




LESA: Longitudinal Elastic Shape Analysis of Brain Subcortical Structures

Zhengwu Zhang, Yuexuan Wu, Di Xiong, Joseph G. Ibrahim, Anuj Srivastava & Hongtu Zhu


To cite this article: Zhengwu Zhang, Yuexuan Wu, Di Xiong, Joseph G. Ibrahim, Anuj Srivastava & Hongtu Zhu (2022): LESA: Longitudinal Elastic Shape Analysis of Brain Subcortical Structures, Journal of the American Statistical Association, DOI: [10.1080/01621459.2022.2102984](https://doi.org/10.1080/01621459.2022.2102984)


To link to this article: <https://doi.org/10.1080/01621459.2022.2102984>

 View supplementary material [↗](#)

 Accepted author version posted online: 20 Jul 2022.

 Submit your article to this journal [↗](#)

 Article views: 108

 View related articles [↗](#)

 View Crossmark data [↗](#)

LESA: Longitudinal Elastic Shape Analysis of Brain Subcortical Structures

Zhengwu Zhang¹, Yuexuan Wu², Di Xiong³, Joseph G. Ibrahim³, Anuj Srivastava^{2*}, and Hongtu Zhu^{1,3,4,5,6*}

¹Department of Statistics and Operations Research, University of North Carolina at Chapel Hill Chapel Hill, North Carolina

²Department of Statistics, Florida State University, Tallahassee, Florida

³Departments of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina

⁴Departments of Genetics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina

⁵Departments of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina

⁶Biomedical Research Imaging Center, University of North Carolina at Chapel, Hill Chapel Hill, North Carolina

*Joint senior authors and correspondences: htzhu@email.unc.edu and anuj@stat.fsu.edu.

This work was partially supported by U.S. NIH grants AG066970, MH120299, MH086633 and MH116527 and U.S. NSF grant 1953087. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH or any other funding agency.

Abstract

Over the past 30 years, magnetic resonance imaging has become a ubiquitous tool for accurately visualizing the change and development of the brain's subcortical structures (e.g., hippocampus). Although subcortical structures act as information hubs of the nervous system, their quantification is still in its infancy due to many challenges in shape extraction, representation, and modeling. Here, we develop a simple and efficient framework of longitudinal elastic shape analysis (LESA) for subcortical structures. Integrating ideas from elastic shape analysis of static surfaces and statistical modeling of sparse longitudinal data, LESA provides a set of tools for systematically quantifying changes of longitudinal subcortical surface shapes from raw structure MRI data. The key novelties of LESA include: (i) it can efficiently represent complex subcortical structures using a small number of basis functions and (ii) it can accurately delineate the spatiotemporal shape changes of the human subcortical structures. We applied LESA to analyze three longitudinal neuroimaging data sets and showcase its wide applications in estimating continuous shape trajectories, building life-span growth patterns, and comparing shape differences among different groups. In particular, with the Alzheimer's Disease Neuroimaging Initiative (ADNI) data, we found that the Alzheimer's Disease (AD) can significantly speed the shape change of ventricle and hippocampus from 60 to 75 years old compared with normal aging.

KEY WORDS: Alzheimer's Disease, Elastic Shape Analysis, Longitudinal shape trajectory, Principal components analysis, Subcortical structures.

1 Introduction

The present study is motivated by using magnetic resonance imaging (MRI) data in longitudinal neuroimaging studies, such as the baby connectome project (Howell et al., 2019) and the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Weiner et al., 2017), to accurately delineate the change and development of the brain subcortical structures (e.g., hippocampus) across time and/or groups. Subcortical structures include the diencephalon, pituitary gland, limbic structures and the basal ganglia, forming a group of diverse neural formations deep within the brain. These structures are not only involved in complex activities, such as memory, emotion, pleasure and hormone production, but also act as information hubs of the nervous system since they relay and modulate information passing to different areas of the brain. As an illustration, Figure 1 shows two extracted subcortical regions, lateral ventricle and hippocampus, from one randomly selected ADNI subject across four time points. After segmenting lateral ventricle and hippocampus across subjects, one may be interested in investigating the quantitative changes of their volumes, three dimensional (3D) surface shapes, and surface areas over time and the effect of some predictors of interest (e.g., disease status) on the shape change. The primary goal of this paper is to develop advanced image processing and statistical tools for characterizing the dynamic change of shapes of subcortical brain regions in the longitudinal setting.

Compared with cross-sectional shape analysis (Styner et al., 2006; Qiu and Miller, 2008; Kurtek et al., 2010), a distinctive feature of longitudinal shape data is that it has a dense spatial dimension, but a sparse temporal dimension (Hyun et al., 2016). Imaging measurements of the same individual often exhibit positive correlation temporally and the strength of the temporal correlation decreases with the time separation. Moreover, due to the inherent biological structure of the human brain, neuroimaging data are spatially correlated in nature and contain spatially contiguous regions. Efficiently dealing with such spatial and temporal dimensions raises at least three challenges. First, since each subject is only measured at a few time points in a typical longitudinal neuroimaging study, it is difficult to accurately reconstruct the

longitudinal profile of subcortical structures at the individual level. Second, most shape representations are in nonlinear manifolds (or rather than their quotient spaces), ruling out the direct application of standard longitudinal data models developed for Euclidean data. Third, the variability in individual growth patterns across subjects is subtle and can be easily overwhelmed by measurement and preprocessing errors.

There are three major types of shape analysis methods in the literature depending on how shapes are represented. The first type of methods uses a set of pre-determined shape features to quantify shape difference (Morra et al., 2009; Thompson et al., 2004; Madsen et al., 2015; Wang et al., 2010; Shi et al., 2015). Some examples include radiomics (van Timmeren et al., 2020) and topological data analysis (Amézquita et al., 2020). A potential issue is that such shape feature vectors only represent partial information about the original structures and thus, it is difficult to map them back to their corresponding shapes. The second type of methods is based on the large deformation diffeomorphic metric mapping (LDDMM) technique (Miller et al., 2002, 2006), which has been applied to longitudinal shapes (Tang et al., 2015; Tward et al., 2017; Lee et al., 2020). In LDDMM, a diffeomorphism as a smooth and bijective infinite-dimensional transformation with an inverse is used to map one shape to another and the size of the diffeomorphism provides a metric, called diffeomorphometry, to quantify shape difference. Such diffeomorphometry is more natural for quantifying pairwise shape changes than modeling the longitudinal shape trajectories (Tang et al., 2015). The third one is based on recent developments of longitudinal data modeling on manifolds (Muralidharan and Fletcher, 2012; Zhang et al., 2018a,b; Dai et al., 2018, 2020). However, since these methods were developed for relatively simple Riemannian manifolds, such as \mathbb{S}^2 , it is nontrivial to extend them to the shape space of surfaces. Therefore, one urgently needs a computationally simple but statistically powerful framework for the analysis of longitudinal subcortical shapes.

This paper aims to develop a *Longitudinal Elastic Shape Analysis* (LESA) framework. Our LESA can efficiently extract and represent shape data from raw MRIs, while addressing the aforementioned statistical challenges in longitudinal shape analysis. We make three important contributions. First, we use a single

parameterization-invariant, elastic Riemannian metric to minimize registration variability, while accounting for large shape variability. In contrast, most shape analysis methods use different metrics (or cost functions) for registration and comparisons (Pizer et al., 2003; Zhao et al., 2014). Second, the use of the elastic Riemannian metric leads to an effective low-dimension Euclidean representation of subcortical shape through using principal component analysis (PCA) in tangent spaces of the shape space. Our numerical data analyses demonstrate that the shape PCA in LESA has better representation power than popular approaches based on spherical harmonics representation (Shen et al., 2009) and sampling points on surfaces (Styner et al., 2006). Moreover, trajectories of longitudinal shapes reduce to those of scalar numbers in Euclidean space, facilitating the use of advanced statistical methods for studying longitudinal shapes (Yao et al., 2005; Wood, 2012; Fan and Gijbels, 2018). Our LESA integrates the developmental patterns of all subjects together, so it avoids large estimation errors caused by the standard two-stage approaches, including the estimation of individual temporal trajectories of shapes and the integration of all estimated shape trajectories (Singh et al., 2015; Fletcher, 2013). Third, the analysis examples and code for LESA along with its documentation are freely accessible from our websites at <https://wuyx5.github.io/LESA/> and <https://github.com/BIG-S2/Longitudinal-Elastic-Shape-Analysis-of-Brain-Subcortical-Structures>.

The remainder of this paper is structured as follows. Section 2 introduces three motivating data sets and their related scientific questions. Section 3 presents all major components of the LESA framework. Section 4 presents the data analysis results for the three motivating data sets. Section 5 concludes the paper with some discussions.

2 Motivating Data Sets and Scientific Questions

Understanding the growth pattern of subcortical structures and the effects of disease on such pattern is extremely important for aging and neuropsychiatric and neurodegenerative disorders. We consider MRIs obtained from three different longitudinal neuroimaging studies: the Alzheimer's Disease Neuroimaging Initiative (ADNI) (Petersen et al., 2010; Weiner et al., 2013; Basaia et al., 2019), the Human

Connectome Project (Glasser et al., 2016), and the OpenPain (Vachon-Preseu et al., 2016).

ADNI data set: We extracted the MRI data set from the ADNI database (adni.loni.usc.edu). The initial goal of ADNI was to test whether MRI, positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). The ADNI has four phases, including ADNI, ADNI-GO, ADNI2, and ADNI3, among which all subjects in ADNI-GO, ADNI2, and ADNI3 were scanned on 3T scanners. We included subjects with T1 MRI images in ADNI-GO and ADNI2 aged between 60 and 90 years old. These MRI T1 images were acquired using MPRAGE sequence with a resolution around $1 \times 1 \times 1.2 \text{ mm}^3$. After the data processing (refer to Sections 3.1 and 3.2), we conducted a careful quality control and removed outlying surfaces. Specifically, we first computed the sampled Karcher mean and then calculated the geodesic distance between each surface with the template. Next, we used a 95% confidence interval to detect potential outlying surfaces. We then visualized the potential outlying surfaces to manually remove the abnormal ones. Finally, we obtained a data set, called ADNIGO2, containing 1045 subjects with 3443 scans for the left lateral ventricle and 974 subjects with 3044 scans for the left hippocampus.

Human Connectome Project (HCP) test-retest data set: The Human Connectome Project (<https://db.humanconnectome.org/>) contains high-quality MRI data from around 1200 healthy young adults aged from 22 to 37. The T1 MRI images were acquired on a 3T Siemens Prisma scanner using multi-band sequence with a resolution of $0.7 \times 0.7 \times 0.7 \text{ mm}^3$. We included all the HCP young-adult subjects. Most of them just have one visit, while a small subset of subjects have two visits, resulting in longitudinal data with two time points. Similar to the ADNIGO2 data set, we applied the same data processing and quality control protocol. We obtained 1113 subjects with 1158 scans for the left lateral ventricle and 1082 subjects with 1125 scans for the left hippocampus.

OpenPain dataset: The OpenPain study (<http://www.openpain.org/>) is a five-year longitudinal study of the transition to chronic back pain. It contains 122 subjects aged

from 21 to 69. MRI scans were collected across four visits (two weeks, three months, six months, one year, and 2-3 years later). OpenPain's T1 MRI images were acquired on a 3T Siemens Trio whole-body scanner using the MPRAGE sequence with a resolution of $1 \times 1 \times 1 \text{ mm}^3$. Similar data processing and quality control were applied, and we ended up with 429 lateral ventricle and hippocampus surfaces from 117 subjects.

Table 1 shows more detailed statistics on the two subcortical regions studied and Figure 2 shows the age distribution of the three data sets. The three data sets together cover a lifespan age ranging from 20 to 90, allowing us to study the lifespan growing pattern during [20, 90] for lateral ventricle and hippocampus. We are particularly interested in the following scientific questions:

- (Q1) How to measure developmental changes in the shape of subcortical regions?
- (Q2) How to quantify the effect of disease or other covariates on subcortical shape changes?

To address (Q1) and (Q2), we need to develop an advanced longitudinal shape analysis pipeline below.

3 Methodology

In this section, we formally introduce LESA. Figure 3 presents a schematic overview of LESA, consisting of four key components: (i) surface extraction and parameterization; (ii) elastic shape analysis of surfaces; (iii) Euclidean representation of surface trajectories; and (iv) trajectory fitting and regression analysis. In the following subsections, we introduce each component in details.

3.1 Subcortical Surface Extraction and Parameterization

To analyze longitudinal subcortical shapes quantitatively, LESA represents each subcortical shape as a parameterized function given by $f : \mathbb{S}^2 \rightarrow \mathbb{R}^3$. This representation brings more convenience in analyzing the shape of subcortical structure, while removing shape confounding transformations, such as translation, rotation, and re-scaling. Our proposed LESA can also handle the parameterization variability, which controls the registration between surfaces; see Section 3.2 for details. Figure 4 illustrates our three-step pipeline to extract a parameterized subcortical shape. The first step is to segment the subcortical region, create a three-dimensional (3D) volume, and fill any holes inside the volume. This step is performed by using the FIRST tool inside the FMRIB Software Library (FSL) (Patenaude et al., 2011). It is done for each MRI T1 image at an individual level without registering them to a template. The second step is to build a surface mesh from the volume and use an area-

preserving, distortion minimizing spherical mapping (Jermyn et al., 2017) to map vertices on the mesh to a unit sphere for spherical parametrization. The third step is to refine our shape representation by improving the sample grid on \mathbb{S}^2 through using a uniformly sampled grid along the polar and azimuthal angles and fitting the corresponding function values in \mathbb{R}^3 . Finally, we obtain a parameterized surface as a mapping from \mathbb{S}^2 to \mathbb{R}^3 as shown in the third column of Figure 4.

3.2 Elastic Shape Analysis of Surfaces

For a given subcortical region, we observe longitudinal surface data f_{ij} for subject i at time point (or age) t_{ij} for $i = 1, \dots, n$ and $j = 1, \dots, m_i$. We use an elastic Riemannian metric to compare all surfaces $\{f_{ij}\}$ under a Riemannian framework, called elastic shape analysis. Such elastic shape analysis refers to a set of comparison methods of shapes of surfaces in a manner that is invariant to rigid motions, global scaling, and re-parameterization. These methods solve for dense, optimal registrations of points across surfaces, *while* comparing their shapes and not as a pre-processing step. Therefore, the registered surfaces reserve the shape heterogeneity, but they minimize the cross-sectional variance. Examples of elastic shape analysis of 3D objects include Younes (2010) and Bauer and Bruveris (2011). In this paper, we take the approach introduced in Jermyn et al. (2012) that uses a specific square-root representation to transform complicated, but important invariant Riemannian metrics into standard Euclidean metrics.

Let \mathcal{F} be the set of surfaces consisting of all smoothed maps $f: \mathbb{S}^2 \rightarrow \mathbb{R}^3$ with a finite \mathbb{L}_2 norm, and Γ be the set of all orientation-preserving diffeomorphisms of \mathbb{S}^2 . For any surface $f \in \mathcal{F}$ and $\gamma \in \Gamma$, the composition $f \circ \gamma$ is simply a re-parameterization of f and has the same shape as f . We consider any two surfaces, f_1 and f_2 such that $f_1(s)$ is registered to $f_2(s)$ for all $s \in \mathbb{S}^2$. If we re-parameterize f_2 by a $\gamma \in \Gamma$, then $f_1(s)$ is now registered to $f_2(\gamma(s))$. Thus, γ here controls the correspondence or registration among points between surfaces. In order to compare *shapes* of surfaces, we need a metric that can invariantly compare surfaces with arbitrary reparameterizations, motivating the following normal vector field representation of surfaces.

For $s \in \mathbb{S}^2$, the vector $n_f(s) = \frac{\partial f}{\partial u}(s) \times \frac{\partial f}{\partial v}(s)$ denotes the normal to f at the point $f(s)$, where $(u, v) \equiv s$ are the local coordinates on \mathbb{S}^2 . Then, the square-root normal field (SRNF) of f is defined to be the normal vector field $q: \mathbb{S}^2 \rightarrow \mathbb{R}^3$ by $q(s) = n_f(s) / \sqrt{|n_f(s)|}$, where $|\cdot|$ denotes the vector norm. As described in Jermyn et al. (2012), the \mathbb{L}^2 -metric under SRNF representation has some critical invariant properties and can be used to compare shapes of surfaces. The essential advantage of using such representation is that it is easy to remove shape-preserving transformations (reparameterizations) from this representation. The SRNF of a surface is already invariant to its translation. Scaling can be separated by re-scaling all surfaces to have unit area: $f(s) = f(s) / \sqrt{\alpha_f}$, where $\alpha_f = \int_{\mathbb{S}^2} |n_f(s)| ds$ is the surface area of f . However, the size of subcortical regions is an important feature, so it will be preserved and analyzed separately.

After appropriately removing scaling and translation, we handle rotation and re-parameterization as follows. Let $SO(3)$ be the rotation group (the set of all 3×3 rotation matrices). Applying a rotation $o \in SO(3)$ and a re-parameterization $\gamma \in \Gamma$ to a surface f is given by $o(f \circ \gamma)$. Thus, the SRNF representation of $o(f \circ \gamma)$ becomes

$o(q \star \gamma) \equiv o \sqrt{J_\gamma} (q \circ \gamma)$, where J_γ is the determinant of the Jacobian of γ . The removal of rotation and re-parameterization leads to the following registration problem:

$$(O^*, \gamma^*) = \underset{o \in SO(3), \gamma \in \Gamma}{\operatorname{argmin}} \|q_1 - o(q_2 \star \gamma)\|, \quad (1)$$

where q_1 and q_2 are the SRNFs of normalized and centered f_1 and f_2 , respectively. The optimal o^* is solved by using Procrustes Analysis, and the optimal γ^* is solved by using a gradient-based optimization over Γ (Jermyn et al., 2012; Kurtek et al., 2010). The minimum value of the objective function, call it d_s , forms the *elastic shape metric* between f_1 and f_2 , and γ^* represents the optimal registration of points across the two surfaces.

3.3 PCA-based Dimension Reduction

The goal of this step is to jointly align all surfaces and then perform principal component analysis (PCA) to obtain their finite-dimensional representations. It allows us to transform the complex shape trajectory into a trajectory in \mathbb{R}^r , leading to a simple downstream analysis.

Let $\{f_{ij}\}$ be a set of normalized surfaces (after removing translation and scaling).

The group alignment of $\{f_{ij}\}$ involves (i) the computation of a template shape and (ii) the pair-wise alignment of every f_{ij} to the template. Specifically, we use the Karcher mean under our elastic shape metric as the template, which is defined as

$$f_\mu = \operatorname{argmin}_f \sum_{i=1}^n \sum_{j=1}^{m_i} d_s(f, f_{ij})^2, \text{ where } d_s(\cdot, \cdot) \text{ denotes the shape metric. We}$$

approximate the optimum using an iterative approach. In each iteration, we register the given surfaces to the current estimate of the mean, and then we update this estimate by a mean of the current registered shapes. In the process of calculating Karcher mean, we also have all f_{ij} s aligned to the Karcher mean. Denote the aligned surface as f_{ij}^* , and $f_{ij}^* = O^*(f_{ij} \circ \gamma^*)$, where $(O^*, \gamma^*) = \operatorname{argmin}_{o \in SO(3), \gamma \in \Gamma} \|q_\mu - o(q_{ij} \star \gamma)\|$, in which q_μ is the SRNF of f_μ .

With the Karcher mean f_μ and aligned shapes f_{ij}^* , we perform dimension reduction in the tangent space at f_μ . Specifically, we compute the shooting vectors or deformations, that take the mean shape f_μ to individual surfaces f_{ij}^* as follows. Although there is an elaborate procedure for computing these deformations using the geometry of the shape space as described in Kurtek et al. (2010), we approximate these deformations by taking simple differences according to $v_{ij} = f_{ij}^* - f_\mu$ for simplicity. This Euclidean metric is different from the \mathbb{L}^2 metric in (1) used for aligning surface shapes. The metric in (1) provides optimal registrations between surfaces, but the subsequent analysis can get computationally expensive. If the underlying variability is small, the results from the two approaches are not that different, motivating us to use the simple Euclidean metric for downstream analyses. Next, we use the Gram-Schmidt procedure to generate an orthogonal basis for the set $\{v_{ij}\}$. Let $\{v'_k\}$ be the new orthogonal basis resulting from the Gram-Schmidt process. Thus, each original shooting vector v_{ij} can be projected onto the basis $\{v'_k\}$ and is represented as a vector of coefficients c_{ij} with its k -th element $c_{ijk} = \langle\langle v_{ij}, v'_k \rangle\rangle$. The original aligned surface f_{ij}^* can be recovered by $f_{ij}^* = f_\mu + \sum_{k=1} c_{ijk} v'_k$. In this step, there is no loss of information since $\{v'_k\}$ is just a new orthogonal basis for the subspace spanned by the shooting vectors $\{v_{ij}\}$ at the tangent space of f_μ .

These coefficients $\{c_{ij}\}$ denote Euclidean representations of original shapes, and we perform PCA in the coefficient space. We calculate the sample covariance matrix as

$$\mathbf{K} = \sum_{i=1}^n \sum_{j=1}^{m_i} c_{ij} c_{ij}^T / (\sum_{i=1}^n m_i - 1) \text{ and its spectral decomposition } \mathbf{K} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T, \text{ where}$$

$\mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2, \dots)$ is a diagonal matrix formed from the eigenvalues of \mathbf{K} and the columns of \mathbf{U} form the eigenvectors of \mathbf{K} . Let \mathbf{u}_j be the j -th column of \mathbf{U} corresponding to the j -th largest eigenvalue of \mathbf{K} . Thus, the Euclidean

representation of f_{ij}^* can now be approximated using the projection: $\tilde{c}_{ij} = \mathbf{U}_r^T c_{ij}$,

where $\mathbf{U}_r = [\mathbf{u}_1, \dots, \mathbf{u}_r]$ is the first r columns of \mathbf{U} , and $\tilde{c}_{ij} \in \mathbb{R}^r$. The k -th element in \tilde{c}_{ij} is denoted as \tilde{c}_{ijk} , representing the k -th principal component (PC) score for the shape surface f_{ij}^* , and its corresponding PC direction is given as $\tilde{v}_k = \sum_{i=1} u_{ki} v'_i$, where

u_{ki} is the i -th elements in \mathbf{u}_k .

Figure 5 (a) shows the Karcher mean of all 3443 left ventricles in the ADNIGO2 dataset discussed in Section 2. Figure 5 (b) shows the cumulative percentage of variance explained by the number of principal components. As shown here, the use of 32 PCs can represent the 95% variation of all surfaces. Figure 5 (c) shows the first PC direction in the shape space by reconstructing the principal geodesic as $f_{\mu} + t\sqrt{\lambda_1} * PC_1$, where PC_1 represents the first principal direction, i.e., \tilde{v}_1 . The PC1 mainly describes the shape change of anterior and posterior ends of the ventricle. In the following ADNI data analysis, significant differences can be observed in these regions between normal controls and AD people. We then bring the temporal labels back (the time of each observation) and plot the area trajectories for 1045 subjects in Figure 5 (d) and PC1 score trajectories in Figure 5 (e).

3.4 Dense Trajectory Fitting and Longitudinal Data Analysis

The goal of this step is to estimate continuous trajectories of shapes for all subjects and conduct longitudinal data analysis based on the outputs of Sections 3.1-3.3, including a surface area trajectory $[\alpha_{i1}, \dots, \alpha_{im_i}]$ and PC score trajectories $[\tilde{c}_{i1}, \dots, \tilde{c}_{im_i}]$ for each subcortical region from subject i . There are two challenges. The first challenge is that we only have sparse observations per subject in longitudinal neuroimaging studies, i.e., each m_i is a small integer. The second challenge is the non-uniform spacing of time points, i.e., the surfaces are observed at different times for different subjects. Due to these challenges, the independent fitting of sparse longitudinal points to trajectories does not work (James et al., 2000).

We develop two approaches for estimating continuous curves by borrowing information from all trajectories in the data set. The first approach is a semi-parametric mixed-effects model and the second approach is the functional data analysis of sparse longitudinal data.

Mixed Effects Model. We first use a semi-parametric mixed effects model to model the k -th PC (or area) trajectory as follows:

$$\tilde{c}_{i,k}(t) = \mu_k + \phi_k(t) + P(t)^T \boldsymbol{\eta}_i + \epsilon_i(t) \quad \text{for } i = 1, \dots, n, \quad (2)$$

where k is the PC index, $\tilde{c}_{i \cdot k}(t_{ij}) = \tilde{c}_{ijk}$, $\phi_k(t)$ is an unknown fixed function of t , $P(t)$ represents the polynomial vector $(1, t, \dots, t^p)^T$, and $\epsilon_i(t)$ is a random noise process with mean zero and variance σ^2 . Moreover, η_i is a $(p+1) \times 1$ vector of random effects. We approximate the fixed effect function $\phi_k(t)$ by using penalized regression spline. The covariance component associated with η_i are estimated using restricted maximum likelihood. Moreover, the fixed effect $\mu_k + \phi_k(t)$ models the mean trajectory for the population and the random effects $P(t)^T \eta_i$ allows for individual variation. The Akaike information criterion is used to select the number of spline basis functions (i.e., p). To fit model (2), we use the *gamm* function provided in publicly available package *mgcv* (Wood, 2012).

The proposed mixed-effects model has many advantages when compared with an independent fit of each trajectory. The model estimates the continuous trajectory $\tilde{c}_{i \cdot k}(t)$ using all observed data rather than just those from the subject i . Therefore, when there is insufficient data for subject i , we can borrow information from all other subjects and still have a reasonable estimate at the individual level. From a theoretical perspective, the maximum likelihood method used to estimate unknown parameters in the model allows different weights to different observations, resulting in estimators with asymptotic optimality properties.

To estimate the covariance component for η_i , we must estimate $(p+1)(p+2)/2$ different parameters. Given the sparse data (sometimes, we only have one to two observations for some individuals), these estimates can be highly variable, and the estimation algorithm may be trapped in local maxima. A possible solution is to employ a more adaptive and representative basis to fit each trajectory. This motivates the use of functional principal component analysis (fPCA) and principal components analysis through conditional expectation (PACE) (Yao et al., 2005).

PACE. The PACE model assumes

$$\tilde{c}_{i \cdot k}(t) = \mu_k(t) + \sum_{p=1}^{\infty} \xi_{ikp} \phi_p(t) + \epsilon_i(t), \quad (3)$$

where $\mu_k(t)$ describes the population mean for the k -th PC trajectory, $\epsilon_i(t)$ is random noise with mean zero and variance σ^2 , and $\sum_{p=1}^{\infty} \xi_{ikp} \phi_p(t)$ models the individual trajectory's deviance from the population mean. Moreover, $\{\phi_p(t)\}$ is the set of basis functions, $\{\xi_{ikp}\}$ is the vector of corresponding coefficients, and $\tau_{i \cdot k} = \mu_k + \sum_{p=1}^{\infty} \xi_{ikp} \phi_p$ denotes the unobserved true k -th PC score trajectory for subject i . The goal of PACE is to estimate $\tau_{i \cdot k}$.

Assuming that $\{\tau_{i \cdot k}\}$ for $i = 1, \dots, n$ are realizations of a stochastic process with the mean function μ_k and covariance function $C_k(\cdot, \cdot)$. Let $C_k(t_1, t_2) = \sum_{p=1}^{\infty} \rho_p \phi_p(t_1) \phi_p(t_2)$ be the eigen-

decomposition of $C_k(\cdot, \cdot)$. By the Karhunen-Loève theorem, with probability one we have

$\tau_{i,k}(t) = \mu_k(t) + \sum_{p=1}^{\infty} \xi_{ip} \phi_p(t)$, indicating that any realization of this stochastic process can be

represented as a linear combination of $\{\phi_p(\cdot)\}$ and coefficients. The basis $\{\phi_p(\cdot)\}$ derived from the eigen-decomposition of $C_k(\cdot, \cdot)$ also facilitates a parsimonious representation of $\tau_{i,k}$ using the first P basis functions (eigenfunctions) in terms of minimum expected mean integrated squared error. That is, if $\{e_p\}_{p=1}^{\infty}$ is a complete orthogonal basis system for

representing any $\tau_{i,k}$, then $\mathbb{E}[\|\tau_{i,k} - \mu_k - \sum_{p=1}^P \langle \tau_{i,k} - \mu_k, e_p \rangle e_p\|^2]$ is minimized by taking

$e_p = \phi_p$ for $p = 1, 2, \dots, P$. PACE utilizes this important theoretical result by estimating an

empirical \hat{C}_k from the given sparse data and utilizing its first P eigenfunctions $\hat{\phi}_p$ to replace

the ϕ_p in (3). Both the empirical population mean $\hat{\mu}_k$ and covariance function \hat{C}_k are estimated using local linear smoothers (Yao et al., 2005; Fan and Gijbels, 2018). To obtain a

good estimate of $\{\xi_{ikp}\}$, PACE assumes that ξ_{ikp} and $\epsilon_i(t)$ are jointly Gaussian. Let

$\tilde{\mathbf{c}}_{ik} = (\tilde{c}_{i1k}, \dots, \tilde{c}_{im,k})^T$, $\boldsymbol{\mu}_{ik} = (\hat{\mu}_k(t_1), \dots, \hat{\mu}_k(t_{m_i}))^T$, and $\boldsymbol{\phi}_{ip} = (\hat{\phi}_p(t_1), \dots, \hat{\phi}_p(t_{m_i}))^T$. Under

Gaussian assumptions, the best estimate of ξ_{ikp} is given by the conditional expectation

$\hat{\xi}_{ikp} = E(\xi_{ikp} | \tilde{\mathbf{c}}_{ik}) = \hat{\rho}_p \boldsymbol{\phi}_{ip}^T \boldsymbol{\Sigma}_{\mathbf{c}_{ik}}^{-1} (\tilde{\mathbf{c}}_{ik} - \boldsymbol{\mu}_{ik})$, where the (a, b) -th element of $\boldsymbol{\Sigma}_{\mathbf{c}_{ik}}^{-1}$ is

$\hat{C}_k(t_a, t_b) + \hat{\sigma}^2 \delta_{ab}$ with $\delta_{ab} = 1$ if $a = b$ and 0 if $a \neq b$. The final estimated trajectory for the k -

th PC score for subject i is given as $\hat{\tau}_{i,k}(t) = \hat{\mu}_k(t) + \sum_{p=1}^P \hat{\xi}_{ikp} \hat{\phi}_p(t)$. The P is selected using the

cross-validation method introduced in Yao et al. (2005). Note that $\hat{\xi}_{ikp}$ is the best estimator

under Gaussian assumptions and best linear prediction of ξ_{ikp} given the information from the i -th subject irrespective of the Gaussian assumptions.

Using either the mixed effects model or the PACE model, we can estimate smooth and

continuous shape trajectories $\hat{f}_i^*(t) = \hat{\alpha}_i(t) * (f_{\mu} + \sum_{k=1}^r \hat{\tau}_{i,k}(t) \tilde{v}_k)$, where $\hat{\alpha}_i(t)$ and $\hat{\tau}_{i,k}$ are,

respectively, the recovered surface area trajectory and the k -th shape PC trajectory for the i -th subject.

Longitudinal Data Analysis. Assuming that there is a set of covariates $\mathbf{x}_i \in \mathbb{R}^K$ from each subject (e.g., gender and disease status), we are interested in learning the effects of \mathbf{x}_i on longitudinal surface trajectories. We refer to this analysis as shape-trajectory-on-scalar regression. Let $f_i^*(t, s)$, in which t indexes the time and s indexes the location on the surface, e.g., $s \in \mathbb{S}^2$, be the shape trajectory. It is assumed that the mean of $f_i^*(t, s)$ is a function of scalar predictive variables, given by

$$E(f_i^*(t, s) | \mathbf{x}_i) = \mu(t, s) + \sum_{j=1}^K x_{ij} \psi_j(t, s), \quad (4)$$

where $\mu(t, \cdot)$ is a 3D shape and $\sum_{j=1}^K x_{ij} \psi_j(t, \cdot)$ deforms $\mu(t, \cdot)$ to the mean of $f_i^*(t, \cdot)$.

With PACE, each sparsely observed shape trajectory $\{f_i^*(t_1, \cdot), \dots, f_i^*(t_{m_i}, \cdot)\}$ is represented as a $rP \times 1$ vector $\mathbf{y}_i = [\hat{\xi}_{i11}, \dots, \hat{\xi}_{i1P}, \dots, \hat{\xi}_{ir1}, \dots, \hat{\xi}_{irP}]^T \in \mathbb{R}^{rP}$, where $[\hat{\xi}_{ik1}, \dots, \hat{\xi}_{ikP}]^T$ comes from the k -th PC score trajectory after applying PACE (or mixed effects model). The elements in \mathbf{y}_i are obtained by two layers of PC analysis, and therefore, they are independent of each other. This nice property significantly reduces the complexity of our regression problem in (4) - we can conduct a simple regression for each element in \mathbf{y}_i separately. With the fitted models, for a given new \mathbf{x}_i , we can easily predict \mathbf{y}_i , the corresponding PC score trajectories, and the shape trajectory.

4 Longitudinal Shape Data Analysis Results

In this section, we carry out a comprehensive data analysis of the three data sets introduced in Section 2 in order to address (Q1) and (Q2).

4.1 Efficient Representation of Surface Shapes in LESA

We compare LESA with spherical harmonic-based point distribution model (SPHARM-PDM) (Styner et al., 2006) in terms of terseness and efficiency of representation. The SPHARM-PDM is a widespread technique in medical shape analysis. The efficiency of a representation is quantified using the number of PCA coefficients needed for representing shapes up to a fixed reconstruction error. Let f be a surface and \hat{f}_r be the reconstructed surfaces with r PCs. We also define the reconstruction error as $\|f - \hat{f}_r\| = \sqrt{\int_{\mathbb{S}} |f(s) - \hat{f}_r(s)|^2 ds}$. For a fix r , the method that leads to smaller reconstruction error is more efficient in shape representation. Figure 6 presents the obtained results for left lateral ventricle surfaces in ADNIGO2 dataset. Similar results are found for other subcortical region and data sets. Figure 6 (a) presents the reconstruction errors of all individual surfaces versus r for the SRNF representation in LESA (blue lines) and SPHARM-PDM (red lines). Figure 6 (b) shows the total distances of all reconstructed surfaces to their original surfaces

under different r , showing that SRNF outperforms SPHARM-PDM in representation efficiency. Figure 6 (c) quantifies this out-performance by the percentage of improvement, indicating that our SRNF framework has much better performance in the sparse cases when only a few PC scores are used to represent the shape.

4.2 From Discrete to Continuous – Fitting Shape Trajectory

We compare PACE with the mixed-effects model (denoted as MGCV from here) in LESA by using them to fit continuous shape trajectories based on the observed discrete data for the left ventricle and hippocampus in ADNIGO2. Figure 7 shows the observed sparse data and the fitted smooth trajectories (with PACE and MGCV) for the surface area (the first row) and PC1 score (the second row). The solid lines in different colors present individual trajectories, whereas the black dashed lines present the mean trajectories. We observe that the mean trajectories fitted by PACE and MGCV follow very similar trajectories, but some individual trajectories fitted by MGCV diverge from the range of observed data significantly, which is probably caused by the high variability of the estimated parameters in MGCV. The third row of Figure 7 illustrates the surface trajectories reconstructed based on

$$f_i(t) = \hat{\alpha}_i(t) * (f_\mu + \sum_{k=1}^r \hat{\tau}_{i,k}(t) \tilde{v}_k), \text{ with } t \in [60, 90] \text{ and } r = 32 \text{ for left ventricle and}$$

64 for left hippocampus, respectively. The surface trajectories built under PACE and MGCV have some agreements with the aging process. Specifically, the left ventricle surface tends to enlarge and its shape mainly deforms at the anterior and posterior ends. The left hippocampus surface tends to shrink and its shape mostly changes in both anterior and posterior ends.

Figure 8 presents trajectory fitting results of three randomly selected individual subjects. Inspecting Figure 8 (a) and (b) reveals that both approaches can capture the patterns of original trajectories and make reasonable predictions. The reconstructed dense individual surface trajectories in panel (c) are also consistent with the raw observations.

Let $\text{MSPE}_\alpha = n^{-1} \sum_{i=1}^n \sqrt{m_i^{-1} \sum_{j=1}^{m_i} \{\alpha_i(t_{ij}) - \hat{\alpha}_i(t_{ij})\}^2}$ be the mean square-root prediction error

(MSPE) of surface area trajectories. Moreover, we also define

$MSPE_{\tau_k} = n^{-1} \sum_{i=1}^n \sqrt{m_i^{-1} \sum_{j=1}^{m_i} \{\hat{\tau}_{i,k}(t_{ij}) - \tilde{c}_{i,k}(t_{ij})\}^2}$ to be the MSPE of PC score trajectories.

To compare trajectory fitting methods, we compute $MSPE_{\alpha}$ and $MSPE_{\tau_k}$, in which we set $n = 1045$ for the left ventricle and $n = 974$ for the left hippocampus. Table 2 shows that results are consistent across the two brain regions: PACE results in better prediction accuracy on the area and PC trajectories. Therefore, in the following data analysis, we will utilize only the PACE method for trajectory fitting.

4.3 Life-span Shape Change

To address (Q1) for the left ventricle and left hippocampus, we integrate the ADNIGO2, HCP test-retest, and the OpenPain datasets into a single data set and then apply LESA to it. Figure 9 shows the observed sparse area trajectories (from all three data sets used in this paper) and their mean trajectories fitted by PACE. The area of the left ventricle keeps increasing after the age of approximately 30 years old. The speed of change is relatively slow before 60 years old, but after 60 years old, the enlargement of the ventricle speeds up. In contrast, the size of the hippocampus reduces with age, while the speed of the shrinking increases after around 60 years old. Figure 9 (c) shows the mean surface trajectories for the two brain regions from the age of 22 to 90 years old. In addition to size change, we observe the changes in shapes due to aging. Specifically, for the left ventricle, the anterior end becomes smoother and fatter with aging, while the posterior end enlarges the most among the whole surface. The whole left hippocampus surface gets thinner with aging, while the anterior and posterior ends atrophy the most.

4.4 Longitudinal Analysis of Shape Trajectories

LESA facilitates simple but effective longitudinal analysis of surface trajectories. We use LESA to analyze the ADNIGO2 data set in order to address questions similar to the Q2 through (i) identifying group differences in longitudinal shape data and (ii) quantifying the contributions of the covariate(s) to the longitudinal shape change.

Group Difference Analysis: In the ADNIGO2 data set, we have three diagnosis groups: AD, mild cognitive impairment (MCI), and normal control (NC). To delineate the group difference, we computed mean trajectories for each of the three groups.

Figure 10 (i) and (ii) presents the mean trajectories of the three groups for the left ventricle and those for the left hippocampus. Within each panel, panel (a) shows the surface area trajectory, panel (b) shows the area changing rate, defined as $100 \times \{\alpha(t_{i+1}) - \alpha(t_i)\} / \alpha(t_i)$, as a description of shape deforming speed with positive numbers representing enlarging and negative numbers representing shrinking, and panel (c) shows the reconstructed shape trajectory. The *shape trajectory* is recovered as $(f_\mu + \sum_{k=1}^r \hat{\tau}_{i-k}(t) \tilde{v}_k)$, which is different from the previous surface trajectory that incorporates the area information. From Figure 10 (i), we observe the following patterns for the left ventricle:

- The AD group has the largest surface area from 60 to 90 years old, followed by the MCI and NC groups.
- The surface area increases with age for all groups, but at different speeds (see Figure 10 panel (b)). Between 60 and 75 years old, the AD group has the largest enlarging speed. The MCI group also enlarges faster than the NC group, but is slower than the AD group. The enlarging speeds of different groups converge at around the age of 75 years old. After 85 years old, due to smaller sample sizes and potential sampling bias, our estimation might have larger variation and so we do not try to interpret it to avoid over-interpretation.
- Most of the shape differences between the AD and NC groups are visible in the anterior and posterior ends. The AD group has the fattest ends, and the NC group has a similar trend in turning fatter with aging but at a much slower pace than the AD group.

From Fig. 10 (ii), for the left hippocampus, we observe the following patterns:

- The AD group has the smallest mean surface area all the time from 60 to 90 years old, followed by the MCI group;
- The surface area tends to shrink with age in most of the time for all three groups. Between 65 and 80 years old, the AD group shrinks the fastest. The shrinkage speeds of different groups converge at around 82 years old. Given the significant acceleration of the AD group's shrinkage speed from

60 to around 70 years old, it seems that the hippocampal atrophy happens much more rapidly for the AD group.

- The atrophy or shrinking happens mainly at the posterior end for all three groups. The posterior end contains a mixture of several essential sub-fields, including CA1, CA1, CA2, and CA4 (DeKraker et al., 2020). The AD group has the sharpest posterior end (severest atrophies), and the NC group deforms the least with aging.

Overall, normal aging, MCI, and AD have a similar effect on the subcortical structure. At 60 years old, the AD group already has a significant shape difference in hippocampi compared with normal controls. However, most subjects were diagnosed with AD after 60 years old, indicating that subcortical brain atrophy may happen long before the clinical diagnosis (Coupé et al., 2019).

Shape-trajectory-on-scalar Regression Analysis:

We are interested in understanding the effects of some predictors of interest on the variability in subcortical shape trajectories by using the ADNIGO2 data set for both the left ventricle and left hippocampus. We included gender, marriage status, education years, diagnostic status (NC, MCI, AD), and ApoE4 type (type 0: e3/e3, type 1: e3/e4, type 2: e4/e4) as covariates of interest. The whole data were split into training (80%) and testing (20%). The training data were used to fit continuous surface trajectories and perform the shape-trajectory-on-scalar regression. We then conducted two sets of analyses: (i) shape prediction accuracy evaluation using the test data; and (ii) controlling the other covariates, only change one covariate to explore its effect on the surface trajectory.

Figure 11 shows the result for the first set of analyses. After training the regression model, we used the covariates in the testing data to predict surface trajectories. To evaluate the prediction accuracy, we defined a metric named average prediction error (APE):

$$APE_i = (1 / m_i) \sum_{j=1}^{m_i} |(f_i(t_{ij}) - \hat{f}_i(t_{ij}))|, \text{ where } f_i(t_{ij}) \text{ is the observed surface at time point } t_{ij} \text{ for}$$

subject i , and $\hat{f}_i(t_{ij})$ is the predicted surface using the regression model. Note that the

regression model only predicts ξ_{ijk} , we need to use these ξ_{ijk} 's to recover the PC score trajectories, and then the surface trajectories. We compared the regression model with a baseline model that uses the mean trajectory f_μ to predict every subject's surface trajectory

\hat{f}_i . Figure 11 (a) presents the percentage of APE improvement compared with the baseline

model. Figure 11 (b) presents some examples of $f_i(t_{ij})$ (original surface), $\hat{f}_i(t_{ij})$

(reconstructed surface based on the regression model), and $f_\mu(t_{ij})$ (mean surface). The

results clearly indicate that the regression model explains part of the variation in the surface trajectories and gives better prediction than the baseline model.

Next, we explore how some covariates of interest would affect the shape trajectory of either the left ventricle or hippocampus. Figure 12 (i) presents some results for the left ventricle. In

the sub-panel (a), we show the predicted area trajectories and PC1 score trajectories by letting gender = female (0), marriage status = married (1), education years = 16, and ApoE type = 1 (e3/e4), and varying the diagnosis status to be AD, MCI, and NC. Comparing the AD group with the MCI and NC groups, we can see that in the AD group the left ventricle tends to be larger and moves along the positive direction of the first shape PC. A similar analysis is performed in the sub-panel (b) by varying the ApoE4 types. We observe that type 2 (with two e4 alleles) has a pretty different effect than types 0 and 1. Specifically, subjects having two e4 alleles tend to have larger left ventricle and more deformation along the first shape PC direction. The sub-panel (c) shows the reconstructed shape trajectories for varying the diagnosis status to be AD, MCI, and NC. We observe that the AD status has more effect on the left ventricle shape trajectory compared with the MCI, and the AD is making the ventricle fatter. Figure 5 (c) presents how an increasing PC1 score changes the ventricle shape.

Figure 12 (ii) shows a similar set of analyses for the left hippocampus. The sub-panel (a) shows how diagnosis status changes the predicted area and PC1 score trajectories. We see that AD makes the hippocampus smaller and changes the shape mainly along the positive PC1 direction. The sub-panel (b) shows how ApoE4 type changes the predicted area and PC1 score trajectories. We observe that the double e4 alleles have a shrinking effect on the hippocampus's size and change the shape along the positive PC1 direction. Existing studies also found that double e4 alleles have significant effect to the volume and shape change in aging (Striepens et al., 2011; O'Dwyer et al., 2012; Li et al., 2016). The sub-panel (c) shows the predicted shape trajectories by varying the diagnosis status, and we see that compared with NC, the AD status contributes to the shape change at the posterior end. Moreover, the difference between NC and MCI is much less the difference between NC and AD.

5 Discussions

This paper introduces a comprehensive LESA framework for statistically analyzing longitudinal brain subcortical regions. LESA contains five major components, including subcortical surface extraction, elastic shape analysis, principal components analysis (PCA) of shapes, continuous shape trajectory fitting, and shape-trajectory-on-scalar regression. We then applied LESA to study the ADNIGO2, HCP, and OpenPain data sets with subjects ranging from 20 to 90 years old and demonstrated several key properties and applications of LESA. First, we illustrated that the elastic shape analysis and PCA in LESA are efficient in creating low-dimensional representations of each shape surface, making statistical modeling much more straightforward. Next, we solved the challenge of estimating a continuous shape trajectory from super sparse longitudinal observations using two advanced functional data analysis techniques - mixed-effects model and PACE. The PACE outperforms the mixed-effects model in the three data sets due to its flexibility and parsimonious representation of the longitudinal data. Another advantage of LESA is that each shape trajectory is eventually represented as a low-dimensional vector with uncorrelated elements (under the PACE model). Consequently, a simple shape-trajectory-on-scalar regression can be developed and applied to study the shape change in ADNI data. The results clearly show that AD has strong adversarial effects on the ventricle and hippocampus.

Applying LESA to the three data sets (totally, 2275 subjects and 9628 shape surfaces), we studied the developmental shape trajectories of left ventricle and left hippocampus in the life-span from 20 to 90 years old. We found that shape change (the atrophy) of these subcortical regions starts very early (~30 years old) and speeds up after 60 years old. Moreover, the AD

further speeds up the atrophy compared with normal aging between 60 and 70 years old. The use of LESA allows us to accurately identify the location of the shape change on the subcortical surfaces. For the left hippocampus (see Figure 10), the atrophy mainly happens at the posterior end, which includes several essential sub-fields, including CA1, CA1, CA2, and CA4 (DeKraker et al., 2020). Moreover, the regression component of LESA estimate the covariates' effect to the surface trajectory. Applying LESA reveals that the AD status and genetic risk (two ApoE4 alleles) all contribute to more severe atrophy of subcortical regions in the aging process.

Although we focus the analysis results of the left ventricle and hippocampus, those for the right ventricle can be found in the Supplementary Materials (Supplementary Figures 5-9). These results are similar to those of left surfaces. In conclusion, LESA is an easy but powerful tool for analyzing longitudinal subcortical surfaces. Implementation of LESA and detailed documentation can be found at <https://wuyx5.github.io/LESA/>.

6 Acknowledgments

The data used in the preparation of this article were obtained from three resources: the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu), the WU-Minn Human Connectome Project (HCP) consortium and the OpenPain Project (OPP). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657), was funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University. The funding for OPP (Principal Investigator: A. Vania Apkarian) was provided by the National Institute of Neurological Disorders and Stroke (NINDS) and National Institute of Drug Abuse (NIDA). OPP data are disseminated by the Apkarian Lab, Physiology Department at the Northwestern University, Chicago.

7 Supplementary Materials

The supplementary materials mainly contain results for subcortical structures on the right side of the brain, e.g., right ventricle and hippocampus, and some additional results for the left ventricle and hippocampus.

References

Amézquita, E. J., Quigley, M. Y., Ophelders, T., Munch, E., and Chitwood, D. H. (2020), "The shape of things to come: Topological data analysis and biology, from molecules to organisms," *Developmental Dynamics*, 249, 816–833.

Basaia, S., Agosta, F., Wagner, L., Canu, E., Magnani, G., Santangelo, R., Filippi, M., and ADNI (2019), "Automated classification of Alzheimer's disease and mild cognitive impairment using a single MRI and deep neural networks," *NeuroImage: Clinical*, 21, 101645.

- Bauer, M. and Bruveris, M. (2011), "A New Riemannian setting for surface registration," in *Proceedings of the 3rd MICCAI Workshop on Mathematical Foundations of Computational Anatomy*, pp. 182–194.
- Coupé, P., Manjón, J. V., Lanuza, E., and Catheline, G. (2019), "Lifespan changes of the human brain in Alzheimer's disease," *Scientific Reports*, 9, 1–12.
- Dai, X., Lin, Z., and Müller, H.-G. (2020), "Modeling sparse longitudinal data on Riemannian manifolds," *Biometrics*, 1–14.
- Dai, X., Müller, H.-G., et al. (2018), "Principal component analysis for functional data on Riemannian manifolds and spheres," *Annals of Statistics*, 46, 3334–3361.
- DeKraker, J., Lau, J. C., Ferko, K. M., Khan, A. R., and Köhler, S. (2020), "Hippocampal subfields revealed through unfolding and unsupervised clustering of laminar and morphological features in 3D BigBrain," *NeuroImage*, 206, 116328.
- Fan, J. and Gijbels, I. (2018), *Local polynomial modelling and its applications: monographs on statistics and applied probability 66*, Routledge.
- Fletcher, P. T. (2013), "Geodesic regression and the theory of least squares on Riemannian manifolds," *International Journal of Computer Vision*, 105, 171–185.
- Glasser, M. F., Smith, S. M., Marcus, D. S., Andersson, J. L., Auerbach, E. J., Behrens, T. E., Coalson, T. S., Harms, M. P., Jenkinson, M., Moeller, S., et al. (2016), "The Human Connectome Project's neuroimaging approach," *Nature Neuroscience*, 19, 1175–1187.
- Howell, B. R., Styner, M. A., Gao, W., Yap, P.-T., Wang, L., Baluyot, K., Yacoub, E., Chen, G., Potts, T., Salzwedel, A., et al. (2019), "The UNC/UMN Baby Connectome Project (BCP): An overview of the study design and protocol development," *NeuroImage*, 185, 891–905.
- Hyun, J. W., Li, Y., Huang, C., Styner, M., Lin, W., Zhu, H., and ADNI (2016), "STGP: Spatio-temporal Gaussian process models for longitudinal neuroimaging data," *Neuroimage*, 134, 550–562.
- James, G. M., Hastie, T. J., and Sugar, C. A. (2000), "Principal component models for sparse functional data," *Biometrika*, 87, 587–602.

Jermyn, I. H., Kurtek, S., Klassen, E., and Srivastava, A. (2012), “Elastic shape matching of parameterized surfaces using square root normal fields,” in *European Conference on Computer Vision (ECCV)*, Springer, pp. 804–817.

Jermyn, I. H., Kurtek, S., Laga, H., and Srivastava, A. (2017), “Elastic shape analysis of three-dimensional objects,” *Synthesis Lectures on Computer Vision*, 12, 1–185.

Kurtek, S., Klassen, E., Ding, Z., Jacobson, S. W., Jacobson, J. L., Avison, M. J., and Srivastava, A. (2010), “Parameterization-invariant shape comparisons of anatomical surfaces,” *IEEE Transactions on Medical Imaging*, 30, 849–858.

Lee, B. C., Tward, D. J., Hu, Z., Trouvé, A., and Miller, M. I. (2020), “Infinitesimal Drift Diffeomorphometry Models for Population Shape Analysis,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 862–863.

Li, B., Shi, J., Gutman, B. A., Baxter, L. C., Thompson, P. M., Caselli, R. J., Wang, Y., and Initiative, A. D. N. (2016), “Influence of APOE genotype on hippocampal atrophy over time—an N= 1925 surface-based ADNI study,” *PloS One*, 11, e0152901.

Madsen, S. K., Gutman, B. A., Joshi, S. H., Toga, A. W., Jack Jr, C. R., Weiner, M. W., Thompson, P. M., and ADNI (2015), “Mapping ventricular expansion onto cortical gray matter in older adults,” *Neurobiology of Aging*, 36, S32–S41.

Miller, M. I., Trouvé, A., and Younes, L. (2002), “On the metrics and Euler-Lagrange equations of computational anatomy,” *Annual Review of Biomedical Engineering*, 4, 375–405.

— (2006), “Geodesic shooting for computational anatomy,” *Journal of Mathematical Imaging and Vision*, 24, 209–228.

Morra, J. H., Tu, Z., Apostolova, L. G., Green, A. E., Avedissian, C., Madsen, S. K., Parikshak, N., Toga, A. W., Jack Jr, C. R., Schuff, N., et al. (2009), “Automated mapping of hippocampal atrophy in 1-year repeat MRI data from 490 subjects with Alzheimer’s disease, mild cognitive impairment, and elderly controls,” *NeuroImage*, 45, S3–S15.

- Muralidharan, P. and Fletcher, P. T. (2012), "Sasaki metrics for analysis of longitudinal data on manifolds," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1027–1034.
- O'Dwyer, L., Lamberton, F., Matura, S., Tanner, C., Scheibe, M., Miller, J., Rujescu, D., Prvulovic, D., and Hampel, H. (2012), "Reduced hippocampal volume in healthy young ApoE4 carriers: an MRI study," *PloS One*, 7, e48895.
- Patenaude, B., Smith, S. M., Kennedy, D. N., and Jenkinson, M. (2011), "A Bayesian model of shape and appearance for subcortical brain segmentation," *NeuroImage*, 56, 907–922.
- Petersen, R. C., Aisen, P., Beckett, L. A., Donohue, M., Gamst, A., Harvey, D. J., Jack, C., Jagust, W., Shaw, L., Toga, A., et al. (2010), "Alzheimer's disease neuroimaging initiative (ADNI): clinical characterization," *Neurology*, 74, 201–209.
- Pizer, S. M., Fletcher, P. T., Joshi, S., Thall, A., Chen, J. Z., Fridman, Y., Fritsch, D. S., Gash, A. G., Glotzer, J. M., Jiroutek, M. R., et al. (2003), "Deformable m-reps for 3D medical image segmentation," *International Journal of Computer Vision*, 55, 85–106.
- Qiu, A. and Miller, M. I. (2008), "Multi-structure network shape analysis via normal surface momentum maps," *NeuroImage*, 42, 1430–1438.
- Shen, L., Farid, H., and McPeck, M. A. (2009), "Modeling three-dimensional morphological structures using spherical harmonics," *Evolution: International Journal of Organic Evolution*, 63, 1003–1016.
- Shi, J., Stonnington, C. M., Thompson, P. M., Chen, K., Gutman, B., Reschke, C., Baxter, L. C., Reiman, E. M., Caselli, R. J., Wang, Y., et al. (2015), "Studying ventricular abnormalities in mild cognitive impairment with hyperbolic Ricci flow and tensor-based morphometry," *NeuroImage*, 104, 1–20.
- Singh, N., Vialard, F.-X., and Niethammer, M. (2015), "Splines for diffeomorphisms," *Medical Image Analysis*.
- Striepens, N., Scheef, L., Wind, A., Meiberth, D., Popp, J., Spottke, A., Kölsch, H., Wagner, M., and Jessen, F. (2011), "Interaction effects of subjective memory impairment and ApoE4 genotype on episodic memory and hippocampal volume," *Psychological Medicine*, 41, 1997–2006.

Styner, M., Oguz, I., Xu, S., Brechbuhler, C., Pantazis, D., Levitt, J. J., Shenton, M. E., and Gerig, G. (2006), "Framework for the Statistical Shape Analysis of Brain Structures using SPHARM-PDM," *Insight J*, 242–250.

Tang, X., Holland, D., Dale, A. M., Younes, L., Miller, M. I., and ADNI (2015), "The diffeomorphometry of regional shape change rates and its relevance to cognitive deterioration in mild cognitive impairment and Alzheimer's disease," *Human Brain Mapping*, 36, 2093–2117.

Thompson, P. M., Hayashi, K. M., De Zubicaray, G. I., Janke, A. L., Rose, S. E., Semple, J., Hong, M. S., Herman, D. H., Gravano, D., Doddrell, D. M., et al. (2004), "Mapping hippocampal and ventricular change in Alzheimer disease," *NeuroImage*, 22, 1754–1766.

Tward, D. J., Sicat, C. S., Brown, T., Bakker, A., Gallagher, M., Albert, M., Miller, M., and ADNI (2017), "Entorhinal and transentorhinal atrophy in mild cognitive impairment using longitudinal diffeomorphometry," *Alzheimer's & Dementia: Diagnosis, Assessment & Disease Monitoring*, 9, 41–50.

Vachon-Presseau, E., Tétreault, P., Petre, B., Huang, L., Berger, S. E., Torbey, S., Baria, A. T., Mansour, A. R., Hashmi, J. A., Griffith, J. W., et al. (2016), "Corticolimbic anatomical characteristics predetermine risk for chronic pain," *Brain*, 139, 1958–1970.

van Timmeren, J. E., Cester, D., Tanadini-Lang, S., Alkadhi, H., and Baessler, B. (2020), "Radiomics in medical imaging—"how-to" guide and critical reflection," *Insights Into Imaging*, 11, 1–16.

Wang, Y., Zhang, J., Gutman, B., Chan, T. F., Becker, J. T., Aizenstein, H. J., Lopez, O. L., Tamburo, R. J., Toga, A. W., and Thompson, P. M. (2010), "Multivariate tensor-based morphometry on surfaces: application to mapping ventricular abnormalities in HIV/AIDS," *NeuroImage*, 49, 2141–2157.

Weiner, M. W., Veitch, D. P., Aisen, P. S., Beckett, L. A., Cairns, N. J., Green, R. C., Harvey, D., Jack, C. R., Jagust, W., Liu, E., et al. (2013), "The Alzheimer's Disease Neuroimaging Initiative: a review of papers published since its inception," *Alzheimer's & Dementia*, 9, e111–e194.

Weiner, M. W., Veitch, D. P., Aisen, P. S., Beckett, L. A., Cairns, N. J., Green, R. C., Harvey, D., Jack Jr, C. R., Jagust, W., Morris, J. C., et al. (2017), "The Alzheimer's Disease

Neuroimaging Initiative 3: Continued innovation for clinical trial improvement,” *Alzheimer’s & Dementia*, 13, 561–571.

Wood, S. (2012), “MGCV: Mixed GAM Computation Vehicle with GCV/AIC/REML smoothness estimation,” *Software Package*.

Yao, F., Müller, H.-G., and Wang, J.-L. (2005), “Functional data analysis for sparse longitudinal data,” *Journal of the American Statistical Association*, 100, 577–590.

Younes, L. (2010), *Shapes and Diffeomorphisms*, Springer.

Zhang, Z., Klassen, E., and Srivastava, A. (2018a), “Phase-amplitude separation and modeling of spherical trajectories,” *Journal of Computational and Graphical Statistics*, 27, 85–97.

Zhang, Z., Su, J., Klassen, E., Le, H., and Srivastava, A. (2018b), “Rate-invariant analysis of covariance trajectories,” *Journal of Mathematical Imaging and Vision*, 60, 1306–1323.

Zhao, Q., Pizer, S., Niethammer, M., and Rosenman, J. (2014), “Geometric-feature-based spectral graph matching in pharyngeal surface registration,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, pp. 259–266.

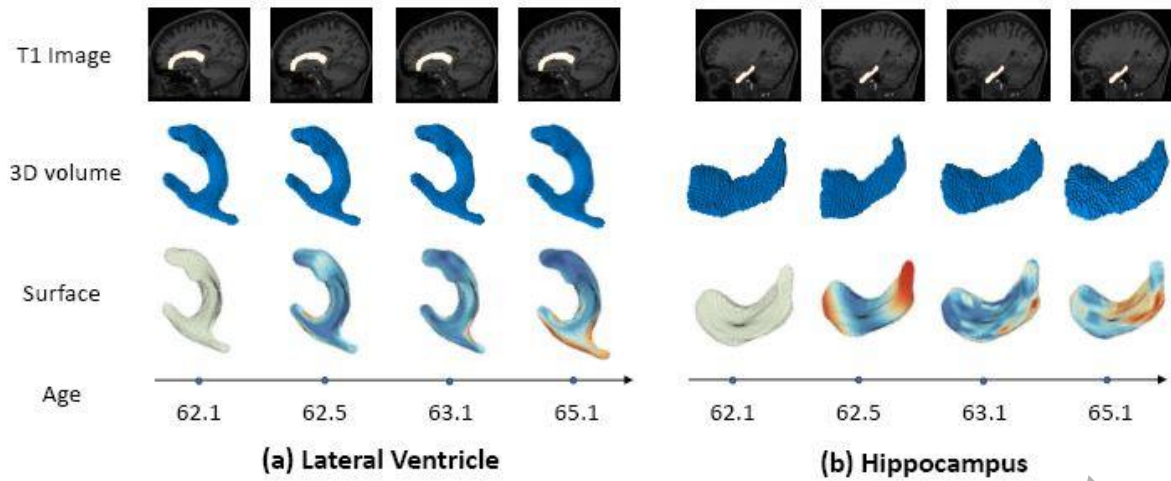


Fig. 1 An example of three different representations of lateral ventricle and hippocampus across four time points for a randomly selected subject. The first row shows the repeated MRI data with segmented lateral ventricle and hippocampus and the second and third rows show their 3D volumes and 3D surface shapes across time, respectively.

Accepted Manuscript

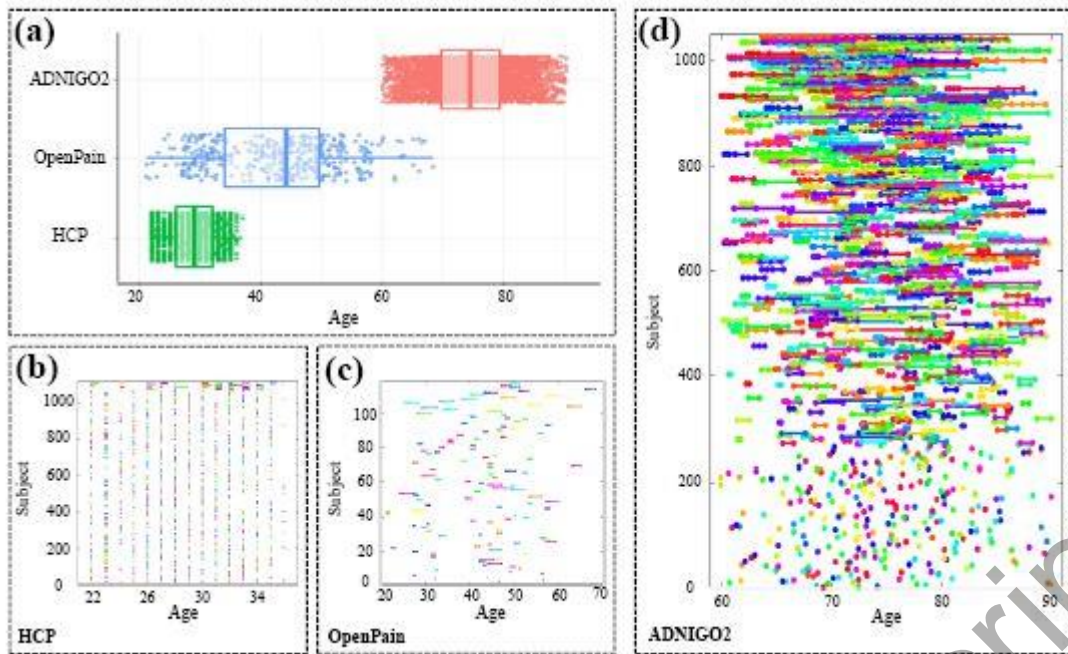


Fig. 2 Panel (a) shows the age distributions of the ADNIGO2, HCP, and OpenPain data sets. Panels (b), (c), and (d) show the temporal information on scans for each subject in the HCP, OpenPain, and ADNIGO2 data sets, respectively.

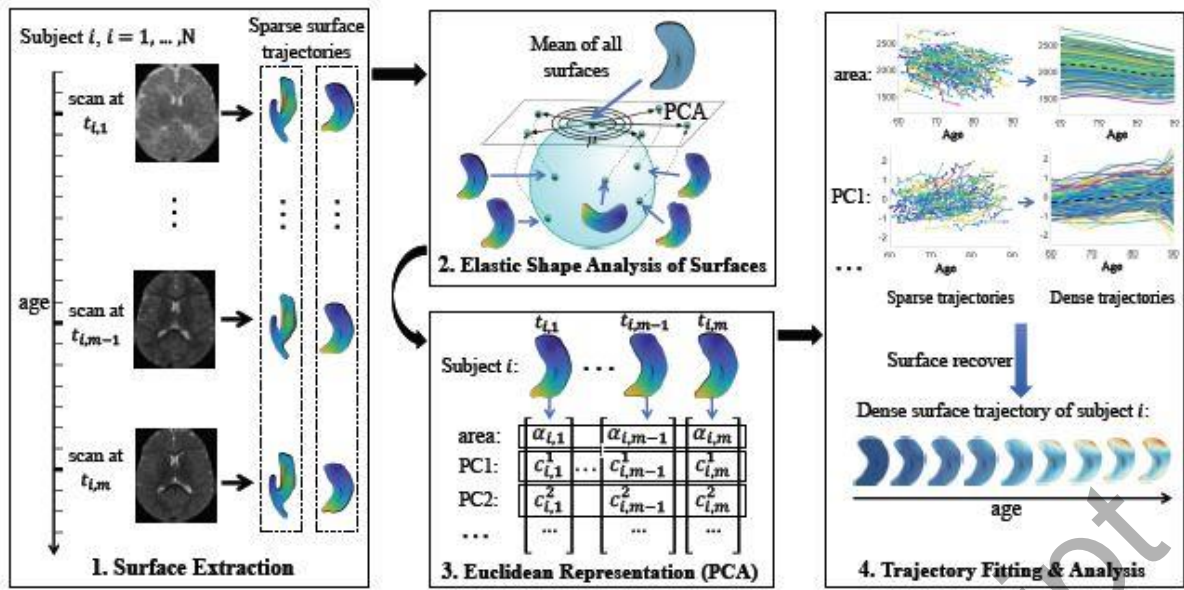


Fig. 3 A schematic overview of LESA consisting of four key components: 1. surface extraction, 2. elastic shape analysis of surfaces, 3. Euclidean representation of shapes, and shape trajectory fitting and regression analysis.

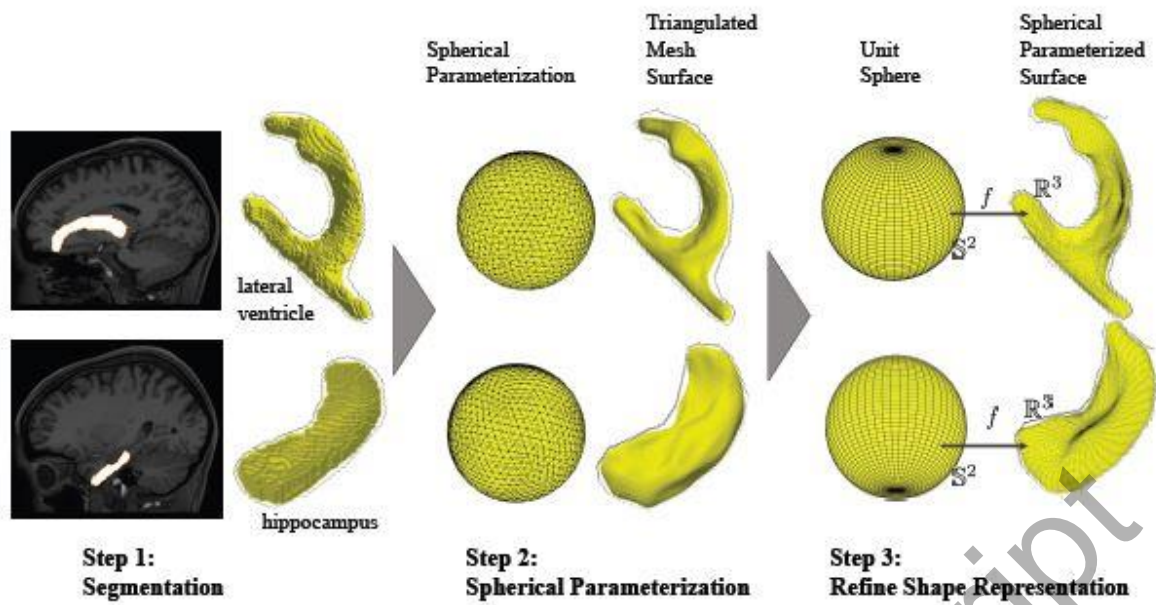


Fig. 4 An illustration on how to extract a parameterized subcortical shape in LESA.

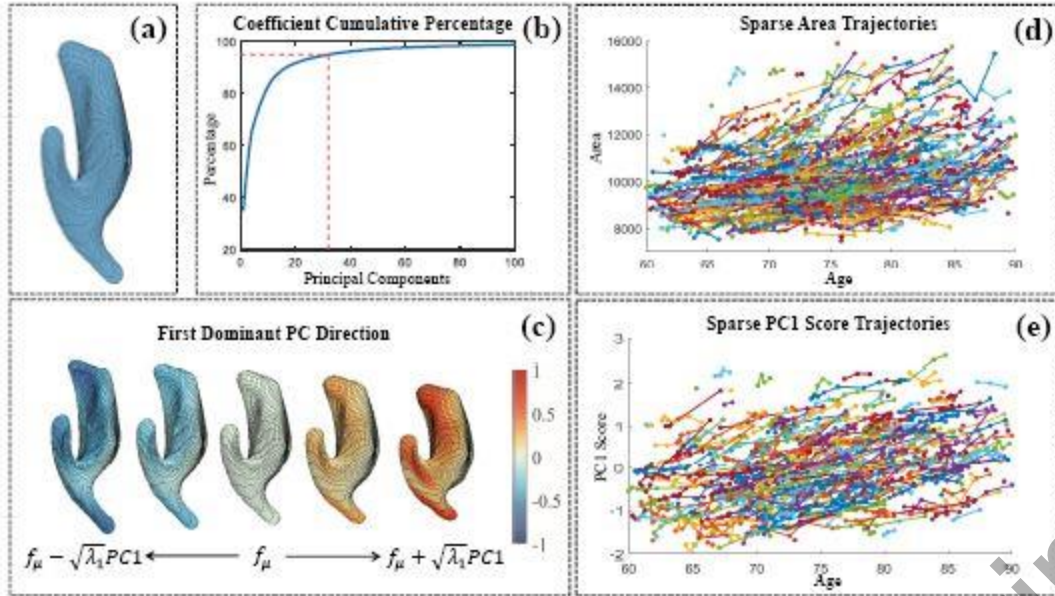


Fig. 5 The PCA results of the ADNIGO2's left ventricle surfaces: (a) the Karcher mean of all left ventricle surfaces; (b) the cumulative percentage of variance explained by the number of PCs; (c) the first dominant PC direction reconstructed as $f_{\mu} + t\sqrt{\lambda_1}PC_1$, in which the five shapes in the front view from left to right correspond to $t = \{-1, -0.5, 0, 0.5, 1\}$ and the color denotes the relative shape change; (d) surface area trajectories, and (e) PC_1 score trajectories.

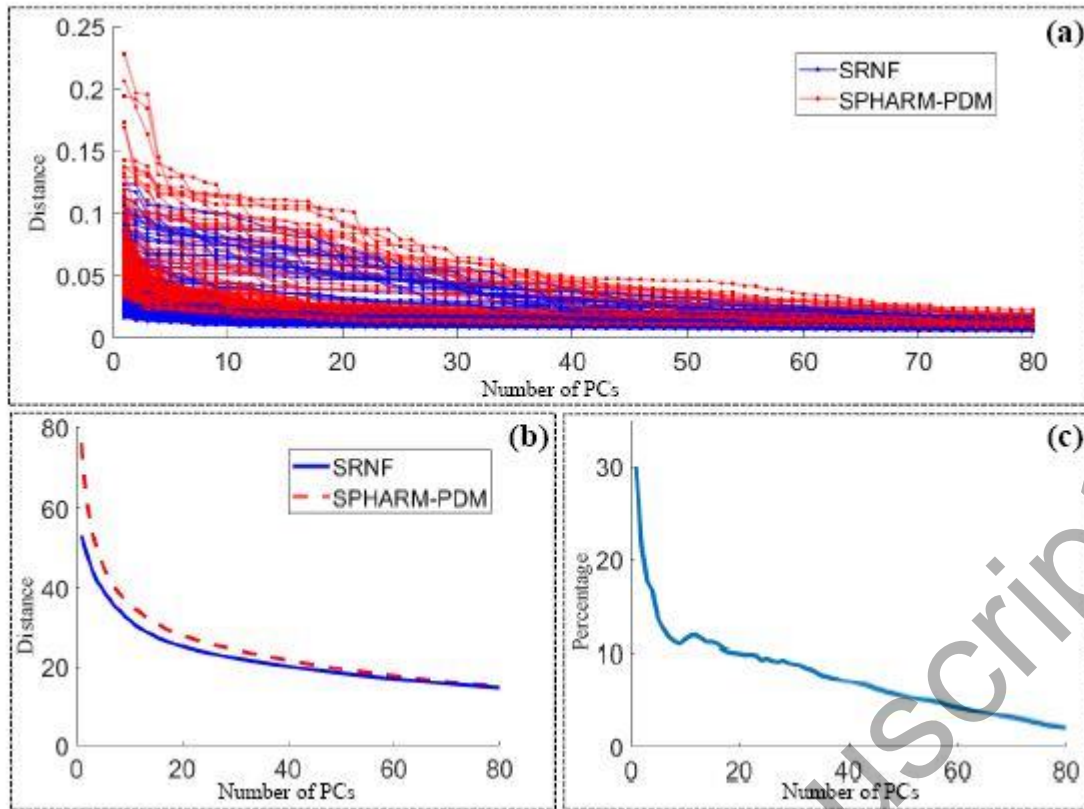


Fig. 6 Comparison of the representation efficiency of the SRNF framework in LESA to that of SPHARM-PDM using left lateral ventricle surfaces from ADNIGO2. (a) Individual surface reconstruction error versus the number of PCs. (b) Total reconstruction error of all surfaces versus the number of PCs, and (c) percentage of performance improvement of our SRNF method over the SPHARM-PDM in Styner et al. (2006).

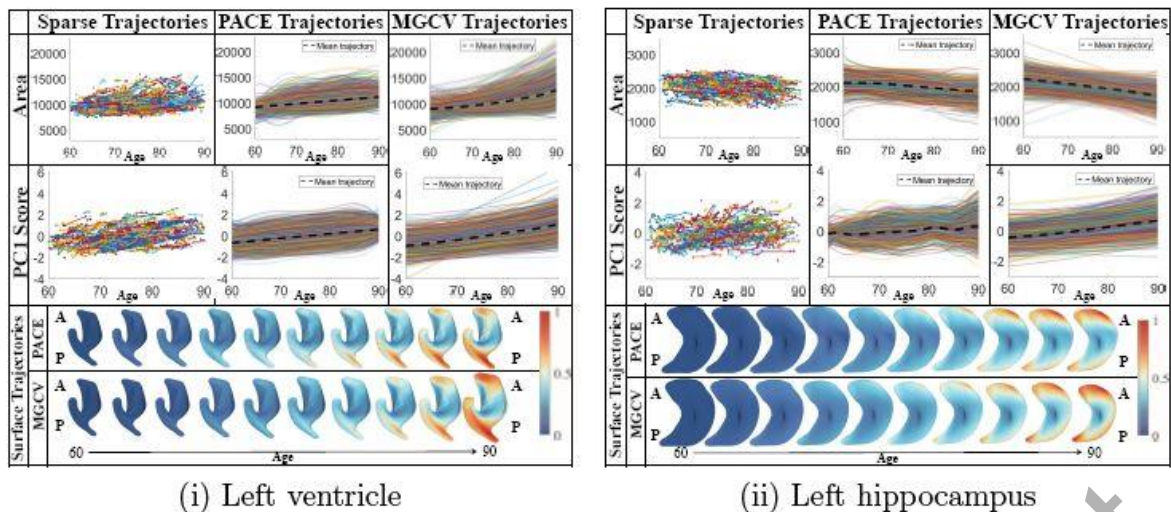


Fig. 7 Trajectory fitting results of LESA from the observed sparse data in ADNIGO2. In the first two rows, the first column shows observed sparse trajectories (for area and PC1 score), and the second and third columns show the continuous trajectories fitted by PACE and MGCV, respectively. In the third and fourth rows, we show the reconstructed mean surface trajectories fitted by PACE and MGCV, respectively, in which color indicates the shape deformation compared with the first shape and we use abbreviations, including A - anterior and P - posterior.

Accepted Manuscript

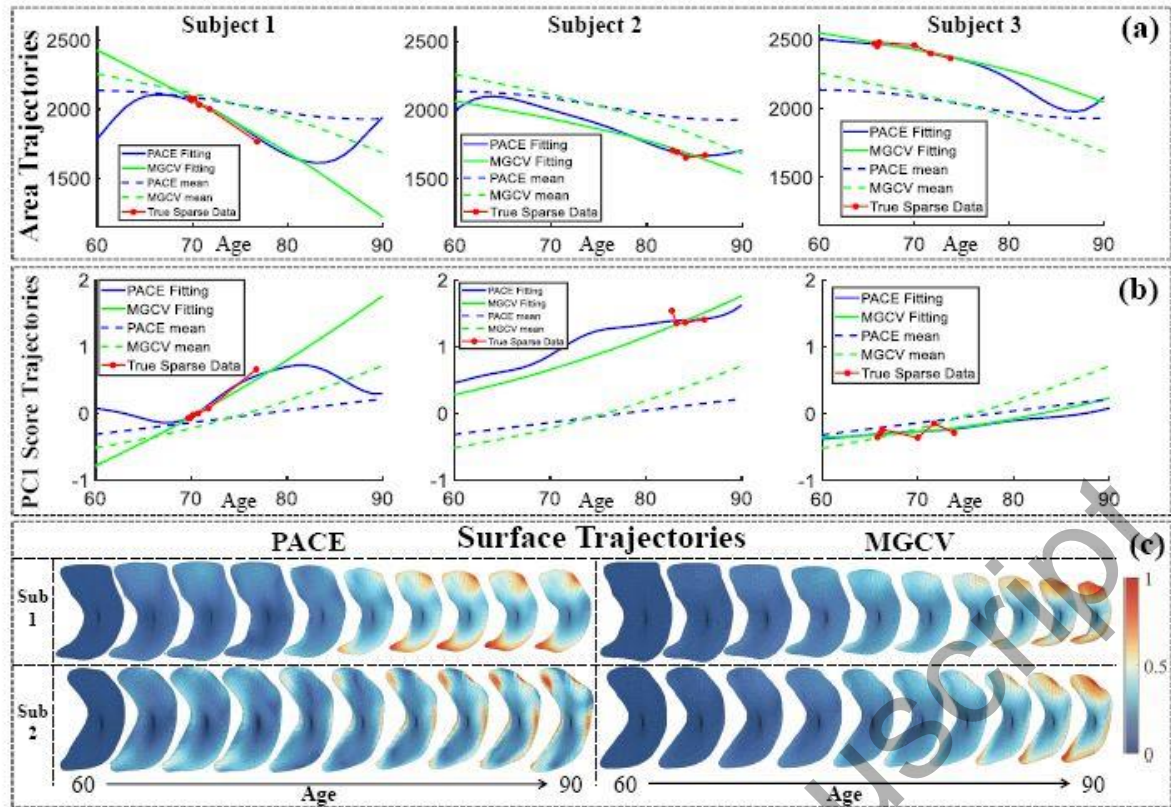


Fig. 8 Individual surface trajectories fitted with LESA in ADNIGO2. Panels (a) and (b) show the raw and fitted trajectories for the surface area and PC1 score, respectively. Panel (c) illustrates the reconstructed surface trajectories based on the fitted surface area and PC score trajectories.

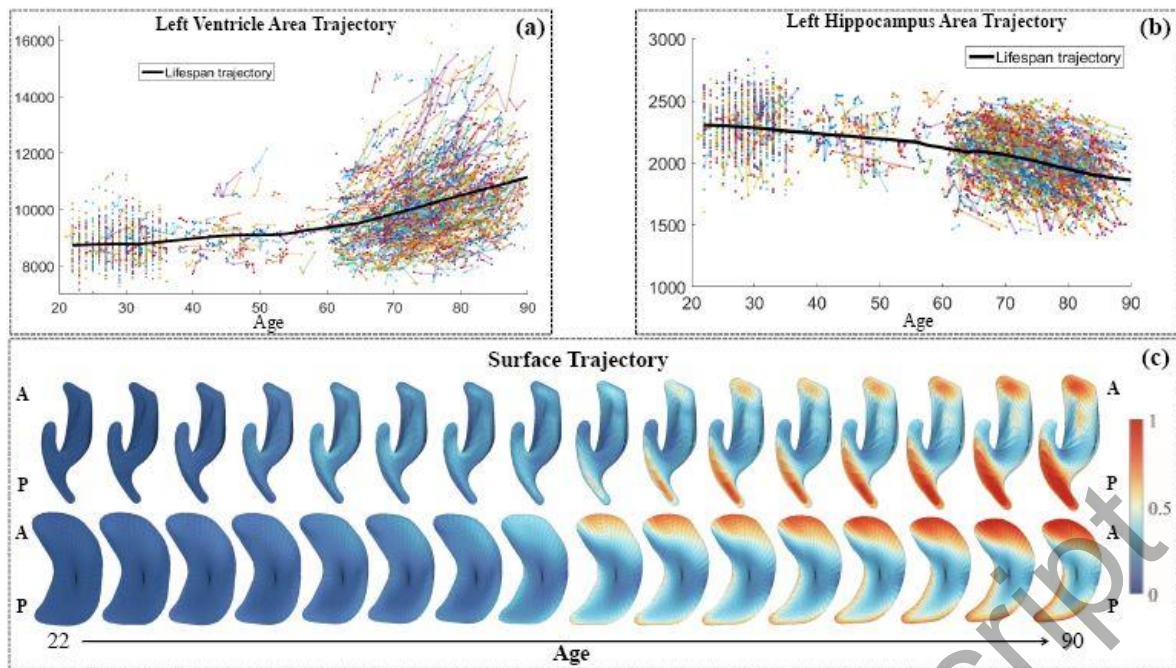


Fig. 9 The life-span growth trajectories from 22 to 90 years old for the left ventricle and left hippocampus. Panels (a) and (b) show the observed sparse data and fitted mean trajectories (black solid line), respectively; and panel (c) shows the reconstructed life-span mean surface trajectories. Color on each surface indicates the surface's deformation size compared with the surface at age 22 years old.

Accepted Manuscript

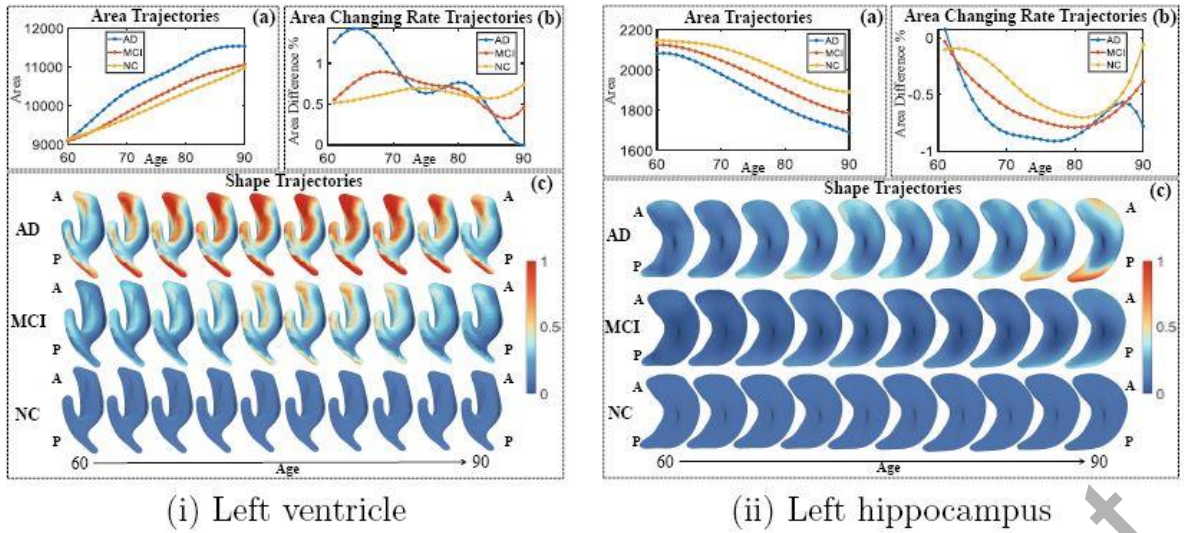
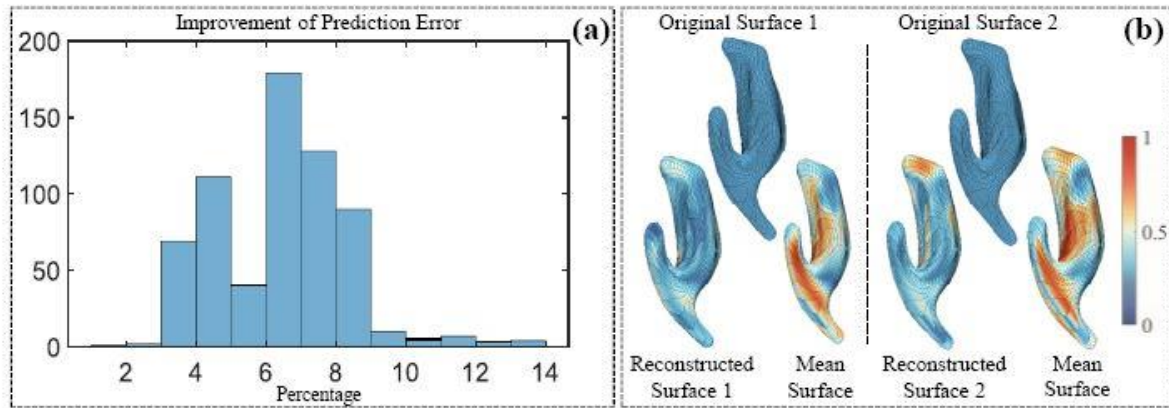
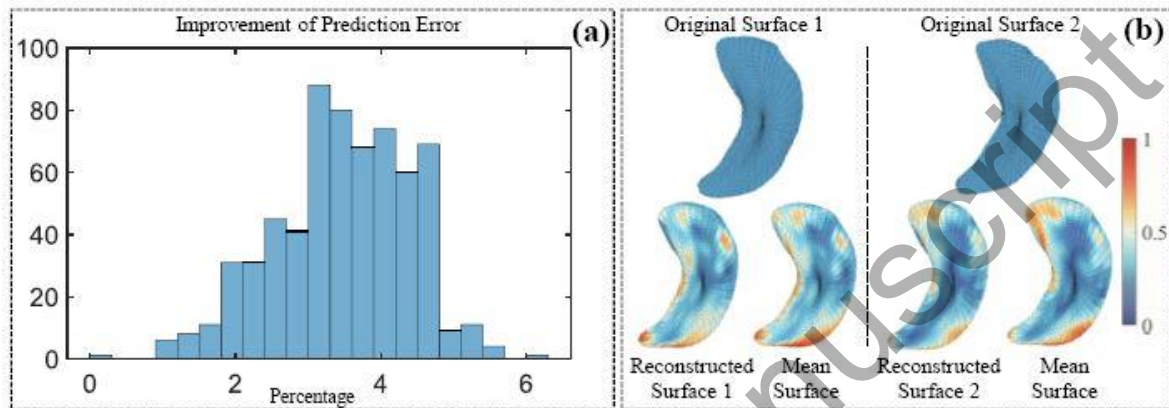


Fig. 10 Comparisons of shape change patterns among AD, MCI and NC using ADNIGO2 data. (a) Mean surface area trajectories of the three groups (blue: AD; red: MCI; yellow: NC); (b) Changing rate of the area trajectories-calculated as $100 * (\alpha(t_{i+1}) - \alpha(t_i)) / \alpha(t_i)$; and (c) reconstructed mean shape trajectories $(f_\mu + \sum_{k=1}^r \hat{\tau}_{i,k}(t) \tilde{v}_k)$. Color on the surface represents shape difference compared with the NC surface at the corresponding time point.

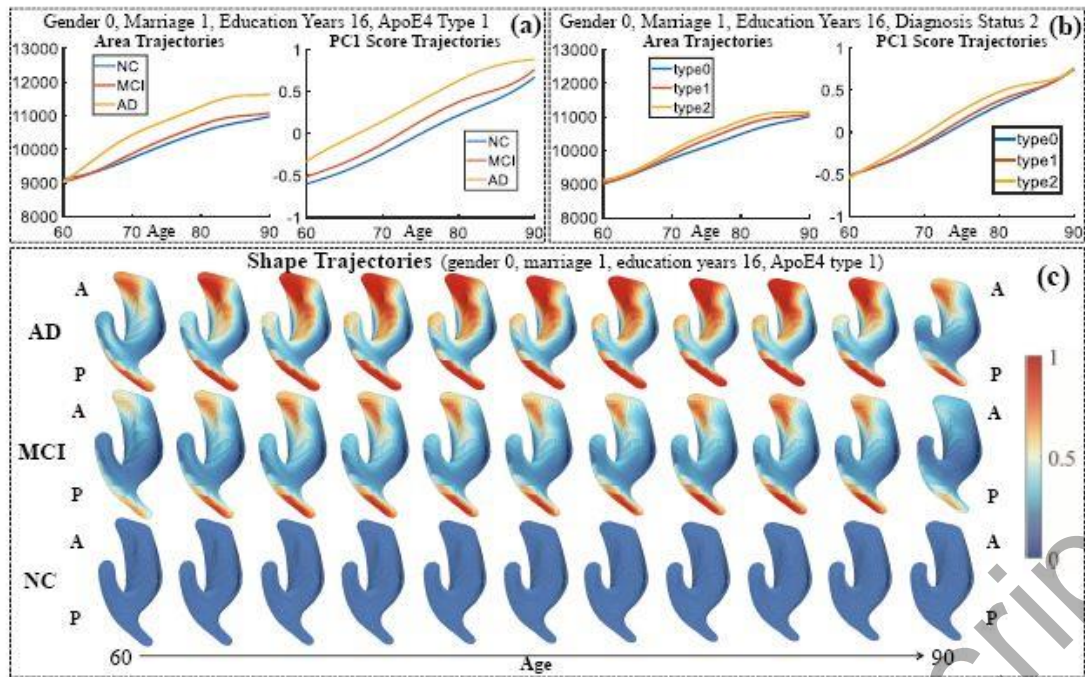


(i) Left ventricle

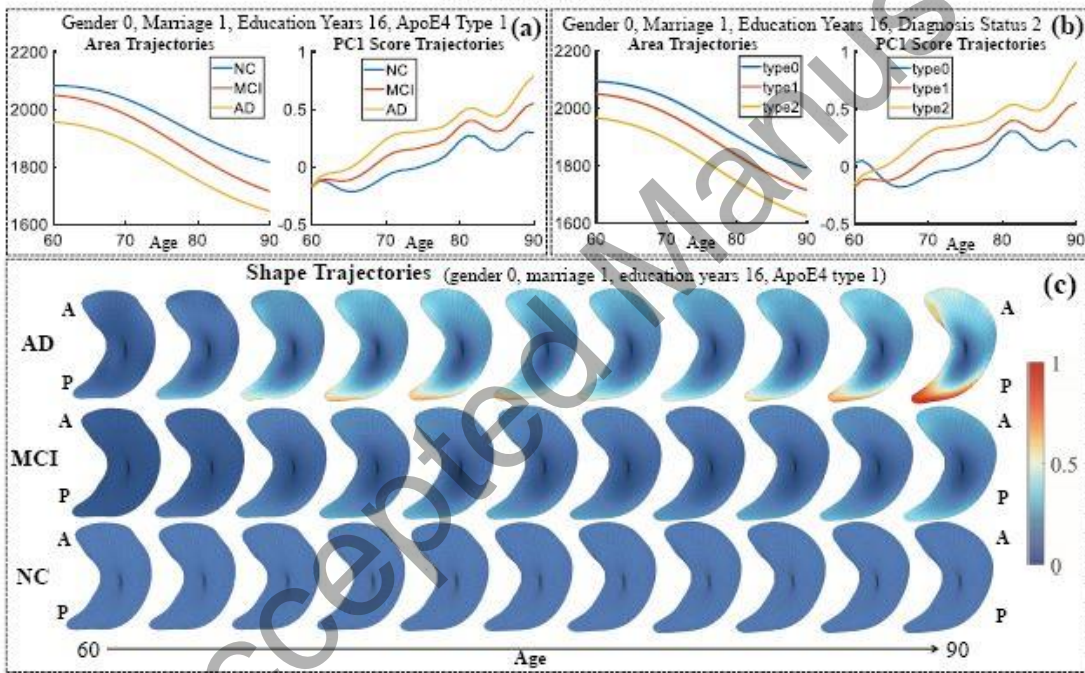


(ii) Left hippocampus

Fig. 11 Evaluation of the shape-trajectory-on-scalar regression in the ADNIGO2 data. Panels (i) and (ii) show results for the left ventricle and the left hippocampus, respectively. Each sub-panel (a) shows the histogram of the percentage of improvement in prediction error when comparing the shape-trajectory-on-scalar regression with the baseline model. Each sub-panel (b) shows some examples of original sparse surface, surface reconstructed by the regression's prediction, and the global mean surface. The color on the reconstructed and mean surfaces indicates their difference to the original surface.



(i) Left ventricle



(ii) Left hippocampus

Fig. 12 Exploration of the covariates' effect to the surface trajectory in the ADNIGO2 data set. In each sub-panel (a), we fixed gender, marriage status, education years and ApoE4 type and varied the diagnosis status. In each sub-panel (b), we fixed the others and varied the ApoE4 type. Each sub-panel (c) shows the reconstructed shape trajectory by only varying the diagnosis status. Color on each surface represents shape deformation compared with the NC surface at the same age.

Table 1 Summary characteristics of the three data sets included in our study.

Regions	Dataset Name	Subject Number	Scan Number	Age Range (Median) (Years)	Gender Ratio (M/F)
Lateral Ventricle	ADNIGO2	1045	3443	[60,90](74.5)	547/498
	HCP	1113	1158	[22,37](29)	503/610
	OpenPain	117	429	(21,69)(44.1288)	65/52
Hippocampus	ADNIGO2	974	3044	[60,90](74.2)	491/483
	HCP	1082	1125	[22,37](29)	486/596
	OpenPain	117	429	(21,69)(44.1288)	65/52

Accepted Manuscript

Table 2 Mean square-root prediction errors (MSPEs) of PACE and MGCV in ADNIGO2

	PACE	MGCV
Area	59.7086	70.537 5
PC1	0.0357	0.0424
PC2	0.0248	0.0256
PC3	0.0448	0.0526
PC4	0.0208	0.0225
PC5	0.0531	0.0580
...
PC Average	0.0383	0.0400
(a) Left ventricle		
	PACE	MGCV
Area	17.7408	22.786 5
PC1	0.0238	0.0239
PC2	0.0462	0.0474
PC3	0.0316	0.0383
PC4	0.0695	0.0852
PC5	0.0272	0.0293
...
PC Average	0.0350	0.0359
(b) Left hippocampus		