

SPHARM-Reg: Unsupervised Cortical Surface Registration using Spherical Harmonics

Seungeun Lee, Seunghwan Lee, Sunghwa Ryu and Ilwoo Lyu

Abstract— We present a novel learning-based spherical registration method, called SPHARM-Reg, tailored for establishing cortical shape correspondence. SPHARM-Reg aims to reduce warp distortion that can introduce biases in downstream shape analyses. To achieve this, we tackle two critical challenges: (1) joint rigid and non-rigid alignments and (2) rotation-preserving smoothing. Conventional approaches perform rigid alignment only once before a non-rigid alignment. The resulting rotation is potentially sub-optimal, and the subsequent non-rigid alignment may introduce unnecessary distortion. In addition, common velocity encoding schemes on the unit sphere often fail to preserve the rotation component after spatial smoothing of velocity. To address these issues, we propose a diffeomorphic framework that integrates spherical harmonic decomposition of the velocity field with a novel velocity encoding scheme. SPHARM-Reg optimizes harmonic components of the velocity field, enabling joint adjustments for both rigid and non-rigid alignments. Furthermore, the proposed encoding scheme using spherical functions encourages consistent smoothing that preserves the rotation component. In the experiments, we validate SPHARM-Reg on healthy adult datasets. SPHARM-Reg achieves a substantial reduction in warp distortion while maintaining a high level of registration accuracy compared to existing methods. In the clinical analysis, we show that the extent of warp distortion significantly impacts statistical significance.

Index Terms— Cortical surface, Spherical CNN, Spherical registration, Spherical warp

I. INTRODUCTION

Neuroimaging data analysis is a challenging task in the field due to the high structural variability across individuals. Neuroimaging data are often understood on 2-manifolds, which requires non-rigid surface registration [1]. Since a bijective

This work was supported in part by the National Research Foundation of Korea (NRF) under RS-2023-00251298, RS-2023-00266120, and RS-2024-00333931 and in part by the Institute for Information & Communications Technology Planning & Evaluation (IITP) under RS-2019-II191906, the Artificial Intelligence Graduate School Program, Pohang University of Science and Technology. (Co-first authors: Seungeun Lee and Seunghwan Lee, Corresponding author: Ilwoo Lyu)

Seungeun Lee is with the AI Research Team, KLeon, Seoul 04637, South Korea (e-mail: seungeun.lee@kleon.io).

Seunghwan Lee is with the Department of Computer Science and Engineering, Pohang University of Science and Technology (POSTECH), Pohang 37673, South Korea (e-mail: shwan@postech.ac.kr).

Sunghwa Ryu is with the Department of Bio and Brain Engineering, Korea Advanced Institute of Science and Technology, Daejeon 34051, South Korea (e-mail: sunghwaryu@kaist.ac.kr).

Ilwoo Lyu is with the Department of Computer Science and Engineering and the Graduate School of Artificial Intelligence, POSTECH, Pohang 37673, South Korea (e-mail: ilwoolyu@postech.ac.kr).

spherical mapping is readily available for cortical surfaces [1], spherical registration is a popular choice to establish shape correspondence. In general, the goals of spherical registration can be summarized as follows: to align the input geometry as closely as possible to the target geometry, to estimate diffeomorphic warp, and to reduce warp distortion. Nevertheless, fewer studies focus on warp distortion to date [2]–[4].

Spherical registration formulates an energy function that governs both geometry matching and warp field regularization. Through direct optimization of this energy function, spherical registration establishes shape correspondence by iteratively refining a smooth warp field. Renowned for their high registration accuracy, these techniques are extensively utilized in many structural data analyses [5]–[15]. Yet, in classical approaches [2], [16]–[21], the resource-intensive optimization process presents a common limitation, posing challenges for large-scale data analyses. Although recent approaches use convolutional neural networks (CNNs) for scalability and further employ a loss function related to distortion [3], [4], their architecture design for handling distortion remains implicit.

Existing approaches typically perform a rigid alignment only once before the subsequent non-rigid alignment. As discussed in [2], a rigid alignment may not be optimal for the following non-rigid alignment, which can potentially introduce unnecessary distortion. On the other hand, if a rigid alignment is optimized together with a non-rigid alignment, it becomes possible to reduce the warp distortion without sacrificing the registration accuracy. Consequently, in downstream tasks, regardless of the tessellation resolution, the sampling density is affected by warp distortion; greater distortion results in more (or fewer) sampling points, which can potentially introduce biases into shape analyses as shown in Fig. 1.

Velocity smoothing on the unit sphere is key to generating a smooth warp field and should be performed in a way that preserves the rotation component, which otherwise complicates its joint optimization. Two popular velocity encoding schemes are (a) polar displacements (*i.e.*, angular changes in latitude and longitude) and (b) tangent vectors in the tangent space at spherical locations. In scheme (a), smoothing is applied to each polar component. In scheme (b), parallel transport transfers vectors onto the tangent space at the smoothing center, followed by smoothing. Yet, neither scheme preserves the original rotation after smoothing. In scheme (a), the signs of latitudinal change vary with direction to the poles: positive toward the north pole and negative toward the south. In scheme (b), the original vector lengths are not preserved (see Fig. 2).

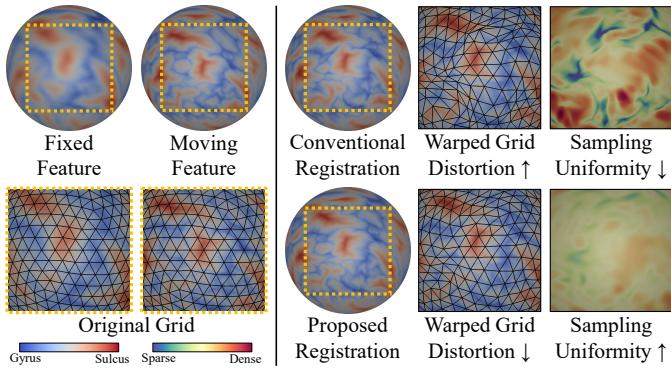


Fig. 1. Warp distortion and its impact on sampling uniformity. In this example, both conventional and our methods achieve comparable registration accuracy but differ in warp distortion. This leads to varying sampling densities between the original and warped grids. In downstream shape analyses, warp distortion causes over- (or under-) sampling in expanded (or shrunk) regions after registration, potentially introducing biases. Joint rigid alignment during optimization helps mitigate warp distortion.

In this paper, we propose a novel learning-based method for cortical surface registration, called SPHARM-Reg. The proposed method is designed to substantially reduce warp distortion while maintaining high registration accuracy. In particular, we present spherical harmonic-based decomposition of the velocity field, enabling joint rigid and non-rigid alignments by simultaneously optimizing the harmonic coefficients. This effectively reduces warp distortion that can arise from a sub-optimal rigid alignment followed by a non-rigid alignment as often caused in conventional approaches. Furthermore, we design a novel velocity encoding scheme that decomposes the velocity field into six independent spherical functions. This addresses the issue of velocity smoothing on the unit sphere in conventional encoding schemes that alter the underlying rotation component after spatial smoothing of the velocity field. The proposed method is built upon our preliminary studies [2], [22] with the following key differences.

- We provide the thorough theoretical foundations of velocity encoding with spectral smoothing as well as harmonic decomposition for joint rigid and non-rigid alignments within a diffeomorphic registration framework.
- We propose an end-to-end compositional learning framework for the warp field that enhances flexibility and improves registration accuracy across different types of geometry while accounting for their variability.
- We comprehensively validate SPHARM-Reg, including ablation studies, region-wise comparisons, and exploration of potential clinical applications.

II. RELATED WORK

A. Classical Spherical Registration

In an early attempt [16], planar projection of spherical polar coordinates was employed to find a non-rigid alignment between moving and fixed spherical data. This approach can be considered as standard image registration, which simplifies cortical surface registration on a regular grid. However, large mapping distortion may occur around the pole, resulting

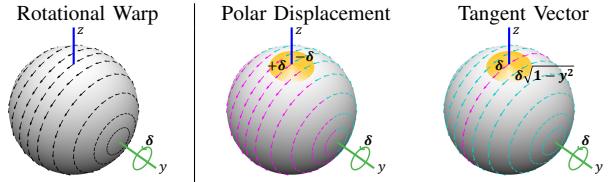


Fig. 2. Velocity smoothing at the pole (z -axis, blue) with a disk kernel (orange) in popular encoding schemes. *Left:* in this example, rotational warp produces counterclockwise velocity (black) with respect to the y -axis (green) by a degree of $\delta \in [0, 2\pi]$. *Middle:* the latitudinal changes in polar displacements are symmetrically encoded in both hemispheres with opposite signs (magenta and cyan); e.g., $+\delta$ and $-\delta$ for $y = 0$. Hence, smoothing results in no latitudinal change at the pole due to the equal balance of positive and negative changes in the kernel. *Right:* at a spherical location (x, y, z) , the magnitude of the tangent vector is $\delta\sqrt{1 - y^2}$. For $y = 0$, the magnitudes of the tangent vectors are maximal (magenta); elsewhere, they are smaller (cyan). In this scheme, parallel transport transfers the vectors in the kernel onto the tangent space at the pole before smoothing. As parallel transport preserves vector length, smoothing inevitably reduces the original magnitude at the pole due to the smaller magnitudes of vectors in the kernel. Neither scheme preserves the original rotation after smoothing.

in over-sampled data points. In this approach, the mapping distortion needs to be carefully addressed in the optimization.

Later, [18] extended the Demons algorithm [23] originally defined on a regular grid to the spherical domain without requiring the projection of spherical coordinates. One of the strengths of this work lies in its ability to provide diffeomorphic trajectories, where both the warp field and its inverse remain smooth and bijective. Compared to [16], this approach offers much faster registration via second-order optimization. However, potential inconsistency arises in the velocity smoothing process, where a rigid alignment turns into non-rigid warp.

In [20], spherical warp was estimated by a composition of rotation matrices. They simplified computing spherical warp as a labeling problem (*i.e.*, a discrete set of rotation candidates). The regularization was enforced by ensuring consistency in rotation labels throughout the optimization process. Their energy function is flexible in choosing multiple features for the registration metric as spherical registration becomes a labeling problem. However, their method lacks a diffeomorphism and instead limits the extent of the warp at each optimization step, which may require post-processing to unroll self-intersection.

More recently, spherical warp was modeled as a set of local rotations, in which each rotation smoothly changes over the unit sphere via spherical harmonic decomposition [2]. They pointed out that most spherical registration methods perform a rigid alignment once before the non-rigid alignment, which may be sub-optimal for the whole optimization. In this context, simultaneous optimization of rigid and non-rigid alignments can help reduce warp distortion. However, their approach models the warp field as a set of displacements rather than velocity vectors, which restricts the flexibility of the warp field.

Landmark-based spherical registration was also proposed in [17], [21], [24]–[26]. Although landmarks effectively address structural ambiguity, their acquisition requires additional effort and can be error-prone. Overall, these non-learning-based methods achieve high registration accuracy and are widely used in structural data analysis. Nevertheless, they often involve costly optimization processes.

B. Learning-based Spherical Registration

Recent advances in CNNs have shown great success in image analyses. Yet, only a few studies have explored CNN-based spherical registration. In early work, [27] extended the spherical registration framework proposed in [16]. They projected spherical polar coordinates onto a plane and adapted the learning-based image registration in Euclidean space [28]. Their approach achieves comparable performance to classical methods, while significantly reducing processing time. Albeit its promise, the approach inherits mapping distortion to Euclidean space, yielding unbalanced sampling around the poles as also noted in [16].

Similar to [27], [29] extended the idea of [28] by incorporating a spatial spherical CNN with the underlying registration framework proposed in [18]. Their approach employs spherical convolution, which avoids the spatial distortion present in [27]. However, the approach faces a singularity issue at the poles due to the predefined order of neighborhoods along a reference direction in the convolution filter. As the local coordinates flip at the poles, the filter orientation also flips. This results in discontinuity in the warp at the poles, yielding significant distortion. Although they fused the warp fields along three orthogonal axes, the pole issue was not fully addressed.

Recently, [3] extended the framework in [20]. With the increase in computing power, their approach can support a larger set of discrete rotations, resulting in more flexible warp than [20]. They applied rotation-equivariant convolutional operations on the unit sphere, which can address the singularity issue at the poles unlike [27], [29]. However, a discrete set of rotations inherently has a limited ability to generate a smooth warp field, depending on the resolution of a spherical tessellation. Later, [4] employed Euler angles for displacement encoding, but this scheme inherits axis-wise interactions; axis-wise angular interpolation leads to poor rotation interpolation [30], posing challenges for neural networks to learn effectively [31]. In [3], [4], warp distortion is mitigated by introducing a distortion loss, but their architecture design for distortion handling remains implicit. Due to the lack of a diffeomorphism in both approaches, any potential self-intersections on the registered sphere must be corrected through post-processing. To date, learning-based methods offer fast processing with high accuracy, but a systematic approach is still needed to achieve diffeomorphic warp with reduced distortion.

III. METHODS

A. Problem Statement

Given K moving geometric features $M : \mathbb{S}^2 \rightarrow \mathbb{R}^K$ and their corresponding target features $F : \mathbb{S}^2 \rightarrow \mathbb{R}^K$ on the unit sphere, the goal is to seek a warp field $\Phi : \mathbb{S}^2 \rightarrow \mathbb{S}^2$ such that $F(\mathbf{x}) = M(\Phi(\mathbf{x}))$ for $\forall \mathbf{x} \in \mathbb{S}^2$. Let \circ denote a function composition operator. Spherical registration minimizes the following energy function:

$$\int_{\mathbb{S}^2} \mathcal{L}_{sim}(F(\mathbf{x}), M \circ \Phi(\mathbf{x})) d\mathbf{x} + \alpha \int_{\mathbb{S}^2} \mathcal{L}_{reg}(\Phi(\mathbf{x})) d\mathbf{x}, \quad (1)$$

where $\mathcal{L}_{sim}(\cdot, \cdot) \in \mathbb{R}$ is a similarity term and $\mathcal{L}_{reg}(\cdot) \in \mathbb{R}$ is a regularization term to control the smoothness of Φ for some non-negative weighting factor $\alpha \in \mathbb{R}^+$.

If both Φ and its inverse mapping are differentiable and bijective, Φ is said to be diffeomorphic, in which the spherical warp has no self-intersection. We denote the tangent space of \mathbb{S}^2 at \mathbf{x} by $T_{\mathbf{x}}\mathbb{S}^2$. A diffeomorphic mapping can then be modeled by introducing a smooth velocity field on the tangent space in a continuous time domain $t = [0, 1]$ as an ordinary differential equation (ODE):

$$\frac{d\Phi(\mathbf{x}, t)}{dt} = v(\Phi(\mathbf{x}, t), t), \quad (2)$$

where $v \in T_{\mathbf{x}}\mathbb{S}^2$ is a time-varying velocity with the initial condition $\Phi(\mathbf{x}, 0) = \mathbf{x}$. The theories of ODEs [32] state that there exist unique trajectories over a velocity field $v(\mathbf{x}(t), t)$ over time t under the Lipschitz continuity in both \mathbf{x} and t . In the case of a stationary ODE, the solution is given by an exponential map, allowing the infinitesimal spherical warp generated by the Lie algebra to be composed:

$$\Phi(\mathbf{x}) = \lim_{n \rightarrow \infty} \Phi_{v/n} \circ \cdots \circ \Phi_{v/n}(\mathbf{x}), \quad (3)$$

where $\Phi_v(\mathbf{x})$ is the exponential map $\exp_{\mathbf{x}}(v)$ at $\mathbf{x} \in \mathbb{S}^2$, which mimics the geodesic at \mathbf{x} in the direction of v over unit time when v is small. For computational simplicity and efficiency, we assume that the velocity field is stationary in this work.

In the remainder of the paper, we describe the computation of the velocity field and its incorporation into a spherical CNN. We begin with the theoretical foundations of the proposed framework in a continuous spherical domain (velocity encoding, velocity field smoothing, and diffeomorphic warp), followed by its discretization into SPHARM-Reg. Fig. 3 illustrates a schematic overview of the proposed method.

B. Velocity Encoding

As discussed, conventional encoding schemes for velocity commonly use either polar displacements or tangent vectors. In both schemes, however, spatial smoothing introduces inconsistency in the sense that a rigid alignment becomes non-rigid, which eventually complicates the joint optimization of rigid and non-rigid alignments. To address this issue, we propose a novel encoding scheme for the velocity field using a rotational representation. Although there are multiple ways to encode a rotation, we follow [31] because this encoding scheme has no parameter-wise interaction and supports a homeomorphic mapping to the 3D rotation group, suitable for machine learning. More importantly, the proposed scheme allows for the smoothing of velocity vectors without altering the inherent rotation component. Specifically, we can recover the rotation matrix $R_{\mathbf{x}}$ at \mathbf{x} from six independent functions $r^{(i)}(\mathbf{x}) : \mathbb{S}^2 \rightarrow \mathbb{R}$, $i = 1, \dots, 6$. We denote two column vectors by $a_1 = [r^{(1)}(\mathbf{x}), r^{(2)}(\mathbf{x}), r^{(3)}(\mathbf{x})]^T$ and $a_2 = [r^{(4)}(\mathbf{x}), r^{(5)}(\mathbf{x}), r^{(6)}(\mathbf{x})]^T$, respectively. Here, we compute each column of $R_{\mathbf{x}}$ through the orthogonalization process [31]:

$$R_{\mathbf{x}} = \left[\begin{array}{c|c|c} b_1 = \frac{a_1}{\|a_1\|} & b_2 = b_3 \times b_1 & b_3 = \frac{a_1 \times a_2}{\|a_1 \times a_2\|} \end{array} \right]. \quad (4)$$

Once these parameters are determined, we have the encoded warp at \mathbf{x} given by

$$\Phi_v(\mathbf{x}) = R_{\mathbf{x}} \cdot \mathbf{x}. \quad (5)$$

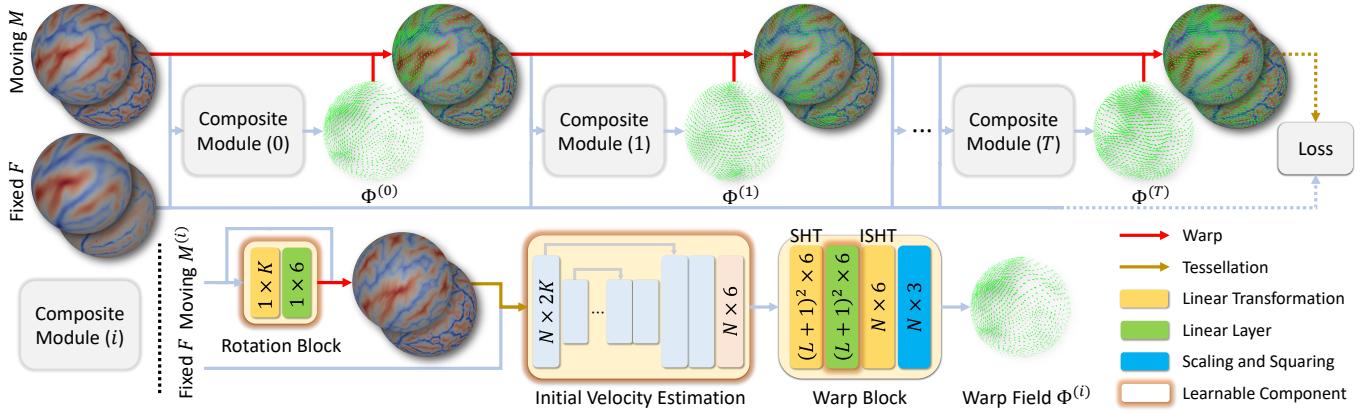


Fig. 3. Schematic overview. SPHARM-Reg executes T compositions of warp fields to perform multi-feature registration on N sampling points. Each composite module consists of the rotation block, initial velocity estimation, and warp block. At the i -th module, the coupled input of the moving geometry $M^{(i)}$ and the fixed target F is used to yield a warp field $\Phi^{(i)}$. The rotation block introduces rotational perturbation on $M^{(i)}$ to improve the initial velocity estimation. The velocity field is then encoded by six independent spherical functions. The warp block further refines the initial velocity by ensuring smoothness via the spherical harmonic transform (SHT) and its inverse transform (ISHT) with harmonic bandwidth L .

The velocity at \mathbf{x} is equivalent to the trajectory over unit time along the great circle from \mathbf{x} to $R_{\mathbf{x}} \cdot \mathbf{x}$. Because the six parameters are functions of spherical locations, the velocity field is viewed as a collection of independent spherical functions.

C. Velocity Field Smoothing

One desirable property of spherical warp is a smooth velocity field for a diffeomorphism. In our encoding scheme, this can be interpreted as smoothing over the six spherical functions. Thanks to the order independence of these functions, a smooth velocity field can be obtained by smoothing over each function in its own spherical space. For the velocity smoothing, one may apply a spatial smoothing filter to these spherical functions. However, the complexity of the smoothing process depends on the tessellation resolution and regularity such as the number of vertices, edge lengths, connectivity, *etc.* Also, the parametric smoothness is implicit rather than explicitly guaranteed. To address these issues, we propose spherical harmonic decomposition to encode our velocity field and encourage its smoothness, inspired by [2]. Here, computing a smooth velocity field can be considered as finding an appropriate set of harmonic coefficients $\{c\} \subset \mathbb{R}$. Once $\{c\}$ is determined, the i -th spherical function $r^{(i)}(\mathbf{x})$, $i = 1, \dots, 6$ can be reconstructed by a linear combination of the irreducible real harmonic basis functions $Y : \mathbb{S}^2 \rightarrow \mathbb{R}$:

$$r^{(i)}(\mathbf{x}) = c_{0,0}^{(i)} Y_{0,0}(\mathbf{x}) + \sum_{l=1}^L \sum_{m=-l}^l c_{l,m}^{(i)} Y_{l,m}(\mathbf{x}), \quad (6)$$

where r is truncated up to degree of L . The first term governs a rigid alignment. Since the reconstruction is band-limited, the spherical functions are inherently smooth. This strategy offers three main advantages. First, the velocity field becomes differentiable because the spherical harmonics consist of trigonometric functions that guarantee C^∞ continuity for $L > 0$. As high-frequency components are truncated by L in the spectral domain, the velocity field becomes smoother as L decreases, in which the smoothness can be easily adjustable.

Second, both rigid ($l = 0$) and non-rigid ($l > 0$) alignments are optimized simultaneously [2] (see Fig. 4). Third, the proposed spectral smoothing is independent of a spherical tessellation unlike conventional spatial smoothing. Any velocity field can be represented by a fixed number of parameters, $6(L+1)^2$, allowing for consistent spectral smoothing. Also, our strategy preserves the rotation component ($l = 0$) of the velocity field.

D. Diffeomorphic Warp

The warp field (or line integral) can be obtained over the velocity field. As discussed, our encoding scheme holds C^∞ continuity, and it is known that a warp field belongs to the same differentiability class as its underlying velocity field [32]. Thus, the proposed warp holds a diffeomorphism, which is smooth and invertible without self-intersection of the warp trajectories. Since our vector field is a collection of rotation matrices, the spherical warp can be computed as a composition of matrix multiplications. For a stationary velocity field, the warp field of Eq. (3) can be efficiently solved in a scaling and squaring fashion on \mathbb{S}^2 as a recursive form [18]:

$$\Phi_{v/2^{(k-1)}}(\mathbf{x}) = \Phi_{v/2^k} \circ \Phi_{v/2^k}(\mathbf{x}), \quad (7)$$

where k is the number of recursion steps. Theoretically, the solution becomes true for a sufficiently large number of k , but we will show that a small number of k is sufficient to achieve stable numerical integration in our work.

E. SPHARM-Reg

We focus on the discretization of the proposed spherical registration framework for its realization as SPHARM-Reg. In this section, we assume that the unit sphere is discretized into a spherical mesh with a set of N vertices represented as $\{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ through a valid spherical tessellation.

1) Rotation block: In existing learning-based approaches, M and F are rigidly aligned before their feature coupling. The registration performance depends on the initial alignment of the two input features that might not be optimal for the

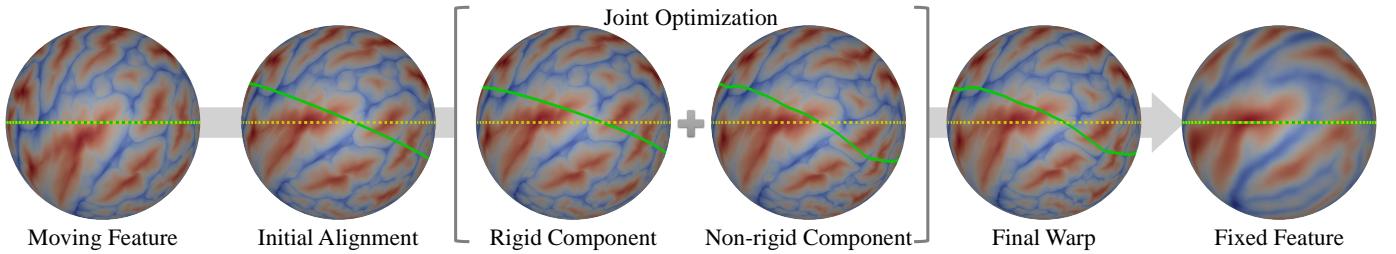


Fig. 4. Registration process and the need for joint optimization. The original equator and its updated position at each step are shown in green. The varying extent of deviation from the reference equator (yellow) indicates the adjusted orientation after the initial alignment, suggesting that the initial alignment may not be optimal. This ensures no warp distortion when combined with the non-rigid component since edge lengths are preserved.

velocity field estimation. Hence, we additionally introduce a learnable rotational perturbation of M to find the optimal feature orientation using six global parameters [31]. To this end, M is converted into a single scalar by a point-wise product with a learnable function $h : \mathbb{S}^2 \rightarrow \mathbb{R}$. The i -th global parameter $g^{(i)} \in \mathbb{R}$, $i = 1, \dots, 6$ is computed via a linear model:

$$g^{(i)} = \sum_{j=1}^N \sum_{n=1}^K M_n(\mathbf{x}_j) \cdot h_n(\mathbf{x}_j) \cdot w^{(i)} + b^{(i)}, \quad (8)$$

where $w^{(i)}, b^{(i)} \in \mathbb{R}$ are learnable parameters. The resulting rotation may not align optimally with F as it depends only on M . Again, the best alignment of M and F is not always ideal for velocity estimation. This stage is used solely to improve the velocity field estimation, independent of the learnable rigid component that is optimized later in our framework.

2) Initial velocity estimation: We feed the rotated M and F into a spherical CNN for the initial velocity estimation; *i.e.*, an initial guess of six spherical functions $r^{(i)}$, $i = 1, \dots, 6$. Although the proposed framework is flexible in choosing spherical CNNs, we use SPHARM-Net [33] as our backbone architecture because SPHARM-Net is based on spectral convolution independent of a spherical tessellation and some resources (*e.g.*, spherical harmonic basis functions) of the architecture are reusable in our framework.

3) Warp block: The initial estimated velocity field is likely neither smooth nor optimal. To refine these spherical functions and control their smoothness, the spherical functions are converted into the harmonic coefficients via the spherical harmonic transform. We have the following decomposition of the i -th spherical function through an inner product:

$$\hat{c}_{l,m}^{(i)} = \int_{\mathbb{S}^2} r^{(i)}(\mathbf{x}) Y_{l,m}(\mathbf{x}) \, d\mathbf{x}. \quad (9)$$

The above integral can be efficiently approximated by a Riemann sum for a sufficiently fine spherical tessellation [33], [34]. By letting $\Delta(\cdot)$ be the area at a spherical location, we have the following discretized form of Eq. (9):

$$\hat{c}_{l,m}^{(i)} \approx \sum_{j=1}^N r^{(i)}(\mathbf{x}_j) Y_{l,m}(\mathbf{x}_j) \Delta(\mathbf{x}_j). \quad (10)$$

Each component is then refined by its own linear model:

$$c_{l,m}^{(i)} = \hat{c}_{l,m}^{(i)} \cdot w_{l,m}^{(i)} + b_{l,m}^{(i)}, \quad (11)$$

where $w_{\cdot,\cdot}^{(i)}, b_{\cdot,\cdot}^{(i)} \in \mathbb{R}$ are learnable parameters. The refined coefficients are recovered back to vertex-wise rotations using Eqs. (6) and (4), and the vertex-wise warp trajectories are computed by the scaling and squaring approach of Eq. (7).

4) Composite module: Surface registration typically aims to align fine-grained geometry such as mean curvature. Direct registration from fine-grained geometry often leads to implausible alignments due to the local minima. Consequently, cascading-registration from coarse- to fine-grained geometry is necessary [2], [16], [18], [24], [35]. In our study, we adopt this approach but focus on multi-feature alignments together rather than individual feature alignments, while we maintain the same spherical tessellation throughout the entire process. This approach incorporates feature-wise interaction and enables end-to-end learning. We propose compositional learning that incorporates the residuals of the warp field. To this end, we stack $T + 1$ sets of the rotation and warp blocks. At the i -th set, we estimate the warp field $\Phi^{(i)}$ from $\Phi^{(i-1)}$. More formally, we have the following process:

$$\Phi(\mathbf{x}) = \Phi^{(T)} \circ \Phi^{(T-1)} \circ \dots \circ \Phi^{(0)}(\mathbf{x}). \quad (12)$$

5) Loss function: To optimize the energy function of Eq. (1), we consider an L^2 similarity metric as widely used in classical methods [2], [16], [18]. For $\forall \mathbf{x} \in \mathbb{S}^2$, we rewrite

$$\mathcal{L}_{sim}(F(\mathbf{x}), M \circ \Phi(\mathbf{x})) = \sum_{n=1}^K \omega_n (F_n(\mathbf{x}) - M_n \circ \Phi(\mathbf{x}))^2, \quad (13)$$

where $\omega_n \in \mathbb{R}^+$ is a weighting factor for the n -th feature. In our velocity encoding, the smoothness of the velocity field is governed by the harmonic truncation. To further reduce warp distortion, we evaluate the extent of warp after registration [2]. For $\forall \mathbf{x} \in \mathbb{S}^2$ and its neighborhoods $\mathcal{N}_{\mathbf{x}} \subset \mathbb{S}^2$, their arc length changes after spherical warp are measured by

$$\mathcal{L}_{reg}(\Phi(\mathbf{x})) = \frac{1}{2} \sum_{\mathbf{y} \in \mathcal{N}_{\mathbf{x}}} (\cos^{-1}(\mathbf{x}^T \mathbf{y}) - \cos^{-1}(\Phi(\mathbf{x})^T \Phi(\mathbf{y})))^2. \quad (14)$$

In this way, the velocity field stays as isometric as possible, which can reduce warp distortion [2]. Finally, the energy function in Eq. (1) is rewritten as the overall loss given by

$$\sum_{i=1}^N \mathcal{L}_{sim}(F(\mathbf{x}_i), M \circ \Phi(\mathbf{x}_i)) + \alpha \sum_{i=1}^N \mathcal{L}_{reg}(\Phi(\mathbf{x}_i)). \quad (15)$$

IV. EXPERIMENTAL DESIGN

A. Imaging Data

We used the HCP dataset [36] (1,113 scans) for training only and the Mindboggle dataset [37] (101 scans) for external validation to assess generalizability. For the clinical study, we used the OASIS dataset [38], of which 316 healthy adults (18–60 years old) were chosen. The cortical surfaces were reconstructed by a standard FreeSurfer pipeline [1] with 32 manually annotated regions of interest (ROIs), following the DKT protocol [37] (see Table I). These labels were used for ROI-wise performance evaluation of spherical registration methods. For M , we extracted two popular geometric features ($K = 2$) from the cortical surfaces: average convexity (*sulc*) that quantifies surface inflation and reflects relatively coarse geometry [1] and mean curvature (*curv*) that captures local bending and reflects relatively fine geometry for optimization and performance evaluation [2], [16], [18]. For \bar{F} , we used the official FreeSurfer template patterns called *fsaverage* [1].

B. Implementation and Training

Let $\omega = \omega_1/\omega_2$, where ω_1 and ω_2 are the weighting factors of \mathcal{L}_{sim} for *sulc* and *curv*, respectively. We chose $L = 40$, $k = 6$, $\alpha = 0.1$, $\omega = 1.4$, $T = 2$, and $C = 64$ for the initial velocity field estimation. We used the Adam optimizer [39] with an initial learning rate of 0.001, decaying by a factor of 0.5 if no improvement is made in four consecutive epochs. In the test phase, we used the learned model parameters at the peak performance on the validation set. We used an icosahedral mesh with a subdivision level of 6 for training. For inference on unseen data, the inferred spherical warp was applied to the original mesh. We used PyTorch for the backend processing. All experiments were conducted on an Intel Xeon 6248R CPU and an NVIDIA GeForce RTX 3090 GPU.

C. Evaluation Metrics

For the similarity evaluation metric, we computed the mean squared error (MSE) for each feature. Because MSE is a part of the optimization in the proposed method, we computed the normalized cross correlation (NCC) between the n -th features, F_n and M_n :

$$\frac{\sum_{i=1}^N (F_n(\mathbf{x}_i) - \bar{F}_n) \cdot (M_n \circ \Phi(\mathbf{x}_i) - \bar{M}_n)}{\sqrt{\sum_{i=1}^N (F_n(\mathbf{x}_i) - \bar{F}_n)^2} \cdot \sqrt{\sum_{i=1}^N (M_n \circ \Phi(\mathbf{x}_i) - \bar{M}_n)^2}}, \quad (16)$$

where $\bar{F}_n = \frac{1}{N} \sum_{i=1}^N F_n(\mathbf{x}_i)$ and $\bar{M}_n = \frac{1}{N} \sum_{i=1}^N M_n \circ \Phi(\mathbf{x}_i)$. NCC quantifies the geometric similarity between the warped input and the target; a higher value indicates a closer match in patterns. For the warp distortion metric, we measured the absolute areal ratio for each spherical location:

$$\exp(|\log(\Delta(\Phi(\mathbf{x}))/\Delta(\mathbf{x}))|). \quad (17)$$

D. Baseline Methods

We used public spherical registration methods: FreeSurfer [16], Spherical Demons (SD) [18], Hierarchical Spherical Deformation (HSD) [2], Multimodal Surface Matching (MSM)

[20], and Deep Discrete Spherical Registration (DDR) [3]. For a fair comparison, we tuned the hyperparameters of the baseline methods to match FreeSurfer's MSE of *curv*, where we adjusted their warp fields to be as smooth as possible. In FreeSurfer, we used default parameter settings. In SD, we set 3 iterations for the velocity field smoothing for all multi-resolution stages (icosahedral subdivision: 4 (*sulc*), 5 (*sulc*), 6 (*sulc*), 7 (*curv*)). In HSD, we set the harmonic bandwidth to 0, 5, 10, 15 for the respective optimization stages (icosahedral subdivision: 4 (*sulc*), 5 (*sulc*), 6 (*sulc*), 6 (*curv*)) by fixing the regularization weight of 230. In MSM, we employed default settings for *sulc*, while reducing the smoothing term and increasing the regularization term for *curv* (icosahedral subdivision: 4, 4, 5, 6 (*sulc*) and 4, 5, 5, 6 (*curv*)). We used NCC as a similarity metric as suggested in MSM for the best performance. In DDR, we followed their coarse-to-fine registration approach for each feature (icosahedral subdivision: 6, 6). We used the HCP dataset for training, and both NCC and MSE were employed for a similarity metric as suggested in DDR for the best performance. In MSM and DDR, we performed registration on *sulc* followed by *curv*. Default settings for other parameters were used for all the baseline methods unless specified otherwise.

V. RESULTS

A. Hyperparameter Choice

We justify our choice of the six hyperparameters: channel size C , harmonic bandwidth L , numerical integration step k , regularization weight α , multi-feature balance ω , and warp field composition T . For each hyperparameter, the other components were fixed to the default settings described in the Section IV-B. We varied C , L , k , α , ω and T in $\{16, 32, 64, 128, 256\}$, $\{10, 20, 30, 40, 50\}$, $\{0, 2, 4, 6, 8, 10\}$, $\{1/25, 1/20, 1/15, 1/10, 1/5\}$, $\{1.0, 1.2, 1.4, 1.6, 1.8, 2.0\}$, and $\{0, 1, 2, 3, 4\}$, respectively. Fig. 5 illustrates their behaviors.

1) *Channel size (C)*: We observed the performance improvement in both registration accuracy and warp distortion as C increases. Such a trend becomes plateau after $C = 128$. This result suggests that expanding the entry channel size can extract useful features up to $C = 128$ but additional channels beyond 128 have a marginal impact on the overall performance. In this study, we chose $C = 64$ to strike a balance between performance and computational efficiency.

2) *Harmonic bandwidth (L)*: A larger harmonic bandwidth allows a more flexible velocity field. As L increases, there is a tendency for registration accuracy to improve while areal distortion increases. This indicates a performance trade-off between the two metrics. It can be observed for $L > 40$ that the performance loss from warp distortion surpasses any improvements in registration accuracy. Consequently, $L = 40$ was chosen to achieve registration performance with minimal sacrifice of areal distortion. This hyperparameter does not affect the number of learnable parameters.

3) *Integration step (k)*: Increasing k improves the numerical precision of the integration, reducing discretization errors and yielding a solution that converges more closely to the true diffeomorphic trajectories. Nevertheless, only marginal

TABLE I

ROIs DEFINED BY THE DESIKAN-KILLIANY-TOURVILLE PROTOCOL [37]. NOTE THAT THE IDs ARE NOT COMPACT IN THE ORIGINAL PROTOCOL.

0. Unknown	2. Caudal Anterior Cingulate	3. Caudal Middle Frontal	5. Cuneus	6. Entorhinal
7. Fusiform	8. Inferior Parietal	9. Inferior Temporal	10. Isthmus Cingulate	11. Lateral Occipital
12. Lateral Orbitofrontal	13. Lingual	14. Medial Orbitofrontal	15. Middle Temporal	16. Parahippocampal
17. Paracentral	18. Pars Opercularis	19. Pars Orbitalis	20. Pars Triangularis	21. Pericalcarine
22. Postcentral	23. Posterior Cingulate	24. Precentral	25. Precuneus	26. Rostral Anterior Cingulate
27. Rostral Middle Frontal	28. Superior Frontal	29. Superior Parietal	30. Superior Temporal	31. Supramarginal
34. Transverse Temporal	35. Insula			

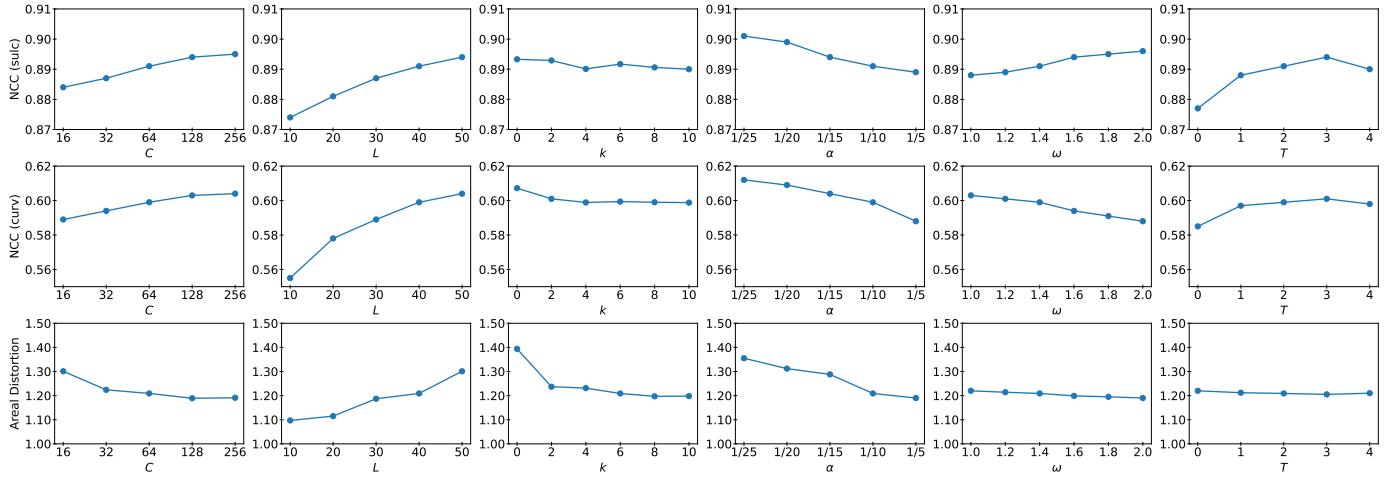


Fig. 5. Hyperparameter tuning. Channel size C : the registration accuracy and warp distortion are improved when C increases. Harmonic bandwidth L : the warp distortion is reduced when L is decreased. Integration step k : the registration accuracy and warp distortion become stable after $k = 6$. Regularization weight α : both registration accuracy and warp distortion decrease as α increases. Multi-feature balance ω : The alignment for *sulc* is improved while warp distortion is reduced. Warp field composition T : registration accuracy and warp distortion reach their peak at $T = 3$.

TABLE II

ABALATION STUDY. THE INCLUSION OF THE PROPOSED BLOCK(S) IN THE ABLATED SPHARM-REG ENHANCES REGISTRATION PERFORMANCE.

Method	NCC (sulc)	NCC (curv)	Dice	Distortion
Ablated SPHARM-Reg	0.876	0.585	0.871	1.275
+ Warp Block	0.882	0.586	0.877	1.244
+ Warp & Rotation Blocks	0.891	0.599	0.881	1.209

differences are observed in both registration accuracy and warp distortion beyond $k = 6$; as k increases up to 6, warp distortion decreases more rapidly than the drop in registration accuracy. Therefore, we chose $k = 6$ (equivalent to 64 integration steps) to strike a balance between computational efficiency and registration performance. In our experiments, this step size was sufficient to avoid self-intersection. This hyperparameter does not affect the number of learnable parameters.

4) Regularization weight (α): A downward trend appears in registration accuracy and warp distortion as α increases. Similar to other registration methods, assigning a large weight on the regularization term reduces registration accuracy as a strong constraint limits the flexibility of the warp field. We set $\alpha = 0.1$ to match FreeSurfer’s MSE of *curv*. This hyperparameter does not affect the number of learnable parameters.

5) Multi-feature balance (ω): A balance between multiple features impacts their own registration accuracy. We found that adjusting ω to be more biased towards coarse-grained features (*i.e.*, *sulc*) tends to reduce warp distortion more effectively. This is likely because coarse-grained features have simple

folding patterns, which does not require rigorous warp. We set $\omega = 1.4$ to match FreeSurfer’s MSE of *curv*. This hyperparameter does not affect the number of learnable parameters.

6) Warp field composition (T): Increasing the number of compositions enhances registration accuracy and effectively reduces warp distortion up to $T = 3$. Meanwhile, the performance begins to decline after $T = 3$. This trend suggests that multiple compositions ($T > 3$) of warp fields begin to overfit to the training data rather than promoting generalization. Despite the best performance at $T = 3$, we chose $T = 2$ to balance performance and inference time in this study.

B. Ablation Study

To validate our design choices, we conducted a quantitative comparison by disabling individual components and evaluating their impact. We ablated the proposed method to only contain the velocity estimation and diffeomorphic warp. For all experiments, we used the default parameter settings as described in the implementation in Section IV-B.

1) Warp block: The warp block can fine-tune the six spherical functions in the spectral domain, which is equivalent to spatial smoothing over each spherical function. Consequently, the warp block refines the initial velocity estimation following its harmonic decomposition. As shown in Table II, refining the harmonic coefficients improves registration accuracy as well as warp distortion. We also emphasize that this enhancement introduces only a modest increase in parameters, specifically $12 \cdot (L + 1)^2$ (*i.e.*, 20,172 extra parameters at $L = 40$).

TABLE III

PERFORMANCE BENCHMARK IN REGISTRATION ACCURACY, WARP DISTORTION, AND AVERAGE RUNTIME ON A SINGLE CPU CORE. THE RUNTIME IS MEASURED ON THE ORIGINAL SPHERICAL TESSELLATION (MEAN: 141,217 VERTICES) FOR THE ENTIRE PROCESS, INCLUDING MODEL INITIALIZATION, RIGID ALIGNMENT, SPHERICAL RE-TESSELLATION, AND SPHERICAL WARP. MODEL INFERENCE (SEC.): 24.02 (DDR) AND 14.78 (SPHARM-REG). **BOLD**: BEST PERFORMANCE. **GRAY**: MATCHED METRICS. **BLUE**: STATISTICAL SIGNIFICANCE ($p < 0.005$).

Method	Areal Distortion				Registration Accuracy				Time (sec.)	Param. (#)
	Mean	Median	$P_{95.4}$	$P_{99.7}$	MSE (curv)	NCC (sulc)	NCC (curv)	Dice		
FreeSurfer [16]	1.401	1.277	1.962	3.418	0.017	0.900	0.587	0.873	439.83	-
SD [18]	1.338	1.228	2.013	3.504	0.016	0.888	0.586	0.871	42.28	-
HSD [2]	1.337	1.228	1.960	3.810	0.016	0.887	0.586	0.871	206.82	-
MSM [20]	1.456	1.267	2.434	6.348	0.016	0.863	0.574	0.857	1006.46	-
DDR [3]	1.523	1.315	2.689	5.738	0.016	0.887	0.574	0.866	30.91	42.52M
SPHARM-Reg	1.209	1.139	1.557	3.248	0.016	0.891	0.599	0.881	21.67	3.59M

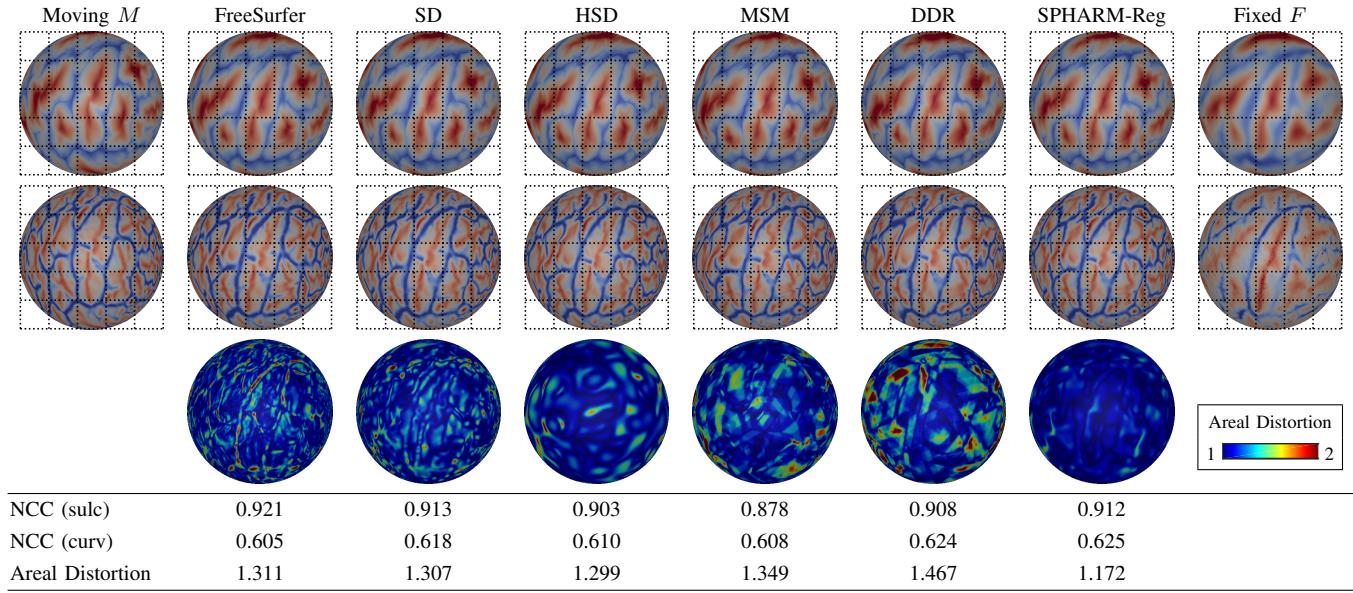


Fig. 6. Qualitative evaluation on an example subject: *sulc* (1st row), *curv* (2nd row), and mean areal distortion (3rd row). All the methods well-align M to F . SPHARM-Reg largely reduces the warp distortion without sacrificing registration accuracy.

2) *Rotation block*: Although the proposed method optimizes the global rotation component ($l = 0$), an explicit rotational adjustment can further improve the overall performance. After adding the warp block to the ablated SPHARM-Reg, introducing rotational perturbation to M by the rotation block adjusts the initial feature alignment to improve the velocity estimation. As summarized in Table II, the rotation block enhances overall registration accuracy and reduces warp distortion. Such improvement suggests that the initial feature alignment between M and F may not be optimal for the initial velocity estimation. Given that the required number of parameters in the rotation block is proportional to that of the spherical tessellation (*i.e.*, 40,962 vertices), such a slight increase of the parameters significantly improves overall performance.

C. Comparison to Baseline Methods

We performed a two-sided paired *t*-test against each baseline method conservatively at $p = 0.005$ to relieve a multi-comparison issue. For warp distortion, we measured the mean, median, and 95.4 and 99.73 percentiles of areal distortion.

1) *Overall performance*: Table III summarizes the overall performance in registration accuracy and warp distortion.

SPHARM-Reg outperforms the baseline methods in both registration accuracy and warp distortion (see Table III) except that FreeSurfer outperforms other methods in *sulc*. However, we were unable to match the MSE of *sulc* while fixing that of *curv*. This would otherwise require sacrificing areal distortion in other methods, resulting in a smaller MSE of *curv*. SPHARM-Reg achieves a reduction in the mean areal distortion of 0.128 (12.8% p reduced distortion) as compared to the second best method (HSD). In comparison to DDR, SPHARM-Reg utilizes only one-tenth of the parameters while accelerating registration. Figs. 6 and 7 show the visual inspection for an example individual and group average, respectively.

2) *ROI-wise analysis*: We investigated ROI-wise warp distortion and Dice coefficient as proposed in [2]. We performed a two-sided paired *t*-test against each baseline method. We then adjusted *p* values after controlling multi-comparison correction using the false discovery rate (FDR) at $q = 0.05$ [40]. In Fig. 8, SPHARM-Reg significantly outperforms the baseline methods and no worse distortion is found in all the ROIs. In most ROIs, SPHARM-Reg significantly outperforms the baseline methods, whereas worse performance is revealed in 3 and 2 ROIs than FreeSurfer and SD, respectively (see Fig. 9) despite their subtle absolute difference.

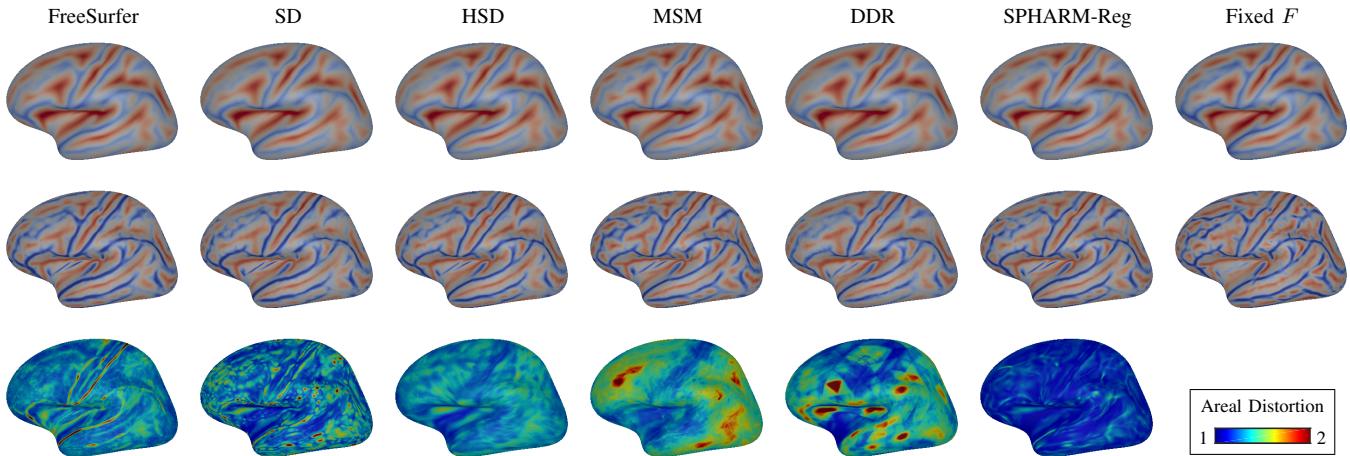


Fig. 7. Average feature and distortion maps: *sulc* (1st row), *curv* (2nd row), and areal distortion (3rd row). All the methods achieve comparable registration accuracy, while SPHARM-Reg enables the lowest areal distortion. The inflated surface is used for better visualization.

D. Cortical Thickness Analysis

We investigate the association between warp distortion and statistical significance in a clinical study. To this end, we employed the OASIS dataset to investigate cortical thickness changes along ages. We fitted a generalized linear model implemented in a standard FreeSurfer package [1] using shape correspondence established by FreeSurfer [16] and SPHARM-Reg. Note that since no ground truth is available, our goal is not to evaluate the superiority of the proposed method but rather to investigate the differences between the methods. The cortical surfaces were aligned to *fsaverage* [1]. Again, we used only the HCP dataset to train SPHARM-Reg. We matched the MSE of *curv* of 0.016 for both methods to ensure fair comparison. FreeSurfer achieves the mean areal distortion of 1.255, whereas SPHARM-Reg achieves that of 1.131, which is 12.4% reduction in distortion. Fig. 10 illustrates cortical thickness changes and their corresponding areal distortion. With an equal extent of MSE, the regions revealed by SPHARM-Reg fully contain those identified by FreeSurfer at the same level of statistical significance.

VI. DISCUSSION

A. Warp Distortion

SPHARM-Reg offers reduced warp distortion while keeping comparable registration accuracy to the baseline methods. This was made possible because SPHARM-Reg explicitly controls a rigid alignment throughout the whole registration process unlike the baseline methods. As observed in our experiments and HSD, the best alignment between moving and target features is not necessarily optimal, which can rather produce unnecessary warp distortion without any improvement in registration accuracy. Although the joint optimization of rigid and non-rigid alignments was originally proposed in HSD, we further enhanced that concept by introducing the rotation and warp blocks. Indeed, there is only subtle improvement in the ablated SPHARM-Reg over HSD in warp distortion (see Tables II and III). Meanwhile, the performance is improved through the specific handling of a rigid alignment in the proposed blocks, which may not be fully addressed by HSD.

B. Registration Accuracy

HSD is limited in its ability to increase harmonic bandwidth due to computational burden during the optimization process. This ultimately leads to an overly smooth warp field, potentially limiting registration accuracy. The issue has been addressed in the proposed learning-based framework through GPU acceleration. This further allows greater flexibility in spherical warp with fast inference. Moreover, inspired by the Spherical Demons framework, SPHARM-Reg computes a velocity field rather than displacements unlike HSD. Since SPHARM-Reg holds a diffeomorphism, no post-processing check is needed for self-intersection. This enables more flexible warp trajectories and accelerates the registration speed. Furthermore, although our current approach utilizes pure cortical geometry only, additional information such as myelination [20] or parcellation maps [28], [29] may improve registration accuracy, which we leave for future work.

Increasing the number of learnable parameters typically improves registration accuracy and reduces warp distortion unless the model becomes too large to fit, which could otherwise lead to overfitting. In our ablation studies, increasing C and T effectively raised the number of learnable parameters, improving overall performance. On the other hand, there is a trade-off between warp distortion and registration accuracy if the number of learnable parameters is fixed. More intensive non-rigid warp can achieve a better geometric alignment but introduces greater warp distortion. As such, varying hyperparameters that do not increase the number of learnable parameters consistently revealed this trade-off.

C. Clinical Applications

In the clinical data analysis, we observed the improved warp distortion while keeping comparable registration accuracy to FreeSurfer. Both approaches reveal a negative association between cortical thickness and ages due to the cerebral atrophy, which is reported in several studies such as [5], [7], [8], [12]. Here, warp distortion can lead to under- or over-sampling in group analyses, where statistical outcomes are potentially biased regarding the extent of warp distortion. Indeed, we found

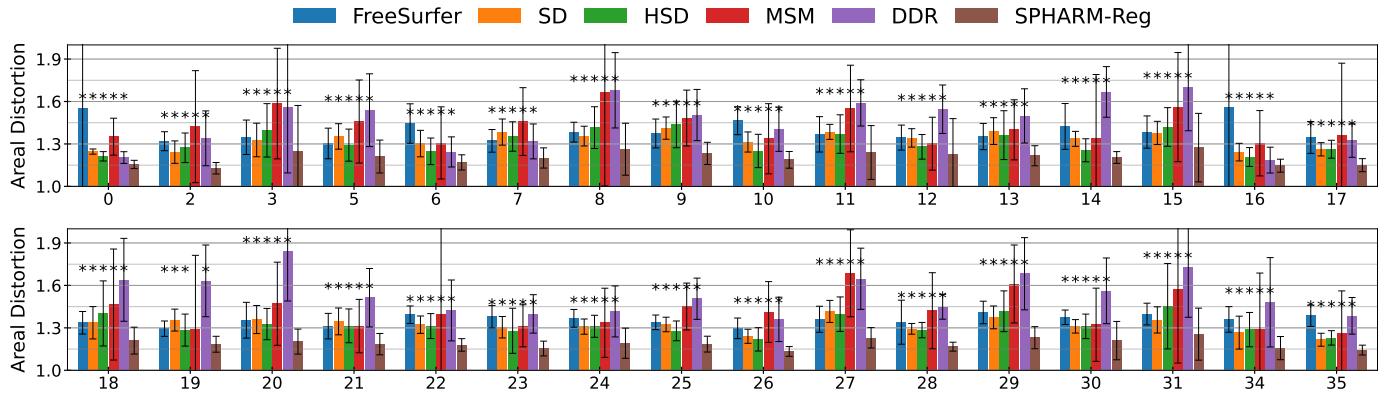


Fig. 8. ROI-wise mean areal distortion. The statistical significance is reported after FDR at $q = 0.05$. The proposed method achieves the lowest areal distortion, while no performance degradation is found against the baseline methods. *: statistical significance with better performance.

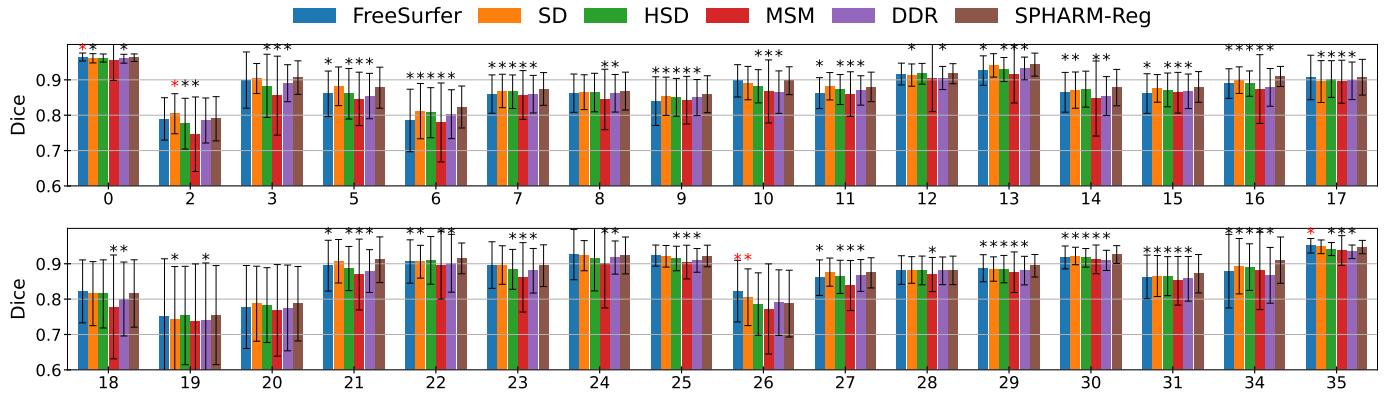


Fig. 9. ROI-wise Dice coefficient. The statistical significance is reported after FDR at $q = 0.05$. SPHARM-Reg significantly improves Dice coefficient for 16, 14, 21, 27, and 28 regions against FreeSurfer, SD, HSD, MSM, and DDR, respectively. Meanwhile, worse performance is revealed in 3 and 2 regions than FreeSurfer and SD. *: statistical significance with better performance. *: statistical significance with worse performance.

a wider area of regions that are associated with more reduced warp distortion than FreeSurfer. Furthermore, we observed a significant reduction in warp distortion such as the frontal lobe, linked to neurodevelopment [8], [11], [12], [14] and functions [9], [10], [15]. By reducing distortion in neurodevelopmental studies, we anticipate that our approach could reveal cortical regions previously obscured by conventional methods that may have been prone to warp distortion causing sampling biases.

D. Potential Improvement

The current approach is designed for pair-wise registration to a single target geometry. Although this is a popular approach in many neuroimaging studies, SPHARM-Reg trained to a specific target may not generalize well across a wide range of ages. This suggests the necessity of additional training for analyzing different age groups, as training biases toward a specific target may occur, otherwise. The target selection issue may be addressed by adopting earlier attempts of group-wise registration for unbiased target estimation [2], [18], [35]. Here, careful consideration will be needed in designing the iterative update of the target during the learning process. It would be feasible to fine-tune the trained SPHARM-Reg by substituting the target geometry rather than starting the training process from scratch. We will explore this avenue in future work.

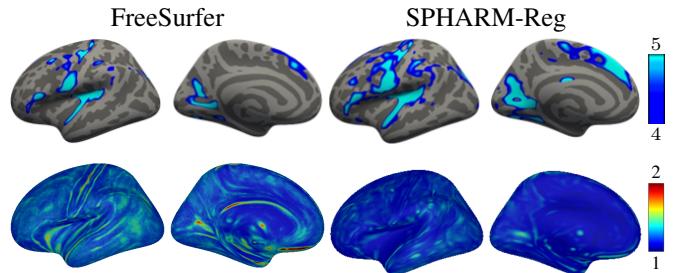


Fig. 10. Cortical thickness changes: revealed regions with the negative logarithm (base 10) of p values (top) and areal distortion (bottom). SPHARM-Reg reveals a wider area of regions with less warp distortion than FreeSurfer. Statistical analyses can be affected by warp distortion.

VII. CONCLUSION

In this paper, we presented a novel spherical harmonics-based cortical surface registration method. SPHARM-Reg generated a smooth velocity field via the harmonic decomposition and enabled simultaneous rigid and non-rigid alignments by optimizing the harmonic coefficients to minimize the unnecessary distortion. Additionally, we decomposed the velocity field into six independent spherical functions for consistent velocity smoothing. In the experiments, we demonstrated that SPHARM-Reg significantly reduces the warp distortion while

preserving the high-level of registration accuracy compared to existing methods. In the clinical study, we showed that warp distortion can affect outcomes of statistical analyses even when registration accuracy remains comparable across different spherical registration methods.

REFERENCES

- [1] B. Fischl, "Freesurfer," *Neuroimage*, vol. 62, no. 2, pp. 774–781, 2012.
- [2] I. Lyu, H. Kang, N. D. Woodward, M. A. Styner, and B. A. Landman, "Hierarchical spherical deformation for cortical surface registration," *Medical image analysis*, vol. 57, pp. 72–88, 2019.
- [3] M. A. Suliman, L. Z. Williams, A. Fawaz, and E. C. Robinson, "A deep-discrete learning framework for spherical surface registration," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2022, pp. 119–129.
- [4] J. Ren, N. An, Y. Zhang, D. Wang, Z. Sun, C. Lin, W. Cui, W. Wang, Y. Zhou, W. Zhang *et al.*, "Sugar: Spherical ultrafast graph attention framework for cortical surface registration," *Medical Image Analysis*, vol. 94, p. 103122, 2024.
- [5] P. Shaw, N. J. Kabani, J. P. Lerch, K. Eckstrand, R. Lenroot, N. Gogtay, D. Greenstein, L. Clasen, A. Evans, J. L. Rapoport *et al.*, "Neurodevelopmental trajectories of the human cerebral cortex," *Journal of neuroscience*, vol. 28, no. 14, pp. 3586–3594, 2008.
- [6] M. Datar, I. Lyu, S. Kim, J. Cates, M. A. Styner, and R. Whitaker, "Geodesic distances to landmarks for dense correspondence on ensembles of complex shapes," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22–26, 2013, Proceedings, Part II* 16. Springer, 2013, pp. 19–26.
- [7] B. A. Zielinski, M. B. Prigge, J. A. Nielsen, A. L. Froehlich, T. J. Abildskov, J. S. Anderson, P. T. Fletcher, K. M. Zygmunt, B. G. Travers, N. Lange *et al.*, "Longitudinal changes in cortical thickness in autism and typical development," *Brain*, vol. 137, no. 6, pp. 1799–1812, 2014.
- [8] I. K. Amlie, A. M. Fjell, C. K. Tamnes, H. Grydeland, S. K. Krogsrud, T. A. Chaplin, M. G. Rosa, and K. B. Walhovd, "Organizing principles of human cortical development—thickness and area from 4 to 30 years: insights from comparative primate neuroanatomy," *Cerebral cortex*, vol. 26, no. 1, pp. 257–267, 2016.
- [9] J. H. Gilmore, R. C. Knickmeyer, and W. Gao, "Imaging structural and functional brain development in early childhood," *Nature Reviews Neuroscience*, vol. 19, no. 3, pp. 123–137, 2018.
- [10] W. I. Voorhies, J. A. Miller, J. K. Yao, S. A. Bunge, and K. S. Weiner, "Cognitive insights from evolutionarily new brain structures in prefrontal cortex," *bioRxiv*, pp. 2020–11, 2020.
- [11] A. R. Zoltowski, I. Lyu, M. Failla, L. E. Mash, K. Dunham, J. I. Feldman, T. G. Woynaroski, M. T. Wallace, L. A. Barquero, T. Q. Nguyen *et al.*, "Cortical morphology in autism: findings from a cortical shape-adaptive approach to local gyration indexing," *Cerebral Cortex*, vol. 31, no. 11, pp. 5188–5205, 2021.
- [12] R. A. Bethlehem, J. Seiditz, S. R. White, J. W. Vogel, K. M. Anderson, C. Adamson, S. Adler, G. S. Alexopoulos, E. Anagnostou, A. Areces-Gonzalez *et al.*, "Brain charts for the human lifespan," *Nature*, vol. 604, no. 7906, pp. 525–533, 2022.
- [13] Y. Liu, S. Bao, D. J. Englot, V. L. Morgan, W. D. Taylor, Y. Wei, I. Oguz, B. A. Landman, and I. Lyu, "Hierarchical particle optimization for cortical shape correspondence in temporal lobe resection," *Computers in biology and medicine*, vol. 152, p. 106414, 2023.
- [14] Z. A. Stoebner, K. Hett, I. Lyu, H. Johnson, J. S. Paulsen, J. D. Long, and I. Oguz, "Comprehensive shape analysis of the cortex in huntington's disease," *Human Brain Mapping*, vol. 44, no. 4, pp. 1417–1431, 2023.
- [15] J. C. Pang, K. M. Aquino, M. Oldehinkel, P. A. Robinson, B. D. Fulcher, M. Breakspear, and A. Fornito, "Geometric constraints on human brain function," *Nature*, vol. 618, no. 7965, pp. 566–574, 2023.
- [16] B. Fischl, M. I. Sereno, R. B. Tootell, and A. M. Dale, "High-resolution intersubject averaging and a coordinate system for the cortical surface," *Human brain mapping*, vol. 8, no. 4, pp. 272–284, 1999.
- [17] D. C. Van Essen, "A population-average, landmark-and surface-based (pals) atlas of human cerebral cortex," *Neuroimage*, vol. 28, no. 3, pp. 635–662, 2005.
- [18] B. T. Yeo, M. R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland, "Spherical demons: fast diffeomorphic landmark-free surface registration," *IEEE transactions on medical imaging*, vol. 29, no. 3, pp. 650–668, 2009.
- [19] I. Lyu, S. H. Kim, J.-K. Seong, S. W. Yoo, A. C. Evans, Y. Shi, M. Sanchez, M. Niethammer, and M. A. Styner, "Group-wise cortical correspondence via sulcal curve-constrained entropy minimization," in *Information Processing in Medical Imaging: 23rd International Conference, IPMI 2013, Asilomar, CA, USA, June 28–July 3, 2013. Proceedings* 23. Springer, 2013, pp. 364–375.
- [20] E. C. Robinson, S. Jbabdi, M. F. Glasser, J. Andersson, G. C. Burgess, M. P. Harms, S. M. Smith, D. C. Van Essen, and M. Jenkinson, "Msm: a new flexible framework for multimodal surface matching," *Neuroimage*, vol. 100, pp. 414–426, 2014.
- [21] I. Lyu, S. H. Kim, J.-K. Seong, S. W. Yoo, A. Evans, Y. Shi, M. Sanchez, M. Niethammer, and M. A. Styner, "Robust estimation of group-wise cortical correspondence with an application to macaque and human neuroimaging studies," *Frontiers in neuroscience*, vol. 9, p. 210, 2015.
- [22] S. Lee, S. Ryu, S. Lee, and I. Lyu, "Unsupervised learning of cortical surface registration using spherical harmonics," in *International Workshop on Shape in Medical Imaging*. Springer, 2023, pp. 65–74.
- [23] T. Vercauteren, X. Pennec, A. Perchant, and N. Ayache, "Diffeomorphic demons: Efficient non-parametric image registration," *NeuroImage*, vol. 45, no. 1, pp. S61–S72, 2009.
- [24] X. Tao, J. L. Prince, and C. Davatzikos, "Using a statistical shape model to extract sulcal curves on the outer cortex of the human brain," *IEEE Transactions on Medical Imaging*, vol. 21, no. 5, pp. 513–524, 2002.
- [25] P. M. Thompson, K. M. Hayashi, E. R. Sowell, N. Gogtay, J. N. Giedd, J. L. Rapoport, G. I. De Zubicaray, A. L. Janke, S. E. Rose, J. Semple *et al.*, "Mapping cortical change in alzheimer's disease, brain development, and schizophrenia," *Neuroimage*, vol. 23, pp. S2–S18, 2004.
- [26] J. Glaunès, M. Vaillant, and M. I. Miller, "Landmark matching via large deformation diffeomorphisms on the sphere," *Journal of mathematical imaging and vision*, vol. 20, no. 1, pp. 179–200, 2004.
- [27] J. Cheng, A. V. Dalca, B. Fischl, L. Zöllei, A. D. N. Initiative *et al.*, "Cortical surface registration using unsupervised learning," *NeuroImage*, vol. 221, p. 117161, 2020.
- [28] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "Voxelmorph: a learning framework for deformable medical image registration," *IEEE transactions on medical imaging*, vol. 38, no. 8, pp. 1788–1800, 2019.
- [29] F. Zhao, Z. Wu, F. Wang, W. Lin, S. Xia, D. Shen, L. Wang, and G. Li, "S3reg: superfast spherical surface registration based on deep learning," *IEEE Transactions on Medical Imaging*, vol. 40, no. 8, pp. 1964–1976, 2021.
- [30] F. S. Grassia, "Practical parameterization of rotations using the exponential map," *Journal of graphics tools*, vol. 3, no. 3, pp. 29–48, 1998.
- [31] Y. Zhou, C. Barnes, J. Lu, J. Yang, and H. Li, "On the continuity of rotation representations in neural networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5745–5753.
- [32] G. Birkhoff and G.-C. Rota, *Ordinary differential equations*. John Wiley & Sons, 1978.
- [33] S. Ha and I. Lyu, "Spharm-net: Spherical harmonics-based convolution for cortical parcellation," *IEEE Transactions on Medical Imaging*, vol. 41, no. 10, pp. 2739–2751, 2022.
- [34] M. K. Chung, "Heat kernel smoothing on unit sphere," in *3rd IEEE International Symposium on Biomedical Imaging: Nano to Macro, 2006*. IEEE, 2006, pp. 992–995.
- [35] O. Lyttelton, M. Boucher, S. Robbins, and A. Evans, "An unbiased iterative group registration template for cortical surface analysis," *Neuroimage*, vol. 34, no. 4, pp. 1535–1544, 2007.
- [36] M. F. Glasser, S. N. Sotiroopoulos, J. A. Wilson, T. S. Coalson, B. Fischl, J. L. Andersson, J. Xu, S. Jbabdi, M. Webster, J. R. Polimeni *et al.*, "The minimal preprocessing pipelines for the human connectome project," *Neuroimage*, vol. 80, pp. 105–124, 2013.
- [37] A. Klein and J. Tourville, "101 labeled brain images and a consistent human cortical labeling protocol," *Frontiers in neuroscience*, vol. 6, p. 171, 2012.
- [38] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner, "Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults," *Journal of cognitive neuroscience*, vol. 19, no. 9, pp. 1498–1507, 2007.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [40] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: a practical and powerful approach to multiple testing," *Journal of the Royal statistical society: series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.