, Y., Jiahao, T.Z., Wang, J., Yushkevich, P.A., Hsieh, M.A., Gee, J.C.: Nodeo: A neural ordinary differential equation based optimization framework for deformable image registration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20804–20813 (2022)
44. Xu, J., Chen, E.Z., Chen, X., Chen, T., Sun, S.: Multi-scale neural odes for 3d medical image registration. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 213–223. Springer (2021)

45. Zhang, X., Stendahl, J.C., Staib, L.H., Sinusas, A.J., Wong, A., Duncan, J.S.: Adaptive correspondence scoring for unsupervised medical image registration. In: European Conference on Computer Vision. pp. 76–92. Springer (2024)
46. Zhang, Y., Pei, Y., Zha, H.: Learning dual transformer network for diffeomorphic registration. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24. pp. 129–138. Springer (2021)
47. Zhuo, Y., Shen, Y.: Diffusereg: Denoising diffusion model for obtaining deformation fields in unsupervised deformable image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 597–607. Springer (2024)