Computers in Biology and Medicine 163 (2023) 107185

Contents lists available at ScienceDirect

Computers in Biology and Medicine

journal homepage: www.elsevier.com/locate/compbiomed

Automatic landmark detection and registration of brain cortical surfaces via quasi-conformal geometry and convolutional neural networks

Yuchen Guo [a], Qiguang Chen [a], Gary P.T. Choi [b], Lok Ming Lui [a,*]

[a] Department of Mathematics, The Chinese University of Hong Kong, Hong Kong
[b] Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

Keywords:
Landmark detection
Surface registration
Human brain mapping
Quasi-conformal geometry
Convolutional neural networks

ABSTRACT

In medical imaging, surface registration is extensively used for performing systematic comparisons between anatomical structures, with a prime example being the highly convoluted brain cortical surfaces. To obtain a meaningful registration, a common approach is to identify prominent features on the surfaces and establish a low-distortion mapping between them with the feature correspondence encoded as landmark constraints. Prior registration works have primarily focused on using manually labeled landmarks and solving highly nonlinear optimization problems, which are time-consuming and hence hinder practical applications. In this work, we propose a novel framework for the automatic landmark detection and registration of brain cortical surfaces using quasi-conformal geometry and convolutional neural networks. We first develop a landmark detection network (LD-Net) that allows for the automatic extraction of landmark curves given two prescribed starting and ending points based on the surface geometry. We then utilize the detected landmarks and quasi-conformal theory for achieving the surface registration. Specifically, we develop a coefficient prediction network (CP-Net) for predicting the Beltrami coefficients associated with the desired landmark-based registration and a mapping network called the disk Beltrami solver network (DBS-Net) for generating quasi-conformal mappings from the predicted Beltrami coefficients, with the bijectivity guaranteed by quasi-conformal theory. Experimental results are presented to demonstrate the effectiveness of our proposed framework. Altogether, our work paves a new way for surface-based morphometry and medical shape analysis.

1. Introduction

Surface registration, the process of finding a 1–1 correspondence between surfaces, has been extensively studied and widely applied to various fields in science, engineering, and medicine. For instance, in medical imaging, surface registration facilitates the comparison between complicated anatomical shapes for disease analysis [1,2]. To ensure the accuracy of surface registration, landmarks representing salient features on the surfaces are frequently used to guide the registration. For relatively simple shapes like human faces, prominent features such as the eyes can be easily identified as landmarks. However, for more complicated shapes such as the highly convoluted brain cortical surfaces, landmark extraction usually requires manual delineations by medical experts, which makes the task much more time-consuming. Besides the correspondence between the labeled landmarks, it is also essential to find an accurate mapping between the overall surfaces with low distortion, which is traditionally done by solving some optimization problems [3–6] and is computationally expensive.

Recently, deep learning has emerged as a powerful tool for various complex tasks. To train a suitable model, a large amount of data are usually required. For image registration, it is easy to get millions of images for training and hence convolutional neural networks (CNNs) have achieved huge success. By contrast, surface registration is more complicated as it involves irregular triangular meshes. In particular, due to the lack of data for training, one usually has to augment datasets manually and consider more complicated networks.

To overcome the above-mentioned issues, in this paper we propose a novel framework for solving the automatic landmark detection and registration problems for brain cortical surfaces by combining quasi-conformal geometry and CNNs (see Fig. 1 for an overview). The contributions of our work are as follows:

(i) We apply quasi-conformal geometry and CNNs for mapping triangular meshes with disk topology.
(ii) Feature landmark curves can be automatically extracted from the given surfaces based on only the labeled starting and ending points of them and the surface curvature.
(iii) Beltrami coefficients can be automatically generated based on the landmark constraints.

* Corresponding author.
E-mail address: lmlui@math.cuhk.edu.hk (L.M. Lui).

https://doi.org/10.1016/j.compbiomed.2023.107185
Received 5 December 2022; Received in revised form 24 May 2023; Accepted 13 June 2023
Available online 17 June 2023

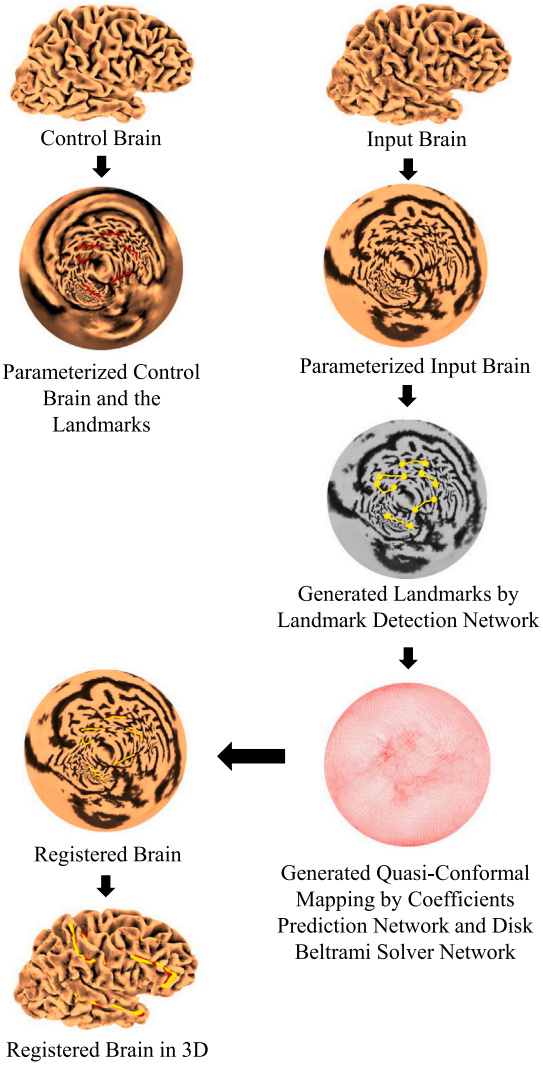

Fig. 1. An overview of our proposed framework. Given a control surface with prescribed landmarks, we first conformally parameterize it onto the unit disk. Then, for any given surface, we compute the disk conformal parameterization of it and use the landmark detection network (LD-Net) to extract the feature landmark curves. We then use the coefficient prediction network (CP-Net) and the disk Beltrami solver network (DBS-Net) to register the disk parameterization with that of the control surface, thereby getting the final registration result between the input surface and the control surface.

(iv) Quasi-conformal mappings can be automatically generated based on the Beltrami coefficients in real time, thereby yielding a landmark-based surface registration.

(v) Once the networks are trained separately, the entire framework can be applied for different input brain surfaces without any additional training.

2. Related works

Surface parameterization and mapping algorithms have been widely studied in recent decades [7,8]. In particular, as conformal parameterizations preserve the local geometry well and can be applied to various fields, many conformal parameterization algorithms have been developed by different research groups [9–12] (see [13] for a survey), with some works focusing on parameterizing brain cortical surfaces [14–18].

More recently, quasi-conformal theory has been found useful for surface parameterization with a wide range of applications [19], such as improving the accuracy of conformal parameterizations and achieving trade-offs between conformality and other prescribed constraints for more flexible parameterizations and registrations. In [20], Lui et al. introduced the use of the Beltrami coefficients and established a 1–1 correspondence between Beltrami coefficients and surface diffeomorphisms. The Beltrami coefficients capture essential properties of the mappings and represent them in a simpler space. Based on the Beltrami coefficients, many efficient conformal and quasi-conformal mapping methods have been developed [21–24].

Landmark-based registration, which aims to obtain a 1–1 correspondence between surfaces with prescribed feature landmarks, has been widely studied [25–29]. In particular, there have been multiple works on using landmark-based quasi-conformal mappings for the analysis of biomedical shapes including teeth [30,31] and human faces [32,33]. For the highly-convoluted brain cortical surfaces, earlier methods primarily focused on landmark-free registration [34–36]. However, it was shown that landmark-free methods often produce misalignment of cortical landmarks [37], which hinders the accurate analysis of certain regions of interest on the cortical surfaces in practice. Therefore, several landmark-constrained mapping methods for brain cortical surfaces have been developed over the past few decades [18,38–40].

Convolutional neural networks (CNNs), a powerful tool extensively studied and used in recent years, perform excellently in imaging science. In 1989, backpropagation was introduced to the network training process and applied to handwritten digit recognition [41]. Nowadays, with the improvement of GPUs and advanced network structures such as ResNet [42], researchers are able to train deeper networks to tackle more complicated problems. Besides the improvement of the network structures, large training datasets such as ImageNet [43] have facilitated the design of more robust networks. More recently, the combination of quasi-conformal geometry and CNNs has been explored [44,45]. In particular, Chen et al. [44] developed a network for generating quasi-conformal maps based on any input Beltrami coefficient for rectangular images.

In recent decades, there has been a growing interest in applying networks for surface registration. One common approach is to treat the surface meshes as graphs and use graph neural networks (GNNs) [46–49]. Besson et al. [50] used GNNs for analyzing brain meshes and predicting sex and age. Besides, customized kernels have also been used to conduct convolutions on the surfaces. In [51,52], Masci et al. used kernels defined on Riemannian manifolds for learning shape correspondence. It is also possible to parameterize the surfaces and perform the network learning in the parameter domain [53–55]. Zhao et al. [56,57] parameterized genus-0 cortical surfaces onto the sphere and developed methods for diffeomorphic spherical surface registration using spherical CNNs in both supervised and unsupervised ways. Cheng et al. [58] developed a learning-based registration framework for cortical surfaces with spherical topology using 2D planar projection and VoxelMorph [59].

The extraction of feature landmarks from complex anatomical shapes is a highly challenging problem. In particular, sulci, which are shallow furrows on the brain surface separating adjacent convolutions, have been extensively studied in medical analysis. In prior brain cortical surface mapping works, the sulcal landmarks (feature points representing the sulci) are usually manually delineated [60–63]. For the automatic extraction of sulci, existing approaches have used pattern recognition system [64,65], supervised learning [66], geometric algorithm [67], geodesic curvature flow [68], geodesic path density map [69], Dense Individualized and Common Connectivity-based Cortical Landmark (DICCCOL)-based method [70] etc.

3. Mathematical background

3.1. Quasi-conformal theory

Conformal maps are angle-preserving homeomorphisms between Riemann surfaces. As they may not exist with the presence of landmark constraints in general, we consider a generalization of them

with bounded conformality distortion known as quasi-conformal maps. Mathematically, $f : \mathbb{C} \to \mathbb{C}$ is *quasi-conformal* if it satisfies the Beltrami equation

$$\frac{\partial f}{\partial \bar{z}} = \mu(z)\frac{\partial f}{\partial z}, \tag{1}$$

where $\mu : \mathbb{C} \to \mathbb{C}$ is a complex-valued function called the *Beltrami coefficient* with $\|\mu\|_\infty < 1$. Infinitesimally, we can also express f by its local parameter around a point p:

$$\begin{aligned} f(z) &\approx f(p) + f_z(p)(z-p) + f_{\bar{z}}(p)\overline{z-p} \\ &= f(p) + f_z(p)((z-p) + \mu(p)\overline{z-p}), \end{aligned} \tag{2}$$

from which we can easily see that f is conformal around a small neighborhood of p if and only if $\mu(p) = 0$. It also shows that f can be expressed as the sum of $f(p)$ and a stretch map $(z-p) + \mu(p)\overline{z-p}$ multiplied by $f_z(p)$, and hence it maps infinitesimal circles to infinitesimal ellipses. All the quasi-conformal distortion is caused by $\mu(p)$, which determines the magnitude of stretch or shrinkage of the ellipses.

Given any Beltrami coefficient μ with $\|\mu\|_\infty < 1$, we can find a corresponding quasi-conformal mapping satisfying the Beltrami equation in the distribution sense [71]:

Theorem 1. *Suppose $\mu : \mathbb{D} \to \mathbb{C}$ is Lebesgue measurable with $\|\mu\|_\infty < 1$. There is a quasi-conformal homeomorphism ϕ from $\mathbb{D}$ to itself, which is in the Sobolev space $W^{1,2}(\mathbb{D})$ and satisfies the Beltrami equation (1) in the distribution sense. Furthermore, by fixing 0 and 1, ϕ is uniquely determined.*

Given an orientation preserving homeomorphism ϕ, we can find the corresponding Beltrami coefficient from (1):

$$\mu_\phi = \frac{\partial \phi}{\partial \bar{z}} \Big/ \frac{\partial \phi}{\partial z}. \tag{3}$$

The Jacobian J of ϕ is related to μ_ϕ as follows:

$$J(\phi) = \left|\frac{\partial \phi}{\partial z}\right|^2 \left(1 - \left|\mu_\phi\right|^2\right). \tag{4}$$

Since ϕ is an orientation preserving homeomorphism, $J(\phi) > 0$ and $\left|\mu_\phi\right| < 1$ everywhere. Hence, we must have $\left\|\mu_\phi\right\|_\infty < 1$. Theorem 1 indicates that under suitable normalization, every μ with $\|\mu\|_\infty < 1$ is associated with a unique homeomorphism. Therefore, a homeomorphism from $\mathbb{C}$ or $\mathbb{D}$ onto itself can be uniquely determined by its associated Beltrami coefficient.

3.2. Surface curvature

Curvature is an important quantity in differential geometry for assessing how a surface deviates from a plane.

We define $N : S \to \mathbb{S}^2 \subseteq \mathbb{R}^3$ to be the *normal map* giving unit vector at each point p. Suppose C is a regular curve on S, p is a point on S and k is the curvature of C at p. We set $\cos\theta = \langle n, N\rangle$, where N is the normal vector to S at p and n is normal to C. $k_n = k\cos\theta$ is called the *normal curvature* of C at p. The *principal curvatures* at p are the maximum and minimum of the normal curvature, denoted as k_1 and k_2 respectively. The *mean curvature* at p is defined to be $H = \frac{1}{2}(k_1 + k_2)$.

3.3. Linear Beltrami Solver (LBS)

In [72], Lui et al. proposed an efficient method called the *Linear Beltrami Solver* (LBS) to reconstruct the associated quasi-conformal homeomorphism from any given Beltrami coefficient. The method is outlined below.

Let M_1 and M_2 be two planar domains and $\mu = \rho + i\tau$, where $i^2 = -1$, be a complex-valued function defined on M_1. The LBS method aims to reconstruct the homeomorphism $f : M_1 \to M_2$ associated with the Beltrami coefficient μ. Let $f = u + iv$. Since

$$\frac{\partial f}{\partial z} = \frac{1}{2}\left(\frac{\partial}{\partial x} - i\frac{\partial}{\partial y}\right), \quad \frac{\partial f}{\partial \bar{z}} = \frac{1}{2}\left(\frac{\partial}{\partial x} + i\frac{\partial}{\partial y}\right), \tag{5}$$

we can derive from (1) that

$$\mu(f) = \frac{(u_x - v_y) + i(v_x + u_y)}{(u_x + v_y) + i(v_x - u_y)}. \tag{6}$$

We can then represent u_x, u_y by v_x, v_y and v_x, v_y by u_x, u_y as follows:

$$\begin{cases} v_y = \alpha_1 u_x + \alpha_2 u_y, \\ -v_x = \alpha_2 u_x + \alpha_3 u_y, \end{cases} \quad \text{and} \quad \begin{cases} -u_y = \alpha_1 v_x + \alpha_2 v_y, \\ u_x = \alpha_2 v_x + \alpha_3 v_y, \end{cases} \tag{7}$$

where $\alpha_1 = \frac{(\rho-1)^2+\tau^2}{1-\rho^2-\tau^2}$, $\alpha_2 = -\frac{2\tau}{1-\rho^2-\tau^2}$, $\alpha_3 = \frac{(\rho+1)^2+\tau^2}{1-\rho^2-\tau^2}$. Since $\nabla \cdot \begin{pmatrix} -v_y \\ v_x \end{pmatrix} = 0$, we have

$$\nabla \cdot \left(A \begin{pmatrix} u_x \\ u_y \end{pmatrix}\right) = 0 \text{ and } \nabla \cdot \left(A \begin{pmatrix} v_x \\ v_y \end{pmatrix}\right) = 0, \tag{8}$$

where $A = \begin{pmatrix} \alpha_1 & \alpha_2 \\ \alpha_2 & \alpha_3 \end{pmatrix}$. By solving the above equations with certain prescribed boundary constraints, the map f can be obtained.

In the discrete case where M_1 is a triangular mesh, we need to restrict that f is piecewise linear on each triangular face T, which can be written as

$$f|_T(x,y) = \begin{bmatrix} u|_T(x,y) \\ v|_T(x,y) \end{bmatrix} = \begin{bmatrix} a_T x + b_T y + r_T \\ c_T x + d_T y + s_T \end{bmatrix}. \tag{9}$$

Hence, the partial derivatives of f at each face T can be denoted as $D_x f(T) = a_T + ic_T$ and $D_y f(T) = b_T + id_T$. Now the gradient $\nabla_T f := (D_x f(T), D_y f(T))^t$ on T can be computed by solving

$$\begin{pmatrix} \vec{v}_1 - \vec{v}_0 \\ \vec{v}_2 - \vec{v}_0 \end{pmatrix} \nabla_T f = \begin{pmatrix} f(\vec{v}_1) - f(\vec{v}_0) \\ f(\vec{v}_2) - f(\vec{v}_0) \end{pmatrix}, \tag{10}$$

where $[\vec{v}_0, \vec{v}_1]$ and $[\vec{v}_0, \vec{v}_2]$ are two edges on T.

The Beltrami coefficient μ is also discretized on the triangular faces. Denote the discretized functions $\alpha_1, \alpha_2, \alpha_3$ on a face T by $\alpha_1^T, \alpha_2^T, \alpha_3^T$. From (7), we have

$$\begin{aligned} -d_T &= \alpha_1^T a_T + \alpha_2^T b_T, \\ c_T &= \alpha_2^T a_T + \alpha_3^T b_T, \end{aligned} \quad \text{and} \quad \begin{aligned} -b_T &= \alpha_1^T c_T + \alpha_2^T d_T, \\ a_T &= \alpha_2^T c_T + \alpha_3^T d_T. \end{aligned} \tag{11}$$

Let $T = [\vec{v}_i, \vec{v}_j, \vec{v}_k]$ and $\vec{w}_I = f(\vec{v}_I)$, where $I = i, j, k$. Suppose $v_I = g_I + ih_I$ and $w_I = s_I + it_I$. From (10), we have

$$\begin{bmatrix} a_T & b_T \\ c_T & d_T \end{bmatrix} \begin{bmatrix} g_j - g_i & g_k - g_i \\ h_j - h_i & h_k - h_i \end{bmatrix} = \begin{bmatrix} s_j - s_i & s_k - s_i \\ t_j - t_i & t_k - t_i \end{bmatrix}. \tag{12}$$

Thus,

$$\begin{bmatrix} a_T & b_T \\ c_T & d_T \end{bmatrix} = \begin{bmatrix} A_i^T s_i + A_j^T s_j + A_k^T s_k & B_i^T s_i + B_j^T s_j + B_k^T s_k \\ A_i^T t_i + A_j^T t_j + A_k^T t_k & B_i^T t_i + B_j^T t_j + B_k^T t_k \end{bmatrix}, \tag{13}$$

where

$$\begin{aligned} A_i^T &= (h_j - h_k)/2 \cdot \text{Area}(T), & B_i^T &= (g_k - g_j)/2 \cdot \text{Area}(T), \\ A_j^T &= (h_k - h_i)/2 \cdot \text{Area}(T), & B_j^T &= (g_i - g_k)/2 \cdot \text{Area}(T), \\ A_k^T &= (h_i - h_j)/2 \cdot \text{Area}(T), & B_k^T &= (g_j - g_i)/2 \cdot \text{Area}(T). \end{aligned} \tag{14}$$

For each vertex v_i, let $\mathcal{N}_i$ be the collection of neighboring faces of v_i. We can see that

$$\sum_{T \in \mathcal{N}_i} A_i^T b_T = \sum_{T \in \mathcal{N}_i} B_i^T a_T; \quad \sum_{T \in \mathcal{N}_i} A_i^T d_T = \sum_{T \in \mathcal{N}_i} B_i^T c_T. \tag{15}$$

Substituting (11) into (15), we obtain the following equations:

$$\sum_{T \in \mathcal{N}_i} (A_i^T[\alpha_1^T a_T + \alpha_2^T b_T] + B_i^T[\alpha_2^T a_T + \alpha_3^T b_T]) = 0, \tag{16}$$

$$\sum_{T \in \mathcal{N}_i} (A_i^T[\alpha_1^T c_T + \alpha_2^T d_T] + B_i^T[\alpha_2^T c_T + \alpha_3^T d_T]) = 0. \tag{17}$$

Therefore, we can solve this linear system to get the xy-coordinates and hence the desired mapping f.

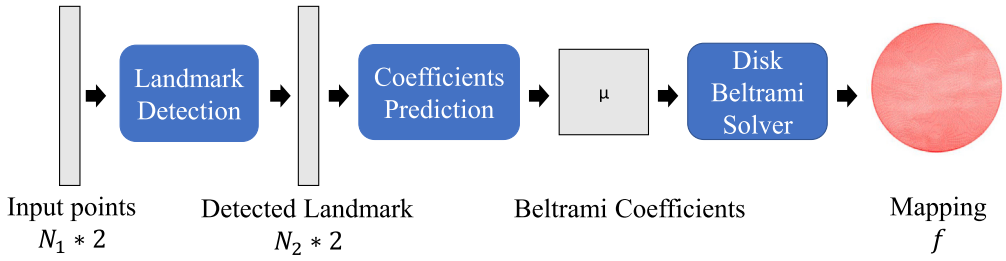

Fig. 2. The architecture of the three proposed networks for the automatic landmark detection and surface registration. The network takes the coordinates of N pairs of endpoints of the sulci as the input and generates a landmark-constrained quasi-conformal mapping on the unit disk as the output.

3.4. Compression by Fourier approximation

Consider the real and imaginary parts of the Beltrami coefficient μ as two channels of the image and denote them by μ_R and μ_I respectively. The Fourier transform of μ_j (where $j = R$ or I) can be expressed as

$$\hat{\mu_j}(m,n) = \frac{1}{N^2}\sum_{k=0}^{N-1}\sum_{l=0}^{N-1}\mu_j(k,l)e^{-i\frac{2\pi km}{N}}e^{-i\frac{2\pi ln}{N}}. \tag{18}$$

The inverse Fourier transform of $\hat{\mu_j}$ can be written as

$$\mu_j(p,q) = \sum_{m=0}^{N-1}\sum_{n=0}^{N-1}\hat{\mu_j}(m,n)e^{i\frac{2\pi pm}{N}}e^{i\frac{2\pi qn}{N}}. \tag{19}$$

In [72], Lui et al. showed that it is possible to compress the Beltrami coefficients by Fourier approximation while preserving the information of homeomorphism and bijectivity. Therefore, keeping only the low-frequency components is sufficient for capturing the majority of deformation in our problem.

4. Proposed framework

In this section, we describe our proposed framework for the automatic landmark detection and surface registration in detail.

Suppose the given brain cortical surfaces are simply-connected open surfaces, i.e. with disk topology. Our strategy is to map them onto the unit disk using conformal parameterization [73], and then design three networks for achieving the automatic landmark detection and registration (see Fig. 2 for an outline): the landmark detection network (LD-Net), the coefficient prediction network (CP-Net), and the disk Beltrami solver network (DBS-Net). For the landmark detection, we first label the starting and ending points of the desired feature curves as the input of the LD-Net, which extracts the complete landmark curves based on the mean curvature of the input surface. Then, we take the output from the LD-Net as the input of the CP-Net for generating a Beltrami coefficient μ. Finally, we obtain a quasi-conformal mapping associated with the Beltrami coefficient μ using the DBS-Net. The architecture of each network is introduced in the following subsections. As a remark, here we focus on simply-connected open surfaces so that we can parameterize them onto the 2D plane and subsequently utilize mapping solvers and learning methods developed for 2D domains, thereby largely simplifying the computation of the registration.

4.1. Landmark detection

In this section, we describe our proposed landmark detection network (LD-Net). The aim of LD-Net is to take the starting and ending points of the desired feature curves as the input and automatically generate the complete landmark curves along prominent features based on the mean curvature of each surface.

For each input surface, we compute its mean curvature at every vertex. We then apply the disk conformal mapping method [73] to parameterize the surface onto the unit disk. In the visualization of the parameterized brain in Fig. 3, the dark regions correspond to the sulci of a cortical surface and the bright regions correspond to the gyri (the ridges on the brain surface).

After getting the disk mesh, we label the locations of the endpoints on it and record their corresponding xy-coordinates. We then rotate the mesh and rescale it based on the distance between two points for normalization (see Fig. 3). In this way, we can generate patches with all starting and ending points at the same location. Then, we concatenate the patch with a new channel recording two points by setting the pixel values of the starting and ending points as 255 and all other pixel values as 0. In our implementation, the first layer is followed by a batch normalization layer. Therefore, the concatenated image is equivalent to a binary image. After passing the concatenated image through the convolution layers, we reshape it and then pass it through the fully connected layers, which ultimately generate the xy-coordinates for the predicted landmarks.

To explain this procedure, note that if the starting and ending points are given, the sulcus connecting them should be contained in a local window and the output should not be affected by any unrelated information or noise. By zooming in and rotating the mesh, we have the starting and ending points of all the training data at the same location, and hence we do not need any additional input of its coordinates for the prediction process.

Loss Function: For the training of the network, we adopt supervised learning by generating and augmenting data. The loss function for this network is as follows:

$$\mathcal{L}_{\mathrm{LD}} = \zeta\sum_{i=1}^{N}\|x_i - \hat{x}_i\|^2 + \kappa\sum_{i=1}^{N-1}\left|\|\hat{x}_{i+1} - \hat{x}_i\|^2 - \overline{\Delta x}\right|, \tag{20}$$

where N is the number of output points for the prediction of the landmark curves, and $x_i, \hat{x}_i$ are respectively the actual point and the output point from the network for $i = 1, \ldots, N$, $\overline{\Delta x} = \frac{1}{N-1}\sum_{i=1}^{N-1}\|\hat{x}_{i+1} - \hat{x}_i\|^2$. Here, the first term aims to improve the accuracy of the landmark prediction, and the second term aims to make the output points more evenly distributed. ζ and κ are two weighting parameters.

4.2. Landmark-based quasi-conformal mapping

As described previously, quasi-conformal mappings can be represented by Beltrami coefficients, and it is easier to generate a 1–1 mapping from the Beltrami coefficients. Therefore, the two main tasks here are: (1) to generate a quasi-conformal mapping based on any input Beltrami coefficient in a fast and robust way, and (2) to generate a Beltrami coefficient that corresponds to a landmark-based registration to be fed into the mapping method.

4.2.1. Disk Beltrami Solver Network (DBS-Net)

Note that the LBS method [72] involves solving linear systems and so the computation may be time-consuming if one has to handle a large set of dense triangular meshes. To overcome this problem, Chen et al. [44] proposed a method to train a network that can generate quasi-conformal mappings from Beltrami coefficients more efficiently. However, their method is only applicable to image registration with

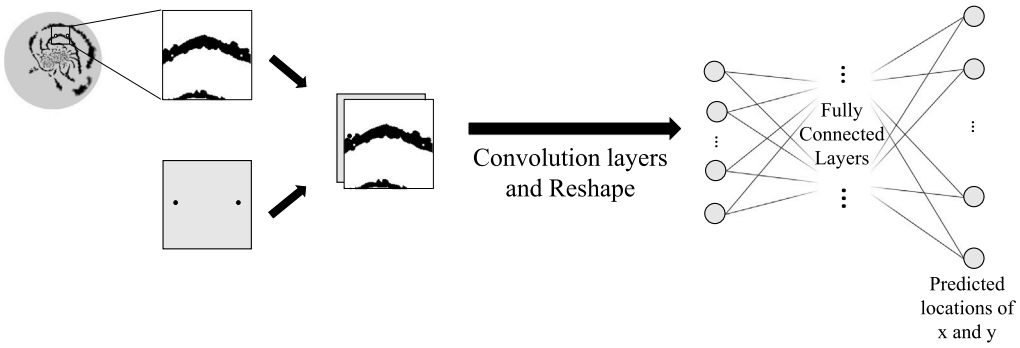

Fig. 3. The structure of the landmark detection network (LD-Net). We first label the starting and ending point on the parameterized mesh. Next, we rescale and rotate the mesh so that the starting and ending points are aligned horizontally in a consistent manner and the sulcus connecting them is contained in a local window. The two-color image is generated from the local patch based on the curvature on each vertex. Then, we concatenate the image with a single channel recording the starting and ending points. The network takes the concatenated image patch as an input, which passes through the convolution layers and fully connected layers, and finally generates the predicted locations.

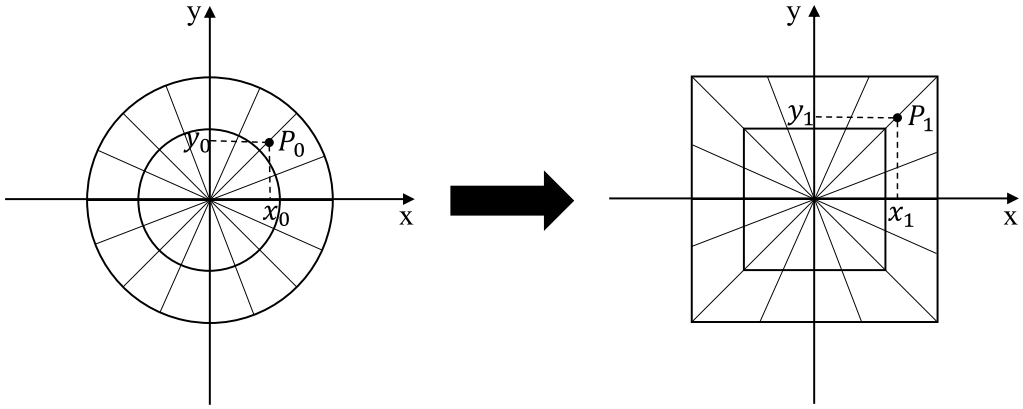

Fig. 4. The circle-to-square transformation in our framework.

the data represented using regular grids. To apply this idea to the parameterized irregular triangular meshes in our case, it is necessary to develop a new method.

Here we propose the *Disk Beltrami Solver Network* (DBS-Net) for generating a quasi-conformal mapping based on any input Beltrami coefficient on the disk. To make use of CNNs in the disk, we need to reshape the disk Beltrami coefficients to a square domain. Here, we consider the following transformation:

$$x_1 = \frac{x_0}{\max\{\cos\theta, \sin\theta\}}, \quad y_1 = \frac{y_0}{\max\{\cos\theta, \sin\theta\}}, \tag{21}$$

where (x_0, y_0), (x_1, y_1) are the coordinates of a point in the disk and its corresponding point in the square, and θ is the angle of the point in the polar coordinate system (see Fig. 4).

After transforming the disk Beltrami coefficients, the DBS-Net uses two routes to determine the inner and boundary coordinates for the quasi-conformal map separately (see Fig. 5).

Route for the Inner Coordinates: Recall that we can use the Fourier approximation to get the majority of information from the Beltrami coefficient μ by keeping its low-frequency components. In our work, we keep 2% of the Fourier coefficients of the Beltrami coefficient. We first apply (18) to obtain the DFT $\hat{\mu}$ of μ. We then remove the high frequency components to get a low-resolution approximation $\hat{\mu}_{tc}$. Next, we apply the following equation to imitate the process of computing the inverse Fourier transform μ_{tc} of $\hat{\mu}_{tc}$ and obtaining the low-resolution Beltrami coefficient $\tilde{\mu}$ associated to the quasiconformal mapping of μ_{tc} on a coarser spatial domain:

$$\tilde{\mu} = M\hat{\mu}_{tc}N = (\hat{\mu}_{tc}^T M^T)^T N, \tag{22}$$

where M and N are learnable. We call this process the *Domain Transform* (DT) process. Finally, we use bilinear interpolation to upsample and conduct convolution right after the interpolation (see Fig. 6).

While the Fourier approximation can preserve the majority of information, some details will still be lost. To overcome this problem, here we use a shortcut to preserve the detailed information and concatenate the feature map generated from the shortcut with the one that was processed after the Fourier approximation. This allows us to process the data in a more efficient way while also having the details in the output.

Route for the Boundary Coordinates: To reconstruct a quasi-conformal map from a Beltrami coefficient, proper boundary conditions are needed. In our case, we need to restrict the boundary points of the quasi-conformal map to remain on the boundary of the unit disk. It is noteworthy that the boundary points are not meant to be fixed. Instead, we only fix one boundary point to remove the freedom of the global rotation and allow all the remaining boundary points to move along the disk boundary to get a low-distortion mapping. To achieve this, here we design a separate route for determining the boundary coordinates.

More specifically, we consider the angles between different boundary points. Since the calculation of each point involves its neighboring points and hence is related to every point on the disk, we need the information of the Beltrami coefficient μ on every face. Therefore, we

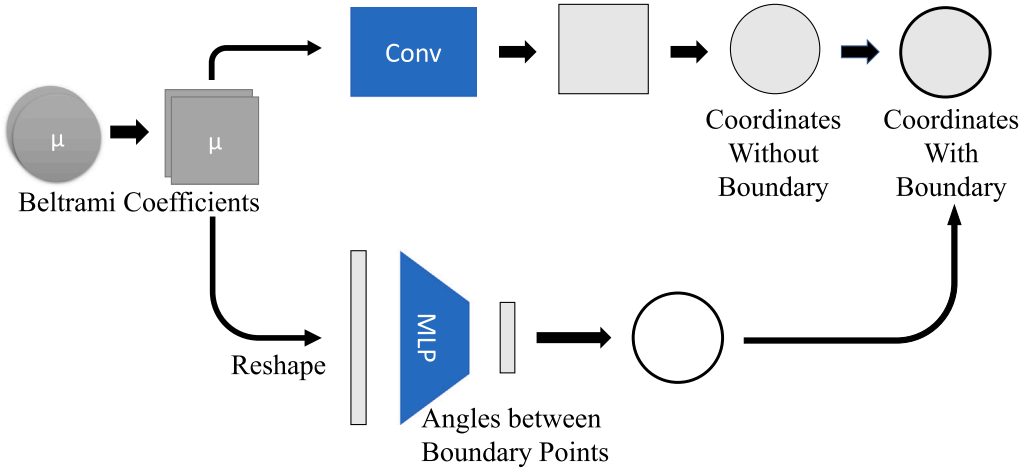

Fig. 5. The flow of the proposed DBS-Net.

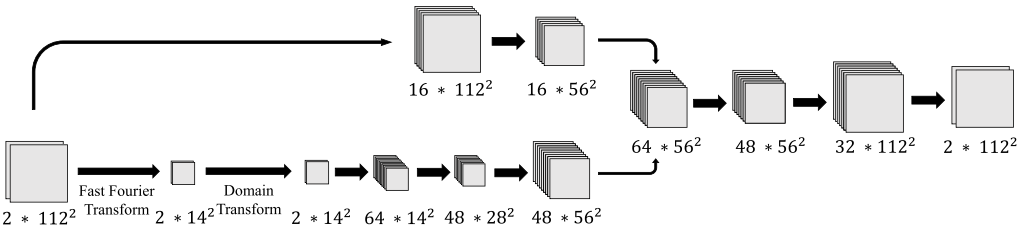

Fig. 6. The convolution for generating the inner coordinates of the quasi-conformal maps.

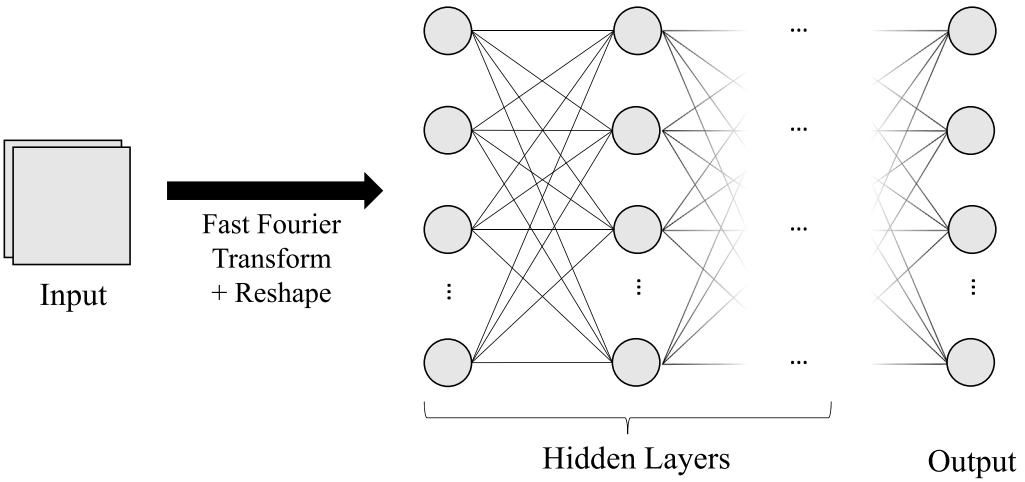

Fig. 7. The process of obtaining the boundary coordinates of the quasi-conformal maps.

reshape the input μ to be a column vector after using the FFT and only keep the low-frequency data. Then, we adopt the MLP to generate the desired angles, thereby determining the boundary coordinates (see Fig. 7).

Loss Function: For the training, we use an unsupervised setting by adopting LBS in computing the loss function.

We first consider replacing a_T, b_T, c_T, d_T in (16) and (17) with the expression in (13) to derive the coefficients for every vertex v_i and all vertices v_l adjacent to v_i. We have

$$c_i = \sum_{T\in\mathcal{N}_i}\left[\alpha_1^T\left(A_i^T\right)^2 + 2\alpha_2^T A_i^T B_i^T + \alpha_3^T\left(B_i^T\right)^2\right], \tag{23}$$

where $\mathcal{N}_i$ is the collection of neighboring faces of v_i. Note that each edge is shared by two faces. Denoting the two faces sharing the edge $[v_i, v_l]$ by T_1, T_2 and the remaining vertices in T_1 and T_2 by v_j and v_k respectively, the coefficient c_l can be expressed as:

$$\begin{aligned} c_l =& \alpha_1^{T_1} A_i^{T_1} A_j^{T_1} + \alpha_2^{T_1}\left(A_i^{T_1} B_j^{T_1} + A_j^{T_1} B_i^{T_1}\right) + \alpha_3^{T_1} B_i^{T_1} B_j^{T_1} + \\ & \alpha_1^{T_2} A_i^{T_2} A_k^{T_2} + \alpha_2^{T_2}\left(A_i^{T_2} B_k^{T_2} + A_k^{T_2} B_i^{T_2}\right) + \alpha_3^{T_2} B_i^{T_2} B_k^{T_2}. \end{aligned} \tag{24}$$

Therefore, we define the following loss function to evaluate the difference between the generated mapping result and the result produced by LBS:

$$\mathcal{L}_{\mathrm{LBS}} = \frac{1}{2N^2}\sum_i\left(|c_i s_i + \sum_{v_l\in\mathcal{V}_i} c_l s_l| + |c_i t_i + \sum_{v_l\in\mathcal{V}_i} c_l t_l|\right), \tag{25}$$

where s_i and t_i are x and y coordinates of v_i respectively, N is the total number of vertices, and $\mathcal{V}_i$ is the collection of all vertices adjacent to

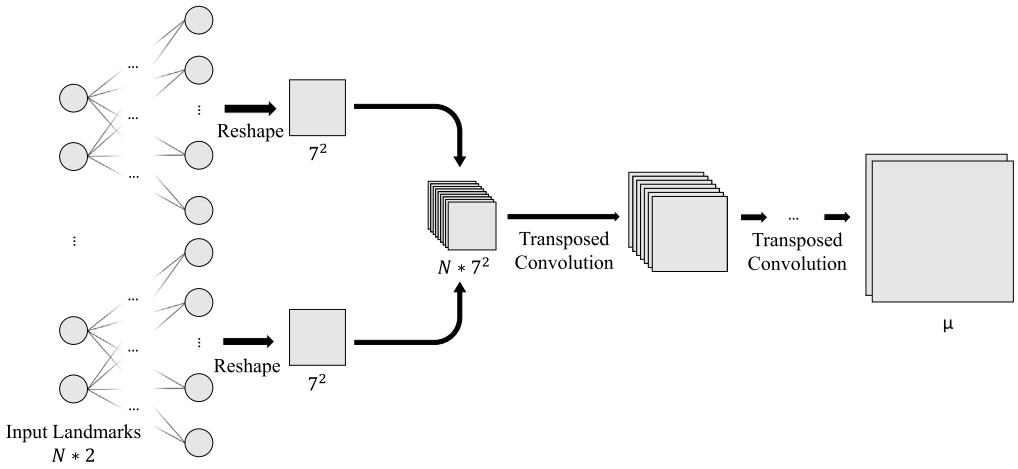

Fig. 8. The structure of the coefficient prediction network (CP-Net). We first process the input landmark locations through MLP and reshape them into separate channels before concatenation. Then we apply the transposed convolution and a tanh-type activation function to generate the Beltrami coefficients.

v_i. This effectively ensures that the output coordinates satisfy (16) and (17).

Moreover, note that the quasi-conformality of a mapping does not change under compositions with conformal maps. With such flexibility, the vertices in the output mapping tend to cluster around the fixed boundary point. To overcome this problem and let the boundary points distribute more evenly on the disk boundary, we introduce another loss function:

$$\mathcal{L}_{\text{Boundary}} = \sum_{i=1}^{N} \frac{1}{\theta_i}, \tag{26}$$

where N is the number of boundary points and θ_i is the angle between two neighboring boundary points with $\sum_{i=1}^{N} \theta_i = 2\pi$. Intuitively, each $\frac{1}{\theta_i}$ term effectively prevents any two neighboring boundary points from being too close to each other, thereby ensuring that the boundary points will be distributed more evenly over the entire disk boundary.

Altogether, we propose the following combined loss function:

$$\mathcal{L}_{\text{DBSNet}} = \eta \mathcal{L}_{\text{LBS}} + \rho \mathcal{L}_{\text{Boundary}}, \tag{27}$$

where η, ρ are two nonnegative balancing parameters. Note that if ρ is too small, $\mathcal{L}_{\text{DBSNet}}$ will be dominated by the first term and the clustering effect will still be present in the mapping result. By contrast, if ρ is too large, the second term in $\mathcal{L}_{\text{DBSNet}}$ will significantly constrain the boundary points. In practice, we find that $\eta = 1$ and $\rho = 10^{-8}$ give satisfactory results. Here, we remark that the magnitude of $\mathcal{L}_{\text{LBS}}$ and $\mathcal{L}_{\text{Boundary}}$ may not be the same. Specifically, with the $1/\theta_i$ term in $\mathcal{L}_{\text{Boundary}}$, its magnitude would be very large if we have a small θ_i. Therefore, while η and ρ appear to be largely different in magnitude, both terms are important.

4.2.2. Coefficient prediction network (CP-Net)

The goal of CP-Net is to take the starting points as the input and output a Beltrami coefficient μ, which can then be used for generating the quasi-conformal mapping using the DBS-Net.

Initially, we use the MLP to process the starting points. We separate the points into groups and use separate MLPs to generate the feature map in the latent space. After that, we reshape the points into channels and perform a transposed convolution on them. Recall that for the DBS-Net, we adopt the bilinear interpolation for upsampling and do convolution right after that, as the information in the feature map is enough for us to get the output mapping. By contrast, the coefficient prediction task involves generating global information from certain sampled local information, while the reshaped landmark coordinates do not contain all information we need for the entire disk. To get a trainable upsampling for generating the Beltrami coefficients, we use the transposed convolution defined by padding and stride to get the global feature from the local feature. This gives a complex number output $\nu(T)$ for each triangle T. Finally, we would like to ensure that the quasi-conformal mappings subsequently generated by the output Beltrami coefficients of this network are bijective, i.e. there will be no folding in the resulting mappings. To achieve this, we follow the idea in [44] and add the following tanh-type activation function $\mathcal{T}$ in the last layer:

$$\mu(T) = \mathcal{T}(\nu(T)) = \tanh(|\nu(T)|)e^{i \arg(\nu(T))}. \tag{28}$$

This ensures that $\|\mu\|_\infty < 1$ and hence the mappings are always bijective. The structure of the CP-Net is shown in Fig. 8.

Loss Function: Recall that $|\mu|$ measures the quasi-conformal distortion of a mapping and hence it is desired to be as small as possible. Therefore, we consider the following term:

$$\mathcal{L}_{\mu} = \frac{1}{N} \sum_{i=1}^{N} \|\mu_i\|_2^2, \tag{29}$$

where N is the number of faces. Next, since $|\nabla \mu|$ controls the smoothness of the mapping, we consider the following term:

$$\mathcal{L}_{\nabla \mu} = \frac{1}{N} \sum_{i=1}^{N} \|\nabla \mu_i\|_2^2. \tag{30}$$

Finally, the landmark mismatch error can be assessed by the following term:

$$\mathcal{L}_{\text{Landmark}} = \frac{1}{M} \sum_{i=1}^{M} \|p_i - \hat{p}_i\|_2^2, \tag{31}$$

where M is the number of landmarks, p_i is the target location, and $\hat{p}_i$ is the network output, $i = 1, \ldots, M$. By combining (29), (30) and (31), we have the following loss function:

$$\mathcal{L}_{\text{CP}} = \alpha \mathcal{L}_{\mu} + \beta \mathcal{L}_{\nabla \mu} + \gamma \mathcal{L}_{\text{Landmark}}, \tag{32}$$

where α, β, γ are nonnegative parameters. Note that for the training of the CP-Net, we need to have a pre-trained DBS-Net for the data generation and for producing the ultimate quasi-conformal mappings from the outputs of the CP-Net. This will be described in Section 5.

5. Experiments

In this section, we describe our implementation of the proposed framework and present the experimental results. 24 human brain cortical surfaces used in our experiments are reconstructed from Magnetic resonance imaging (MRI) images from the Open Access Series of Imaging Studies (OASIS) [74] using FreeSurfer [75]. Each surface

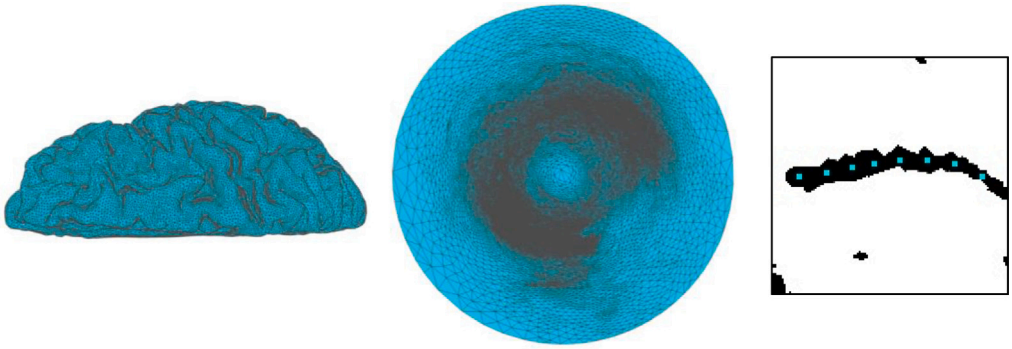

Fig. 9. (Left) A brain cortical surface mesh reconstructed from MRI images. (Middle) The disk conformal parameterization of the surface mesh. (Right) A two-color image is then generated from a local patch of the disk after rotation and rescaling, which is subsequently concatenated with a single channel recording the starting and ending points. The blue points are the labeled landmarks.

is represented as a triangular mesh consisting of about 45000 vertices and 90,000 faces, with a cut introduced at the lateral ventricle to make it simply-connected. Six sulci, including the Central Sulcus (CS), Post-central Sulcus (PostCS), Superior Frontal Sulcus (SFS), Inferior Frontal Sulcus (IFS), Superior Temporal Sulcus (STS), and Inferior Temporal Sulcus (ITS), are used as landmark curves for the registration. For testing the performance of the landmark detection network, a set of sulcal landmarks are also manually delineated and verified by medical experts as the benchmark.

5.1. Training the landmark detection network

Our goal is to ensure that the LD-Net can produce landmark curves accurately along the sulci. To start with, we need to generate the training data for landmark detection (see Fig. 9 for a sample). We first prepare multiple template brains and manually label densely distributed landmarks along the sulci as the prepared data. The brains are then parameterized onto the unit disk using the disk conformal parameterization method in [73]. Then, we randomly select the starting and ending points on those sulci and generate the local patch and its corresponding labels by selection or interpolation. More specifically, after selecting the starting and ending points, we rescale the mesh based on the distance between the two points and rotate it horizontally, so that the starting and ending points for different patches are normalized to be at the same location. Since the sulci of each brain are densely labeled, we can generate plenty of patches at different places and scales, which enables us to produce enough data for training.

After getting a well-trained network, we test it on the dataset with medically verified sulcal landmarks. Some examples are presented in Fig. 10, from which we can see that the sulci can be accurately extracted. More quantitatively, we compare our network output and the medically verified landmarks, and the resulting mean difference and the standard deviation are 0.0089 and 0.0062 respectively. In Table 1, we further consider different choices of the parameters in the LD-Net and assess the landmark error and the landmark distribution. Note that the aim of LD-Net is to get the detected landmarks and ensure that they are evenly distributed. If we set κ as zero, then rigid error, which is the difference between the outputs and the labels, could be small but the difference in gaps is large. However, if we set κ as 0.1, then we can effectively reduce the difference between gaps while keeping the rigid error almost unchanged. This demonstrates the effectiveness of our proposed LD-Net.

Table 1
The effect of different choices of parameters on the landmark detection network (LD-Net). Here, the rigid error and the difference between gaps are computed using the second term in (20).

Parameter	Rigid Error	Difference between Gaps
$(\zeta, \kappa) = (1, 0)$	5.28	6.12
$(\zeta, \kappa) = (1, 1)$	16.6	4.42
$(\zeta, \kappa) = (1, 0.1)$	5.87	1.78

5.2. Experiments on disk Beltrami Solver network

5.2.1. Training

At the start of our training for the DBS-Net, we need sufficiently many Beltrami coefficient data on the disk for training. These data should be random and should contain as many scenarios as possible. To achieve this, we also use data from ImageNet [43] to augment the input Beltrami coefficients. In the augmenting process, we first convert all images into grayscale images and randomly select one as the real part and another as the imaginary part. We then smooth them and add random noise to them. The resulting two-channel image then serves as the input of DBS-Net, with the value for each pixel used for computing the LBS term in the loss function (27). Note that with the large number of images in ImageNet, we can generate a large variety of Beltrami coefficients for the training process. Consequently, the trained network can handle any type of data effectively. We can then apply it for brain cortical surface registration.

Some example outputs of the DBS-Net are shown in Fig. 11, from which it can be observed that the error in the Beltrami coefficient is very small. This shows that the DBS-Net is capable of generating quasi-conformal mappings from the input Beltrami coefficients accurately. For a more quantitative analysis, we compare the performance of LBS [72] and our DBS-Net in terms of the error in the norm of the Beltrami coefficients $|\mu|$ and the computation time (see Table 2). Since LBS is used as our training loss, it can be observed that the error achieved by DBS-Net is not as low as that by LBS directly. Nevertheless, the computation time of DBS-Net is shorter than that of LBS by over 50%, and the resulting quasi-conformal mapping generated by DBS-Net is already good enough to be combined with CP-net for disk registration. Another improvement achieved by DBS-Net is that we only need to fix one point on the boundary in using DBS-Net, while LBS requires fixing every point on the boundary. Consequently, DBS-Net is more flexible than LBS in generating the corresponding quasi-conformal mappings.

5.2.2. Ablation study

Next, we study the necessity of having the Domain Transform (DT) process, the shortcut path for the network and the boundary term in training loss. We first consider replacing the Domain Transform process with convolution layers and training the network by the same training method. We call this method DBS-Net (No DT) and test it on the same dataset. As indicated by the large error in $|\mu|$ in Table 2, this method cannot even generate mappings without foldings and hence is undesirable. We also consider using more convolution layers to replace the Domain Transform process, which we refer to as DBS-Net (No DT2) in

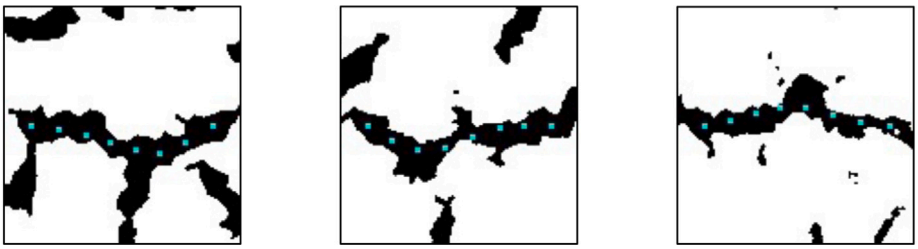

Fig. 10. Some sample outputs from the landmark detection network (LD-Net). The blue points show the detected landmarks.

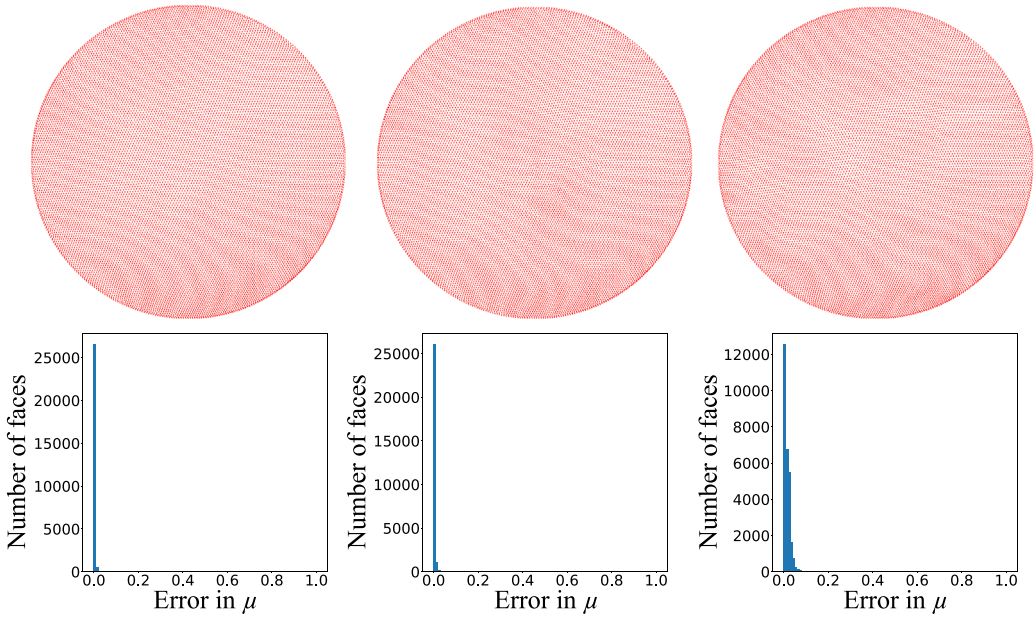

Fig. 11. Three sample outputs of the disk Beltrami solver network (DBS-Net) with histograms of the error in the Beltrami coefficient μ.

Table 2
A comparison between the proposed DBS-Net and the linear Beltrami solver (LBS) [72] and the ablation study of DBS-Net on the test dataset. For each method, the mean error and standard deviation in the norm of the Beltrami coefficients $|\mu|$ are recorded. All experiments were performed on a computer with the Intel CPU i7-11800H processor, and the time is the average time cost for each input.

Method	Mean Error in $\|\mu\|$	SD Error in $\|\mu\|$	Time (s)
DBS-Net	1.54×10^{-2}	8.9×10^{-3}	1.57×10^{-2}
LBS [72]	5.7×10^{-3}	5.9×10^{-3}	3.8×10^{-2}
DBS-Net (No DT)	4.17	2.86	1.88×10^{-2}
DBS-Net (No DT2)	5.14	3.76	2.01×10^{-2}
DBS-Net (No Shortcut)	2.43×10^{-1}	3.93×10^{-2}	1.29×10^{-2}

Table 2, to alleviate the influence of the depth of the network. However, the performance is still unsatisfactory. In addition, we remove the shortcut in DBS-Net, which we refer to as DBS-Net (No Shortcut), and train the network by the same training method. As shown in Table 2, while this method can generate mappings without folding, the error in $|\mu|$ is large and hence the performance is unsatisfactory. As for the boundary term in the loss function, the aim is to prevent the points on the boundary from concentrating at one part on the disk. In Fig. 12, we show a sample result from the training without the boundary term loss. It can be observed that even though the error in the Beltrami coefficient μ is small, the output mapping squeezes at a small region, which is not satisfactory. From the above studies, we can see that the Domain Transform process, the shortcut path in DBS-Net and the boundary term in loss function are important for generating satisfactory quasi-conformal mappings.

5.3. Training the coefficient prediction network

After getting a well-trained LD-Net for getting the landmark curves and a well-trained DBS-Net for generating quasi-conformal mappings based on any input Beltrami coefficients, the remaining step is to train the CP-Net to generate the Beltrami coefficients from the landmark constraints.

To achieve this, we first need to prescribe the target positions of the landmarks. To generate the training data, we randomly choose two single-channel images and concatenate them to be two-channel data that acts as the Beltrami coefficient. We then use it as the input of the DBS-Net to generate a quasi-conformal mapping for perturbing the landmarks, and take the positions of the perturbed landmarks as the input of the CP-Net. All landmarks are discretized and reshaped into one row in the format of $(x_1, y_1, x_2, y_2, \ldots, x_n, y_n)$ before passing through the network. This ensures that there is at least one admissible mapping without folding, i.e. $\|\mu\|_\infty < 1$. We use the distance between the resulting positions of the input points and the target positions to calculate the loss function (32). Note that (32) also contains two terms for the norm of the Beltrami coefficients μ and $\nabla\mu$, which suppress the angle distortion.

5.4. Landmark-based brain cortical surface registration

With all three proposed networks trained, we apply them for landmark-based brain cortical surface registration. We first conformally parameterize every brain onto the unit disk using [73] and then apply the three networks for the landmark-based registration, where one brain from the dataset is selected as the control brain. As illustrated

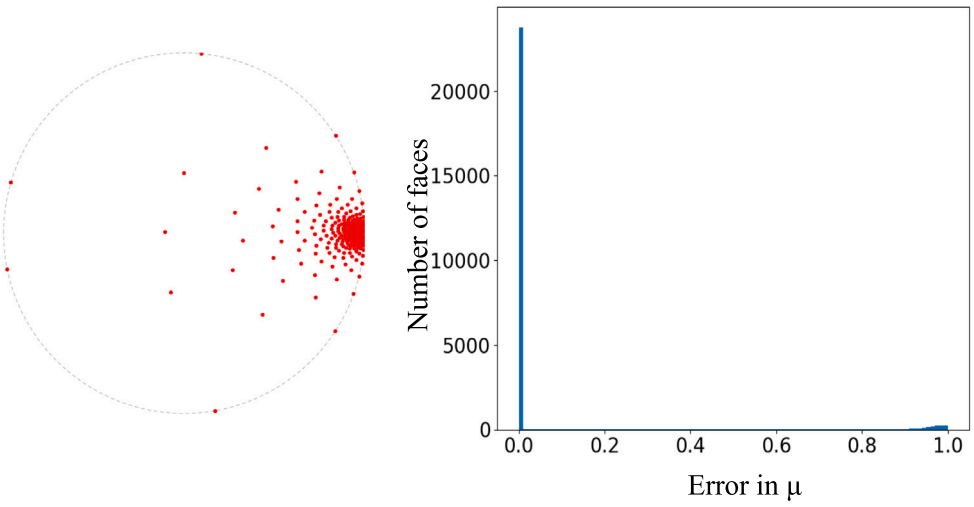

Fig. 12. The left plot shows the output location of the points from the DBS-Net trained without the boundary term. Here, the boundary of the circle is denoted by the dashed line. The right plot shows the corresponding histogram of the error in the Beltrami coefficient μ.

Table 3
A comparison between our framework with different parameters α, β, γ in (32) and other methods for landmark registration in 24 brain cortical surfaces. For the landmark-constrained optimized harmonic mapping approach [18,73], λ is the soft landmark matching parameter in [18]. The mean and standard deviation of the quasi-conformal distortion $|\mu|$ and the landmark error e_{landmark} are assessed. It is noteworthy that our method is capable of achieving a smaller quasi-conformal distortion while maintaining a low landmark error comparable to that obtained using the optimized harmonic map method with a large λ.

Method	Mean $\lvert\mu\rvert$	SD $\lvert\mu\rvert$	Mean e_{landmark}	SD e_{landmark}	Time (s)
Our method $((\alpha,\beta,\gamma)=(1,1,10))$	0.78×10^{-2}	0.23×10^{-3}	3.71×10^{-2}	10.61×10^{-3}	0.59
Our method $((\alpha,\beta,\gamma)=(1,1,10^2))$	1.25×10^{-2}	1.74×10^{-3}	2.03×10^{-2}	3.90×10^{-3}	0.57
Our method $((\alpha,\beta,\gamma)=(1,1,10^3))$	1.95×10^{-2}	3.11×10^{-3}	1.96×10^{-2}	2.79×10^{-3}	0.59
Our method $((\alpha,\beta,\gamma)=(1,1,10^4))$	1.87×10^{-2}	2.78×10^{-3}	1.14×10^{-2}	2.59×10^{-3}	0.53
Our method $((\alpha,\beta,\gamma)=(1,1,20^5))$	4.30×10^{-2}	5.71×10^{-3}	0.62×10^{-2}	1.70×10^{-3}	0.55
Optimized harmonic map $(\lambda=0.1)$	1.00×10^{-2}	2.72×10^{-3}	4.10×10^{-2}	12.1×10^{-3}	7.64
Optimized harmonic map $(\lambda=0.5)$	2.90×10^{-2}	8.37×10^{-3}	2.24×10^{-2}	6.50×10^{-3}	14.96
Optimized harmonic map $(\lambda=1)$	3.98×10^{-2}	11.5×10^{-3}	1.19×10^{-2}	3.43×10^{-3}	32.58
Optimized harmonic map $(\lambda=5)$	5.16×10^{-2}	14.8×10^{-3}	0.51×10^{-2}	0.39×10^{-3}	46.14
TPS [76]	8.31×10^{-2}	26.8×10^{-3}	9.30×10^{-2}	27.2×10^{-3}	0.27
QCLR [29]	5.77×10^{-2}	16.4×10^{-3}	0	0	27.17

in Fig. 1, for every brain, we select the starting and ending points of each sulcus as the input and obtain the quasi-conformal map on the unit disk. We then apply the inverse map of the disk parameterization to map the disk back to the control brain in 3D. This completes the landmark-based brain registration. As shown in Fig. 13 and Fig. 14, the mappings match the sulcal landmarks accurately and possess very low quasi-conformal distortion.

For a more quantitative analysis of all mapping results, we record the quasi-conformal distortion in terms of the norm of the Beltrami coefficients $|\mu|$, the landmark error and the computational time in Table 3. It is noteworthy that by setting different training parameters α, β, γ, we can achieve a lower landmark error or a lower quasi-conformal distortion depending on the goal of the task. For comparison, note that landmark-constrained optimized harmonic mapping methods have been developed for registering genus-0 brain cortical surfaces [18, 39], while the brain surfaces we handle in our framework are simply-connected open surfaces. Therefore, here we consider combining the disk conformal map method [73] with the landmark-constrained optimized harmonic mapping method [18] and compare the performance of this combined approach with our proposed framework. Specifically, we first apply the disk conformal map method to map the brain surfaces onto a unit disk. Then, we solve for the landmark-constrained harmonic map as in [18], with the original triangular boundary constraint replaced with a disk boundary constraint. Finally, we apply the bijectivity preservation method in [18] to enforce the bijectivity of the mapping, which finally gives in the registration result. It can be observed that our proposed method achieves a significant improvement in the computational time by over 90% when compared with the optimized harmonic map approach with different soft landmark matching parameters λ, while the quasi-conformal distortion and landmark error of the resulting registrations are comparable with those obtained using the optimized harmonic map approach. We further compare our method with the thin-plate spline (TPS) method [76] and the quasi-conformal landmark registration (QCLR) algorithm [29]. It can be observed that the TPS method gives a quasi-conformal distortion and a larger landmark error when compared to our method. While the QCLR method can achieve exact landmark matching, it gives a significantly larger quasi-conformal distortion and is also more computationally expensive. In Table 4, we further use the t-test to compare the quasi-conformal distortion and the landmark error obtained using our method and the other methods for all brains. The results suggest that the improvement achieved by our method is statistically significant.

Also, in Fig. 15 we analyze the Beltrami coefficients μ of the resulting mappings for all brains. The histogram of $|\mu|$ shows that our framework achieves a minimal quasi-conformal distortion, with the majority of the $|\mu|$ values being close to 0. By taking the average of $|\mu|$ for all resulting mappings, we can again see that the quasi-conformal distortion is minimal in the entire unit disk. We further plot the standard deviation of $|\mu|$ evaluated on each triangular face, from which we can see a relatively large variation in $|\mu|$ near the landmarks, while the variation in $|\mu|$ far away from the landmarks is very small.

6. Discussion and conclusion

In this paper, we have established a framework with three networks for automatic landmark detection and surface registration of anatomical shapes with disk topology. Firstly, the LD-Net allows us to efficiently

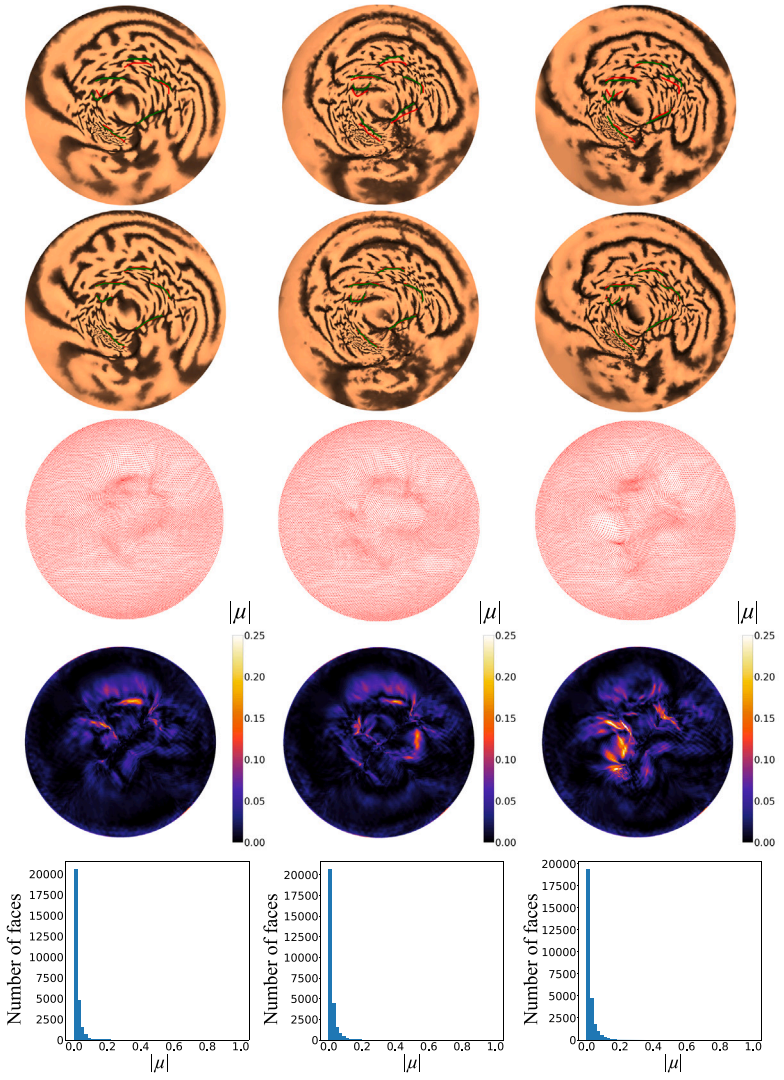

Fig. 13. Some sample outputs of the proposed framework. Each column shows one experiment. The first row shows the initial disks produced using disk conformal parameterization [73] color-coded by the surface curvature, where the green curves are the landmarks detected by the LD-Net and the red ones are the target locations of the landmarks. The second row shows the quasi-conformal maps on the disk produced by the CP-Net and the DBS-Net, from which it can be observed the landmarks are well-aligned. The third row shows an alternative representation of the mapping results. The fourth row shows the norm of the Beltrami coefficients $|\mu|$ on the unit disk. The fifth row shows the histograms of $|\mu|$.

Table 4
A comparison between our method and other methods using the t-test. Here, we record the p-value of the t-test for the quasi-conformal distortion $|\mu|$ and the landmark error e_{landmark} The value in the table is the p-value. The hypothesis is that H_0: population mean of Dataset1 is equal to Dataset2.

	Quasi-conformal distortion $\lvert\mu\rvert$			Landmark error e_{landmark}		
Dataset1 \ Dataset2	Optimized harmonic map ($\lambda = 1$)	QCLR	TPS	Optimized harmonic map ($\lambda = 1$)	QCLR	TPS
Our Method ($(\alpha, \beta, \gamma) = (1, 1, 10^2)$)	2.95×10^{-11}	1.66×10^{-12}	4.85×10^{-12}	6.27×10^{-10}	3.88×10^{-18}	2.52×10^{-12}
Our Method ($(\alpha, \beta, \gamma) = (1, 1, 10^4)$)	3.57×10^{-9}	2.61×10^{-11}	2.79×10^{-11}	0.63	1.44×10^{-16}	2.85×10^{-13}

extract continuous feature landmarks from highly convoluted surfaces based on the endpoints of the desired feature curves and the surface curvature. Then, the CP-Net generates a Beltrami coefficient based on the detected landmarks. Finally, the DBS-Net, trained in unsupervised mode, produces a quasi-conformal mapping on the unit disk based on the generated Beltrami coefficients, which yields the surface registration result. The mappings produced by our framework possess low geometric distortion and preserve the bijectivity. Experimental results have demonstrated the effectiveness of our framework for brain surface registration. More broadly, as our method works for any simply-connected open surfaces, it can also be applied to other disk-type structures such as human faces and tooth occlusal surfaces, or any simply-connected region of interest extracted from arbitrary biomedical shapes.

For future work, we plan to extend the framework for anatomical surfaces with other topologies. In particular, with the large number of prior approaches for surfaces with spherical topology, a natural next step is to develop a method that can be directly used for registering

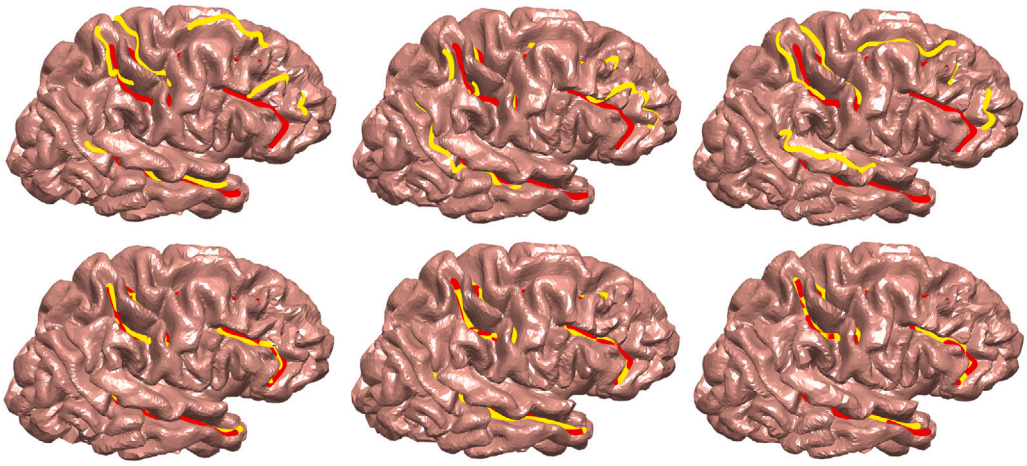

Fig. 14. Examples of 3D brain registration results. Three different brain surfaces are mapped to a template brain surface with or without using our proposed network. Each column shows the results for one surface. The top row shows three registration results without using our proposed network, where the registered sulcal landmark curves are highlighted in yellow and the target landmarks are in red. It can be observed that the landmarks are not well-aligned, as a small landmark mismatch on the parameterized planar domain will already lead to a large landmark mismatch in the final 3D registration result. The bottom row shows the registration results obtained by our framework, with the landmarks well-aligned.

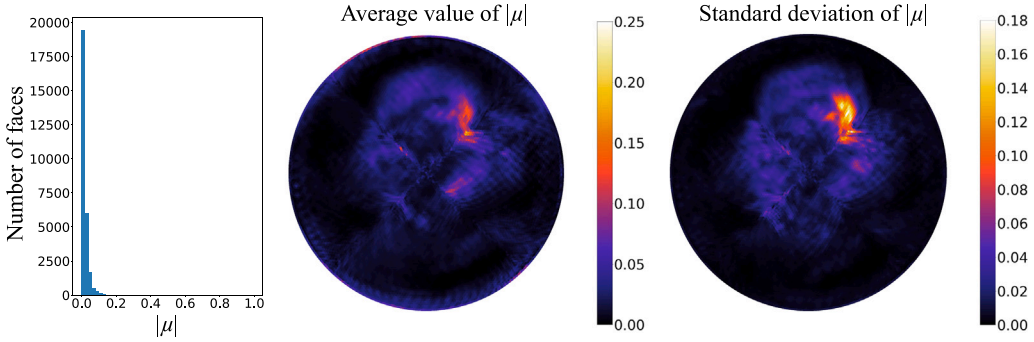

Fig. 15. Analysis of the Beltrami coefficient μ of all mappings generated by our framework. The first plot shows the histogram of all face-based $|\mu|$ values. The second and third plots show the average value and the standard deviation of $|\mu|$ evaluated on each triangular face on the unit disk for all mappings.

spherical surfaces and combine it with other tools such as the spherical harmonics and other available functions in FreeSurfer [75] for biomedical shape analysis.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The code is available on GitHub at https://github.com/ycguo97/BrainReg.

Acknowledgments

This work was supported in part by the National Science Foundation, United States under Grant No. DMS-2002103 (to Gary P. T. Choi), and HKRGC GRF under Project ID 14305919 (to Lok Ming Lui).

References

[1] 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 23 (2004) S2–S18.

[2] W. Zeng, J. Marino, K.C. Gurijala, X. Gu, A. Kaufman, Supine and prone colon registration using quasi-conformal mapping, IEEE Trans. Vis. Comput. Graphics 16 (6) (2010) 1348–1357.

[3] P. Thompson, A.W. Toga, A surface-based technique for warping three-dimensional images of the brain, IEEE Trans. Med. Imaging 15 (4) (1996) 402–417.

[4] B. Fischl, M.I. Sereno, A.M. Dale, Cortical surface-based analysis. II: Inflation, flattening, and a surface-based coordinate system, NeuroImage 9 (2) (1999) 195–207.

[5] B. Fischl, N. Rajendran, E. Busa, J. Augustinack, O. Hinds, B.T. Yeo, H. Mohlberg, K. Amunts, K. Zilles, Cortical folding patterns and predicting cytoarchitecture, Cereb. Cortex 18 (8) (2008) 1973–1980.

[6] E.C. Robinson, S. Jbabdi, M.F. Glasser, J. Andersson, G.C. Burgess, M.P. Harms, S.M. Smith, D.C. Van Essen, M. Jenkinson, MSM: A new flexible framework for multimodal surface matching, NeuroImage 100 (2014) 414–426.

[7] M.S. Floater, K. Hormann, Surface parameterization: A tutorial and survey, Adv. Multiresolution Geom. Model. (2005) 157–186.

[8] A. Sheffer, E. Praun, K. Rose, Mesh parameterization methods and their applications, Found. Trends Comput. Graph. Vis. 2 (2) (2007) 105–171.

[9] M. Desbrun, M. Meyer, P. Alliez, Intrinsic parameterizations of surface meshes, Comput. Graph. Forum 21 (3) (2002) 209–218.

[10] B. Lévy, S. Petitjean, N. Ray, J. Maillot, Least squares conformal maps for automatic texture atlas generation, ACM Trans. Graph. 21 (3) (2002) 362–371.

[11] P. Mullen, Y. Tong, P. Alliez, M. Desbrun, Spectral conformal parameterization, Comput. Graph. Forum 27 (5) (2008) 1487–1494.

[12] M.-H. Yueh, W.-W. Lin, C.-T. Wu, S.-T. Yau, An efficient energy minimization for conformal parameterizations, J. Sci. Comput. 73 (1) (2017) 203–227.

[13] N. Lei, F. Luo, S.-T. Yau, X. Gu, Computational conformal geometric methods for vision, in: Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging, Springer, 2021, pp. 1–52.

[14] S. Angenent, S. Haker, A. Tannenbaum, R. Kikinis, Conformal geometry and brain flattening, in: MICCAI, Springer, 1999, pp. 271–278.

[15] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, M. Halle, Conformal surface parameterization for texture mapping, IEEE Trans. Vis. Comput. Graphics 6 (2) (2000) 181–189.
[16] M.K. Hurdal, K. Stephenson, Cortical cartography using the discrete conformal approach of circle packings, NeuroImage 23 (supp-S1) (2004) S119–S128.
[17] X. Gu, Y. Wang, T.F. Chan, P.M. Thompson, S.-T. Yau, Genus zero surface conformal mapping and its application to brain surface mapping, IEEE Trans. Med. Imaging 23 (8) (2004) 949–958.
[18] P.T. Choi, K.C. Lam, L.M. Lui, FLASH: Fast landmark aligned spherical harmonic parameterization for genus-0 closed brain surfaces, SIAM J. Imaging Sci. 8 (1) (2015) 67–94.
[19] G.P.T. Choi, L.M. Lui, Recent developments of surface parameterization methods using quasi-conformal geometry, in: Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging, Springer, 2023, pp. 1483–1523.
[20] L.M. Lui, T.W. Wong, W. Zeng, X. Gu, P.M. Thompson, T.F. Chan, S.-T. Yau, Optimization of surface registrations using Beltrami holomorphic flow, J. Sci. Comput. 50 (3) (2012) 557–585.
[21] L.M. Lui, K.C. Lam, S.-T. Yau, X. Gu, Teichmüller mapping (T-map) and its applications to landmark matching registration, SIAM J. Imaging Sci. 7 (1) (2014) 391–426.
[22] G.P.-T. Choi, L.M. Lui, A linear formulation for disk conformal parameterization of simply-connected open surfaces, Adv. Comput. Math. 44 (1) (2018) 87–114.
[23] G.P.T. Choi, Y. Leung-Liu, X. Gu, L.M. Lui, Parallelizable global conformal parameterization of simply-connected surfaces via partial welding, SIAM J. Imaging Sci. 13 (3) (2020) 1049–1083.
[24] G.P.T. Choi, Efficient conformal parameterization of multiply-connected surfaces using quasi-conformal theory, J. Sci. Comput. 87 (3) (2021) 1–19.
[25] F.L. Bookstein, Thin-plate splines and the decomposition of deformations, IEEE Trans. Pattern Anal. Machine Intel. 16 (6) (1989) 567–585.
[26] S.C. Joshi, M.I. Miller, Landmark matching via large deformation diffeomorphisms, IEEE Trans. Image Process. 9 (8) (2000) 1357–1370.
[27] J. Glaunès, M. Vaillant, M.I. Miller, Landmark matching via large deformation diffeomorphisms on the sphere, J. Math. Imaging Vision 20 (1) (2004) 179–200.
[28] Y. Yin, E.A. Hoffman, K. Ding, J.M. Reinhardt, C.L. Lin, A cubic B-spline-based hybrid registration of lung CT images for a dynamic airway geometric model with large deformation, Phys. Med. Biol. 56 (1) (2011) 203.
[29] K.C. Lam, L.M. Lui, Landmark-and intensity-based registration with large deformations via quasi-conformal maps, SIAM J. Imaging Sci. 7 (4) (2014) 2364–2392.
[30] G.P.T. Choi, H.L. Chan, R. Yong, S. Ranjitkar, A. Brook, G. Townsend, K. Chen, L.M. Lui, Tooth morphometry using quasi-conformal theory, Pattern Recognit. 99 (2020) 107064.
[31] G.P.T. Choi, D. Qiu, L.M. Lui, Shape analysis via inconsistent surface registration, Proc. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci. 476 (2242) (2020) 20200147.
[32] T.W. Meng, G.P.-T. Choi, L.M. Lui, TEMPO: Feature-endowed Teichmüller extremal mappings of point clouds, SIAM J. Imaging Sci. 9 (4) (2016) 1922–1962.
[33] H.-L. Chan, H.-M. Yuen, C.-T. Au, K.C.-C. Chan, A.M. Li, L.-M. Lui, Quasi-conformal geometry based local deformation analysis of lateral cephalogram for childhood OSA classification, 2020, arXiv:2006.11408.
[34] X. Geng, D. Kumar, G.E. Christensen, Transitive inverse-consistent manifold registration, in: Biennial International Conference on Information Processing in Medical Imaging, Springer, 2005, pp. 468–479.
[35] I. Eckstein, A.A. Joshi, C.-C.J. Kuo, R. Leahy, M. Desbrun, Generalized surface flows for deformable registration and cortical matching, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2007, pp. 692–700.
[36] B.T. Yeo, M.R. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, P. Golland, Spherical demons: Fast diffeomorphic landmark-free surface registration, IEEE Trans. Med. Imaging 29 (3) (2009) 650–668.
[37] D. Pantazis, A. Joshi, J. Jiang, D.W. Shattuck, L.E. Bernstein, H. Damasio, R.M. Leahy, Comparison of landmark-based and automatic methods for cortical surface registration, Neuroimage 49 (3) (2010) 2479–2493.
[38] Y. Wang, L.M. Lui, T.F. Chan, P.M. Thompson, Optimization of brain conformal mapping with landmarks, in: MICCAI, Springer, 2005, pp. 675–683.
[39] L.M. Lui, Y. Wang, T.F. Chan, P. Thompson, Landmark constrained genus zero surface conformal mapping and its application to brain mapping research, Appl. Numer. Math. 57 (5–7) (2007) 847–858.
[40] L.M. Lui, S. Thiruvenkadam, Y. Wang, P.M. Thompson, T.F. Chan, Optimized conformal surface registration with shape-based landmark matching, SIAM J. Imaging Sci. 3 (1) (2010) 52–78.
[41] Y. LeCun, B. Boser, J.S. Denker, D. Henderson, R.E. Howard, W. Hubbard, L.D. Jackel, Backpropagation applied to handwritten zip code recognition, Neural Comput. 1 (4) (1989) 541–551.
[42] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: CVPR, 2016, pp. 770–778.
[43] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, ImageNet: A large-scale hierarchical image database, in: CVPR, IEEE, 2009, pp. 248–255.
[44] Q. Chen, Z. Li, L.M. Lui, A learning framework for diffeomorphic image registration based on quasi-conformal geometry, 2021, arXiv:2110.10580.
[45] H. Law, G.P.T. Choi, K.C. Lam, L.M. Lui, Quasiconformal model with CNN features for large deformation image registration, Inverse Prob. Imaging 16 (4) (2022) 1019–1046.
[46] D.I. Shuman, B. Ricaud, P. Vandergheynst, Vertex-frequency analysis on graphs, Appl. Comput. Harmon. Anal. 40 (2) (2016) 260–291.
[47] M. Defferrard, X. Bresson, P. Vandergheynst, Convolutional neural networks on graphs with fast localized spectral filtering, Adv. Neural Inf. Process. Syst. 29 (2016).
[48] M. Niepert, M. Ahmed, K. Kutzkov, Learning convolutional neural networks for graphs, in: International Conference on Machine Learning, PMLR, 2016, pp. 2014–2023.
[49] F.P. Such, S. Sah, M.A. Dominguez, S. Pillai, C. Zhang, A. Michael, N.D. Cahill, R. Ptucha, Robust spatial filtering with graph convolutional neural networks, IEEE J. Sel. Top. Signal Process. 11 (6) (2017) 884–896.
[50] P. Besson, T. Parrish, A.K. Katsaggelos, S.K. Bandt, Geometric deep learning on brain shape predicts sex and age, Comput. Med. Imaging Graph. 91 (2021) 101939.
[51] J. Masci, D. Boscaini, M. Bronstein, P. Vandergheynst, Geodesic convolutional neural networks on Riemannian manifolds, in: Proceedings of the IEEE International Conference on Computer Vision Workshops, 2015, pp. 37–45.
[52] D. Boscaini, J. Masci, E. Rodolà, M. Bronstein, Learning shape correspondence with anisotropic convolutional neural networks, Adv. Neural Inf. Process. Syst. 29 (2016).
[53] Z. Wu, G. Li, L. Wang, F. Shi, W. Lin, J.H. Gilmore, D. Shen, Registration-free infant cortical surface parcellation using deep convolutional neural networks, in: MICCAI, Springer, 2018, pp. 672–680.
[54] F. Zhao, S. Xia, Z. Wu, D. Duan, L. Wang, W. Lin, J.H. Gilmore, D. Shen, G. Li, Spherical U-net on cortical surfaces: Methods and applications, in: International Conference on Information Processing in Medical Imaging, Springer, 2019, pp. 855–866.
[55] S.-B. Seong, C. Pae, H.-J. Park, Geometric convolutional neural network for analyzing surface-based neuroimaging data, Front. Neuroinform. 12 (2018) 42.
[56] F. Zhao, Z. Wu, L. Wang, W. Lin, J.H. Gilmore, S. Xia, D. Shen, G. Li, Spherical deformable U-Net: Application to cortical surface parcellation and development prediction, IEEE Trans. Med. Imaging 40 (4) (2021) 1217–1228.
[57] F. Zhao, Z. Wu, F. Wang, W. Lin, S. Xia, D. Shen, L. Wang, G. Li, S3Reg: Superfast spherical surface registration based on deep learning, IEEE Trans. Med. Imaging 40 (8) (2021) 1964–1976.
[58] J. Cheng, A.V. Dalca, B. Fischl, L. Zöllei, Cortical surface registration using unsupervised learning, NeuroImage 221 (2020) 117161.
[59] G. Balakrishnan, A. Zhao, M.R. Sabuncu, J. Guttag, A.V. Dalca, VoxelMorph: A learning framework for deformable medical image registration, IEEE Trans. Med. Imaging 38 (8) (2019) 1788–1800.
[60] M. Ono, S. Kubik, C.D. Abernathey, Atlas of the Cerebral Sulci, Thieme Medical, New York, 1990.
[61] E.R. Sowell, P.M. Thompson, D. Rex, D. Kornsand, K.D. Tessner, T.L. Jernigan, A.W. Toga, Mapping sulcal pattern asymmetry and local cortical surface gray matter distribution in vivo: Maturation in perisylvian cortices, Cereb. Cortex 12 (1) (2002) 17–26.
[62] J. Hill, D. Dierker, J. Neil, T. Inder, A. Knutsen, J. Harwell, T. Coalson, D. Van Essen, A surface-based analysis of hemispheric asymmetries and folding of cerebral cortex in term-born human infants, J. Neurosci. 30 (6) (2010) 2268–2276.
[63] S.H. Joshi, R.P. Cabeen, A.A. Joshi, B. Sun, I. Dinov, K.L. Narr, A.W. Toga, R.P. Woods, Diffeomorphic sulcal shape analysis on the cortex, IEEE Trans. Med. Imaging 31 (6) (2012) 1195–1212.
[64] J.-F. Mangin, D. Riviere, A. Cachia, E. Duchesnay, Y. Cointepas, D. Papadopoulos-Orfanos, D.L. Collins, A.C. Evans, J. Régis, Object-based morphometry of the cerebral cortex, IEEE Trans. Med. Imaging 23 (8) (2004) 968–982.
[65] G. Auzias, O. Colliot, J.A. Glaunes, M. Perrot, J.-F. Mangin, A. Trouve, S. Baillet, Diffeomorphic brain registration under exhaustive sulcal constraints, IEEE Trans. Med. Imaging 30 (6) (2011) 1214–1227.
[66] Z. Tu, S. Zheng, A.L. Yuille, A.L. Reiss, R.A. Dutton, A.D. Lee, A.M. Galaburda, I. Dinov, P.M. Thompson, A.W. Toga, Automated extraction of the cortical sulci based on a supervised learning approach, IEEE Trans. Med. Imaging 26 (4) (2007) 541–552.
[67] C.-Y. Kao, M. Hofer, G. Sapiro, J. Stern, K. Rehm, D.A. Rottenberg, A geometric method for automatic extraction of sulcal fundi, IEEE Trans. Med. Imaging 26 (4) (2007) 530–540.
[68] A.A. Joshi, D.W. Shattuck, R.M. Leahy, A method for automated cortical surface registration and labeling, in: International Workshop on Biomedical Image Registration, Springer, 2012, pp. 180–189.
[69] A. Le Troter, G. Auzias, O. Coulon, Automatic sulcal line extraction on cortical surfaces using geodesic path density maps, NeuroImage 61 (4) (2012) 941–949.
[70] S. Zhang, R. Wang, Z. Han, S. Yu, H. Gao, X. Jiang, T. Zhang, A DICCCOL-based K-nearest landmark detection method for identifying common and consistent 3-hinge gyral folding landmarks, Chaos Solitons Fractals 158 (2022) 112018.
[71] F.P. Gardiner, N. Lakic, Quasiconformal Teichmüller Theory, no. 76, American Mathematical Soc., 2000.

[72] L.M. Lui, K.C. Lam, T.W. Wong, X. Gu, Texture map and video compression using beltrami representation, SIAM J. Imaging Sci. 6 (4) (2013) 1880–1902.
[73] P.T. Choi, L.M. Lui, Fast disk conformal parameterization of simply-connected open surfaces, J. Sci. Comput. 65 (3) (2015) 1065–1090.
[74] D.S. Marcus, A.F. Fotenos, J.G. Csernansky, J.C. Morris, R.L. Buckner, Open access series of imaging studies: Longitudinal MRI data in nondemented and demented older adults, J. Cogn. Neurosci. 22 (12) (2010) 2677–2684.
[75] B. Fischl, FreeSurfer, NeuroImage 62 (2) (2012) 774–781.
[76] G. Donato, S. Belongie, Approximate thin plate spline mappings, in: Computer Vision—ECCV 2002: 7th European Conference on Computer Vision Copenhagen, Denmark, May 28–31, 2002 Proceedings, Part III 7, Springer, 2002, pp. 21–31.