Medical Image Analysis 90 (2023) 102974

Contents lists available at ScienceDirect

Medical Image Analysis

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

Learning joint surface reconstruction and segmentation, from brain images to cortical surface parcellation

Karthik Gopinath *, Christian Desrosiers, Herve Lombaert

ETS Montreal, Canada

ARTICLE INFO

Keywords:
Surface reconstruction
Cortical parcellation

ABSTRACT

Reconstructing and segmenting cortical surfaces from MRI is essential to a wide range of brain analyses. However, most approaches follow a multi-step slow process, such as a sequential spherical inflation and registration, which requires considerable computation times. To overcome the limitations arising from these multi-steps, we propose SegRecon, an integrated end-to-end deep learning method to jointly reconstruct and segment cortical surfaces directly from an MRI volume in one single step. We train a volume-based neural network to predict, for each voxel, the signed distances to multiple nested surfaces and their corresponding spherical representation in atlas space. This is, for instance, useful for jointly reconstructing and segmenting the white-to-gray-matter interface and the gray-matter-to-CSF (pial) surface. We evaluate the performance of our surface reconstruction and segmentation method with a comprehensive set of experiments on the MindBoggle, ABIDE and OASIS datasets. Our reconstruction error is found to be less than 0.52 mm and 0.97 mm in terms of average Hausdorff distance to the FreeSurfer generated surfaces. Likewise, the parcellation results show over 4% improvements in average Dice with respect to FreeSurfer, in addition to an observed drastic speed-up from hours to seconds of computation on a standard desktop station.

1. Introduction

Brain surface analysis requires the accurate reconstruction and segmentation of cortical surfaces from MRI volumes (Querbes et al., 2009; Glasser et al., 2016; Billot et al., 2023). Standard surface processing pipelines for reconstructing cortical surfaces (Fischl et al., 2004; Dahnke et al., 2013; Kim et al., 2005; Kriegeskorte and Goebel, 2001; Shattuck and Leahy, 2002) and hippocampus (Styner et al., 2006; Puonti et al., 2016) follow a sequence of costly operations that often include: white matter segmentation, surface mesh generation from the segmentation masks, mesh smoothing and projection to a sphere, topological correction of the projected mesh, and fine-tuning of re-projected mesh on the segmented volume. The cortical surfaces are segmented into neuroanatomical parcels in a subsequent and highly-expensive step. Such segmentation can take several hours to finish, which involves the re-projection of each surface to a sphere via a metric-preserving inflation process, registration to a spherical atlas (Fischl et al., 1999; Klein and Tourville, 2012) and cortical parcellation using atlas labels (Desikan et al., 2006).

Recently, Henschel et al. (2020) developed a framework called FastSurfer using deep learning that accelerates the processing times for brain segmentation and spectral embedding for registration to a spherical atlas. Despite reducing computation times considerably, this pipeline still performs reconstruction and segmentation in two consecutive steps. To overcome this limitation, Cruz et al. (2021), Ma et al. (2021), Hoopes et al. (2022) and Bongratz et al. (2022) proposed a deep learning based models for cortical surface reconstruction. This method draws inspiration from Park et al. (2019), and samples points on a reference grid of arbitrary resolution to reconstruct a surface without the need for an explicit segmentation step. However, this process is highly expensive in terms of both computation and memory for detailed surfaces with hundreds of thousands of points. Additionally, DeepCSR only performs surface reconstruction, and cannot be used for parcellation which is one of the most time-costly operations in conventional neuroimaging pipelines (Fischl et al., 2004). Similarly, (Ma et al., 2021) proposed a 3D deep learning framework for pial surface reconstruction. Approaches that directly operate (Wu et al., 2019; López-López et al., 2020) or learn on surface data (Gopinath et al., 2019; Lombaert et al., 2015; Gopinath et al., 2020) have been used for cortical parcellation, but are designed to process single surfaces separately for each subjects. Spectral embeddings of surface meshes in a low-dimensional space can be exploited to predict cortical parcellation labels (Lombaert et al., 2015; Germanaud et al., 2012). However, a major limitation of these early learning approaches is that mesh nodes are considered separately instead of jointly. Recent work has proposed using graph convolutional

* Correspondence to: ETS Montreal, Computer and Software Engineering, 1100 Notre Dame St. W., Montreal QC, H3C 1K3, Canada.
E-mail address: karthik.gopinath.1@ens.etsmtl.ca (K. Gopinath).

https://doi.org/10.1016/j.media.2023.102974
Received 28 August 2021; Received in revised form 14 September 2023; Accepted 16 September 2023
Available online 22 September 2023

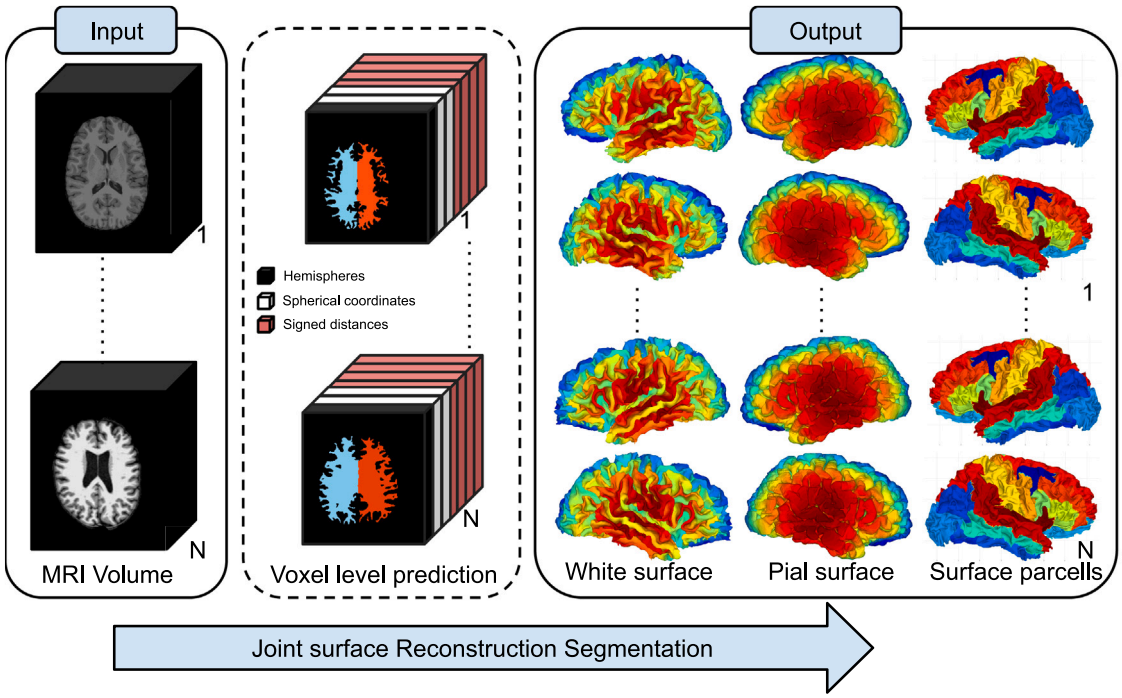

Fig. 1. Overview of the algorithm – Our proposed method for joint brain surface reconstruction and segmentation from images. On the left is input: MRI volume X. In the middle are the intermediate learned voxel level predictions from the 3D CNN model. These include for each voxel v: hemisphere label $h_v(X)$, azimuthal angle $\gamma_v(X)$ and polar angle $\phi_v(X)$ in the spherical atlas space, signed distance to white matter surface $d_v^w(X)$, and signed distance to pial surface $d_v^p(X)$. On the right are the reconstructed white and pial surface along with cortical parcels for each hemisphere of the brain.

networks (GCN) (Gopinath et al., 2019; Wu et al., 2019; He et al., 2020; Gopinath et al., 2020) to exploit the connectivity information of a mesh graph. Similarly, DBPN (Zhang and Wang, 2019) introduced a two-stage spatial graph convolution network and ASEGAT (Li et al., 2022) proposed to use an anatomically constrained squeeze-and-excitation self attention graph network to perform a cortical parcellation in the original mesh space. While such strategy provides an accurate and faster parcellation of the cortical surface, it is highly sensitive to the quality of the surface reconstruction step. For instance, small errors or holes in the reconstructed cortical mesh may causing graph convolution methods to fail at parcellation tasks. In a similar manner, SPHARM-Net (Ha and Lyu, 2022a) uses a spherical harmonics-based convolutional neural network for vertex-wise cortical parcellation. Furthermore, recent parcellation methods such as (Zhao et al., 2019, 2021; Parvathaneni et al., 2019; Ha and Lyu, 2022b) rely on inflated spheres of the cortical surface to perform spherical convolution-based parcellation. Extracting these cortical meshes, inflating them and registering them to spheres are typically computationally expensive steps, often ignored by these methods.

We propose SegRecon, a novel deep learning model for the joint reconstruction and parcellation of cortical surfaces. Our end-to-end model works directly on MRI volumes and predicts a dense set of surface points along with their corresponding parcellation labels (see Fig. 1). A 3D CNN based UNet (Çiçek et al., 2016) predicts for each input voxel of the volume, the brain hemisphere, its signed distance to the nested surfaces (white matter and pial surfaces) of that hemisphere, and the spherical coordinates in the registered atlas space. We use this predicted nested signed distance surfaces for surface reconstruction and the spherical coordinates for surface parcellation. By learning to solve this multi-task problem, the network can be used to reconstruct and segment brain surfaces efficiently and in a topologically-accurate manner (Bazin and Pham, 2007).

The main contributions of our work are the following:

- To our knowledge, we propose the first deep learning model for the joint reconstruction of multiple nested surfaces and their segmentation, with an application on brain surfaces. This contrasts with existing approaches, which either perform surface reconstruction and segmentation in separate steps (Henschel et al., 2020), are limited to reconstruction (Cruz et al., 2021), or require a pre-generated mesh as input (Gopinath et al., 2019; Wu et al., 2019; Lombaert et al., 2015);
- Compared to the current surface reconstruction learning approaches (such as DeepCSR), the proposed method implements a fully-convolutional architecture that densely predicts the location of all input voxels relative to cortical surfaces, in a single feed-forward pass. Our method also leverages a novel surface reconstruction loss that controls the minimum and maximum distance between white matter and pial surfaces (i.e., cortical thickness), thereby ensuring that these surfaces never cross;
- We present a comprehensive set of experiments involving three publicly-available datasets, i.e., MindBoggle (Klein et al., 2017), OASIS (Marcus et al., 2007) and ABIDE-I (Di Martino et al., 2014), that compare the surface reconstruction and segmentation accuracy of our method against several baselines. Our results demonstrate the major advantages of our method over standard brain surface analysis pipelines are recent approaches for cortical parcellation. With respect to the widely-used FreeSurfer software, our method generates surfaces with an average Hausdorff distance less than 0.52 mm, while boosting the parcellation Dice by 4.3% and being several orders of magnitude faster.

The proposed work is a substantial extension of our MICCAI article (Gopinath et al., 2021). The major extensions are: firstly, the reconstruction of multiple nested surfaces extracts both white matter and pial surfaces; secondly, a novel surface reconstruction loss ensures that the white and pial surfaces never cross; finally, the evaluation of the surface reconstruction is made on two independent test sets with varying factors, such as image acquisition, processing parameters, age, and cortical surface alterations.

In the next section, we present our proposed joint reconstruction and segmentation approach, describing in detail the network architecture, training losses and inference steps. The performance of our

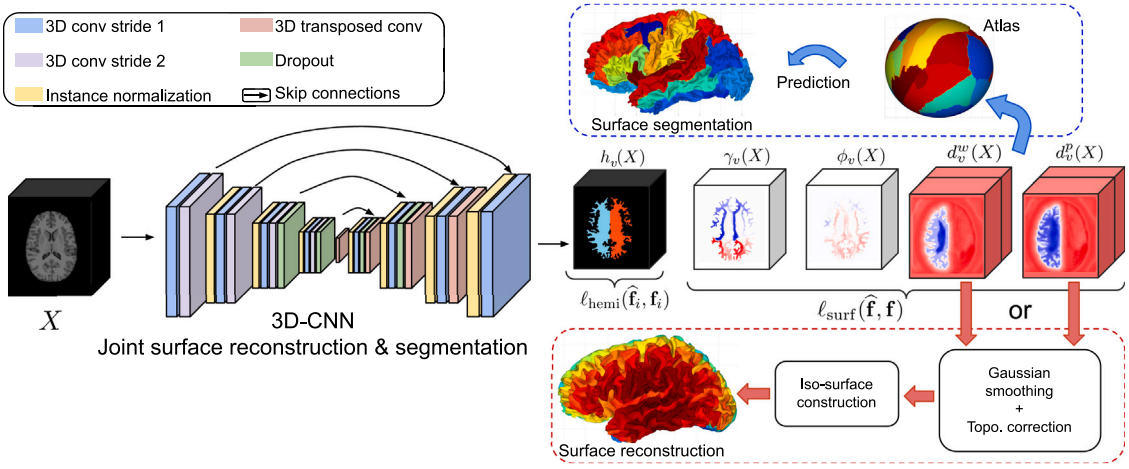

Fig. 2. Overview of SegRecon architecture: The 3D-CNN model takes as input MRI volume X for joint learning of the signed distance to white-to-gray matter interface, gray-matter-to-CSF interface and its corresponding spherical coordinates in the registered atlas space. (*Red*) The cortical white matter or pial surface is reconstructed by applying Gaussian smoothing and topological correction on the predicted signed distance map prediction $d_v^w(X)$ or $d_v^p(X)$, followed by iso-surface reconstruction via the Marching Cubes algorithm. (*Blue*) In parallel, the predicted spherical atlas coordinates ($\gamma_v(X)$, $\phi_v(X)$) and hemisphere label ($h_v(X)$) are used to propagate atlas parcellation labels to near-surface voxels v. An illustration of the left hemisphere white matter surface is shown here.

method is then evaluated on the MindBoggle (Klein et al., 2017), OASIS (Marcus et al., 2007) and ABIDE-I (Di Martino et al., 2014) datasets. The ablation study and comparison to the state-of-art in our experiments demonstrate the important benefits of our method.

2. Method

An overview of SegRecon is shown in Fig. 2 with the end-to-end surface construction and segmentation steps illustrated. Let $\mathcal{D} = \{(X_i, S_i^W, S_i^P, Y_i)\}_{i=1}^n$ be a training set where each example is composed of: a 3D volume $X_i \in \mathbb{R}^{|\Omega|}$ with voxel set $\Omega \subset \mathbb{Z}^3$, a white matter surface $S_i^W \in \mathbb{R}^{m\times 3}$ defined by m points, a pial surface $S_i^P \in \mathbb{R}^{n\times 3}$ defined by n points, and a one-hot encoded segmentation $Y_i \in [0,1]^{m\times c}$ of the white matter surface, where c is the number of segmentation classes. The goal is to learn a model f parameterized by θ which maps an input 3D volume X to a white matter surface S^W with corresponding parcel labels Y, and a pial surface S^P.

One of the main challenges in this task comes from the disparity between the well-defined grid space of images X and the domain of surfaces S^W and S^P, where the number of points can vary from one surface to another and points can lie anywhere in 3D space. In Cruz et al. (2021), this problem is solved by giving as input to model f both the image X and a query point $p \in \mathbb{R}^3$ in the template space. The model then predicts if p belongs to the surface in X or, alternatively, its distance to this surface. To reconstruct a surface at inference time, the model is queried over a fixed reference grid. While this strategy allows reconstructing a surface at arbitrary resolution, it suffers from two important drawbacks. First, since the template points which can be in the hundreds of thousands are queried independently, reconstructing a surface requires significant time and computation. Moreover, unlike dense prediction approaches, this strategy does not exploit the spatial relationship between points. Last, because feature maps need to be computed for the whole 3D volume X, it also needs a large amount of memory.

To overcome these drawbacks, we instead learn a model that densely projects voxels of the input volume X to a spherical atlas space. Specifically, f maps each voxel $v \in \Omega$ to a vector

$$f_v(X) = \left[d_v^w, d_v^p, \phi_v, \gamma_v, h_v^{lh}, h_v^{rh}, h_v^{bg}\right], \tag{1}$$

where d_v^w is the signed distance from v to its nearest surface point, such that $d_v^w \leq 0$ if v is inside the surface else $d_v^w > 0$. Similarly, d_v^p is the signed distance from v to its nearest pial surface. ϕ_v, γ_v are the polar angle and azimuthal angle of $v \in \Omega$ defining its position in the spherical atlas, and h_v^{lh}, h_v^{rh}, $h_v^{bg} \in [0,1]$ are the probabilities that v is in the left hemisphere, right hemisphere and background, respectively. Here, polar and azimuthal angles are normalized so to lie in the $[-1,1]$ range. A further topological correction step (Bazin and Pham, 2007) over the predicted surface points prevents the extraction of critical points yielding topological defects. The resulting white and pial surfaces are defined implicitly as the 0-levelset of their respective distance map and can be efficiently reconstructed using an iso-surface extraction algorithm such as the Marching Cubes (Lorensen and Cline, 1987).

2.1. Training the model

Denote $\hat{\mathbf{f}}_i = f(X_i)$ as the predicted vector for an image X_i and let $\mathbf{f}_i$ be the corresponding ground-truth. To train the model, we use the following loss function

$$\mathcal{L}(\theta; \mathcal{D}) = \sum_{i=1}^{n} \ell_{\text{wsurf}}(\hat{\mathbf{f}}_i, \mathbf{f}_i) + \ell_{\text{psurf}}(\hat{\mathbf{f}}_i, \mathbf{f}_i) + \lambda_1 \ell_{\text{hemi}}(\hat{\mathbf{f}}_i, \mathbf{f}_i) + \lambda_2 \ell_{\text{thick}}(\hat{\mathbf{f}}_i), \tag{2}$$

The first loss term, ℓ_{wsurf}, ensures that the signed distance of voxels to the white matter surface, as well as their position in the spherical atlas space, are well predicted. Dropping index i for simplicity, it is defined as

$$\ell_{\text{wsurf}}(\hat{\mathbf{f}}, \mathbf{f}) = \sum_{v\in\Omega} \mathbb{1}_{|d_v^w|\leq\epsilon} \cdot \Big[(\hat{d}_v^w - d_v^w)^2 + \min\{(\hat{\phi}_v - \phi_v)^2, (1+\hat{\phi}_v - \phi_v)^2\} + \min\{(\hat{\gamma}_v - \gamma_v)^2, (1+\hat{\gamma}_v - \gamma_v)^2\} \Big]. \tag{3}$$

where $\mathbb{1}_P$ is the indicator function, equal to 1 if predicate P is true else, 0 otherwise. We only consider voxels within a distance of ϵ to the nearest white matter surface point in order to focus learning on relevant points close to our surface. This is achieved with function $\mathbb{1}_{|d_v^w|\leq\epsilon}$ in Eq. (3). Additionally, we consider the non-uniqueness of spherical coordinates (e.g., $-\pi \equiv \pi$) by computing, for each angle, the minimum L_2 distance from the predicted angle or this angle plus 1 to the ground-truth. Using the minimum function of the two cyclic angles aids in optimizing the training of both polar and azimuth angles.. The distance d_v^w is, therefore, defined between the center of the voxel v in image space and the nearest point on white matter surface S^W. In this work, we use the white matter surface mesh generated by FreeSurfer for training. The sign of d_v^w is determined using the white-matter segmentation mask, with voxels inside the white matter having

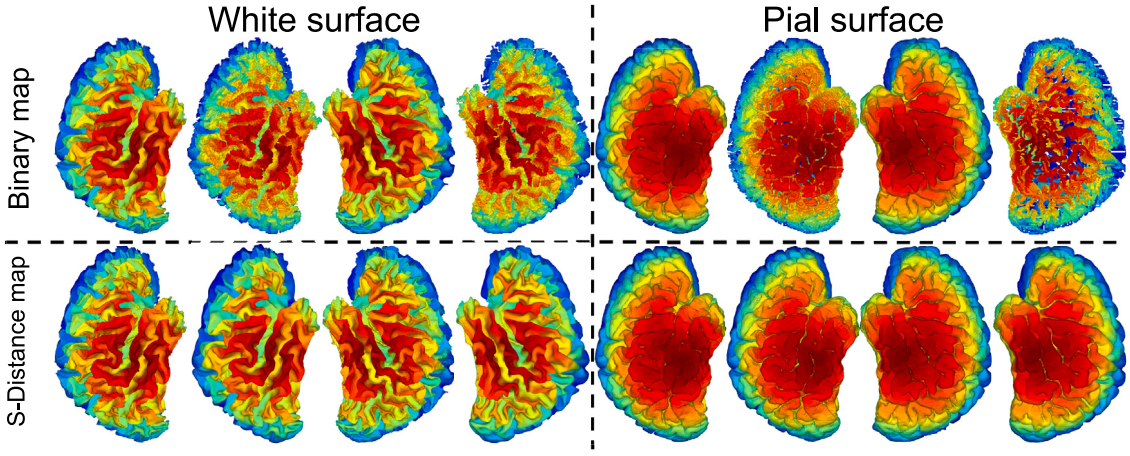

Fig. 3. Surface reconstruction visualization: Comparison of cortical white and pial surface reconstructed by our method (even columns) with FreeSurfer generated surfaces (odd columns). The first row shows the surfaces reconstructed by our method using binary mask prediction. A reconstruction error of 2.8 mm in Hausdorff distance (HD) is seen with the use of binary mask (BW) vs. a reconstruction error of 0.9 mm in Hausdorff distance (HD) with the use of a signed distance (SD) map. The reconstruction in the last column for pial surface highlights the downside of using BW masks with irregular surface reconstruction.

a negative distance. Likewise, the ground-truth spherical coordinates ϕ_v and γ_v are obtained using FreeSurfer (Fischl et al., 2004) with the Desikan–Killiany–Tourville (DKT) atlas (Klein and Tourville, 2012).

The second loss term, ℓ_{psurf}, ensures that the signed distance of voxels to the nearest pial surface is predicted accurately. We define it as

$$\ell_{\text{psurf}}(\hat{\mathbf{f}},\mathbf{f}) = \sum_{v\in\Omega} \mathbb{1}_{|d_v^p|\le\epsilon} \cdot (\hat{d}_v^p - d_v^p)^2. \tag{4}$$

where d_v^p is the distance defined between the center of the voxel v and its closest point on pial surface S^P obtained by FreeSurfer pial meshes. The sign of the distance d_v^p is estimated using the brain segmentation mask with voxels inside the brain mask having negative distance. Similar to ℓ_{wsurf} in Eq. (3), $\mathbb{1}_P$ is the indicator function used to restrict the training to the useful voxels within a distance of ϵ to the closest pial surface.

The third term, ℓ_{hemi} enables the network to predict if a voxel v lies in the left hemisphere (lh), in the right hemisphere (rh) or is outside both (bg). This prediction is necessary since the surface atlas is defined separately for each hemisphere. Here, we use cross-entropy as loss function:

$$\ell_{\text{hemi}}(\hat{\mathbf{f}},\mathbf{f}) = -\sum_{v\in\Omega} \sum_{c\in\{lh,rh,bg\}} h_v^c \log \hat{h}_v^c. \tag{5}$$

The ground-truth hemisphere masks are once again obtained from FreeSurfer.

Since the white matter and pial surfaces are reconstructed from two separate predictions, it may happen that predicted surfaces are near to the ground-truth while still violating anatomical constraints. For example, in very thin regions of the cortex, the reconstructed surfaces may overlap or even cross each other. To avoid this problem, we add a last term to the loss function, ℓ_{thick}, which controls the minimum and maximum distance between the surfaces:

$$\ell_{\text{thick}}(\hat{\mathbf{f}}) = \sum_{v\in\Omega} \mathbb{1}_{d_v^p\le 0} \cdot \left[\max(\hat{d}_v^p - \hat{d}_v^w + t_{min}, 0) + \max(\hat{d}_v^w - \hat{d}_v^p - t_{max}, 0)\right] \tag{6}$$

where t_{min} and t_{max} are the minimum and maximum allowed inter-surface distances (cortical thickness). These hyperparameters can be set based on the dataset ground-truth or some clinical reference. For instance, a cortical thickness range from 1 to 4.5 mm is reported in Fischl and Dale (2000). We use similar values in this work: $t_{min} = 0$ and $t_{max} = 5$ mm. Effectively, this prevents surfaces from crossing each other or separating beyond 5 mm. As defined in Eq. (6), this penalty is only calculated for voxels inside the pial surface, i.e., voxels v such that $\mathbb{1}_{d_v^p\le 0}$.

2.2. Surface reconstruction and segmentation

Once the network is trained, it can be used to reconstruct and segment surfaces directly from a test volume X. First, we feed the volume to the network to obtain a prediction vector for all voxels. Since the network is fully-convolutional, this can be done efficiently in a single feed-forward pass. Next, we apply a small-width Gaussian filter on the predicted 3D white matter surface distance map $\hat{d}^w$ using a single convolution operation and employ a topological correction step (Bazin and Pham, 2007) to overcome any defects in the surface. The same steps are followed to extract the 3D pial surface using distance map $\hat{d}^p$. The surface is reconstructed using the Marching Cubes algorithm (Lorensen and Cline, 1987) on the 0-levelset of its predicted signed distance map, smoothed with a Gaussian kernel and topologically corrected with the method of Bazin and Pham (2007).

To segment the surface, we first compute the near-surface voxels in each hemisphere as follows:

$$S^c = \left\{v \in \Omega \mid |d_v| \le \epsilon \wedge c = \arg\max_{c'} \hat{h}_v^{c'}\right\}, \quad c \in \{lh, rh\}. \tag{7}$$

Next, we find the nearest-neighbor to a given reference atlas R^c for all the near-surface voxels $v \in S^c$ using their predicted angles $\hat{\phi}_v$ and $\hat{\gamma}_v$. The segmentation labels from this reference atlas R^c are later projected back to the near-surface voxels S^c. A majority voting across multiple atlases is eventually applied to obtain the final parcellation labels.

2.3. Implementation details

The overall architecture of SegRecon is shown in Fig. 2. As an input, we provide the skull-stripped, intensity normalized 3D T1-MRI volume. We use a 3D-UNet architecture similar to Çiçek et al. (2016) in order to map the input voxel to a point in the spherical atlas space.

We apply a softmax activation in the first three output channels to predict the probability of a voxel belonging to the background, left hemisphere, or right hemisphere. The polar and azimuthal angles, $\hat{\phi}_v$ and $\hat{\gamma}_v$, are predicted with a tanh activation. The FreeSurfer generated cortical meshes and the corresponding registered spheres are used to pre-compute the surface distances and interpolate the ground-truth spherical coordinates for each voxels respectively. The last two output

Table 1
Performance of surface reconstruction: The reconstruction error (mm) measured between surface meshes generated by our SegRecon method and FreeSurfer generated meshes.

Experiment	Left White Matter			Right White Matter			Left Pial			Right Pial		
	CD	AAD	HD	CD	AAD	HD	CD	AAD	HD	CD	AAD	HD
SD map	1.760(0.23)	0.337(0.02)	0.850(0.06)	1.772(0.20)	0.346(0.02)	0.876(0.05)	2.234(0.38)	0.419(0.03)	0.835(0.07)	2.287(0.29)	0.438(0.02)	0.896(0.09)
BW map	5.019(0.35)	0.630(0.04)	3.104(0.60)	4.954(0.27)	0.632(0.04)	3.032(0.47)	5.034(0.34)	0.806(0.04)	2.730(0.60)	4.970(0.27)	0.829(0.04)	2.509(0.47)
OASIS CN	1.820(0.68)	0.360(0.05)	0.805(0.12)	1.726(0.45)	0.354(0.04)	0.806(0.16)	2.949(0.91)	0.508(0.07)	1.405(0.24)	2.893(0.60)	0.512(0.06)	1.424(0.24)
OASIS AD	2.003(0.75)	0.396(0.05)	0.843(0.11)	1.973(0.70)	0.385(0.05)	0.819(0.09)	2.975(0.78)	0.554(0.08)	1.374(0.24)	2.948(0.57)	0.552(0.06)	1.389(0.22)
ABIDE CN	1.995(0.80)	0.357(0.05)	0.934(0.17)	1.920(0.52)	0.357(0.04)	0.936(0.15)	3.182(1.29)	0.470(0.13)	1.385(0.54)	3.350(1.13)	0.494(0.12)	1.536(0.54)
ABIDE ASD	1.946(0.70)	0.353(0.04)	0.936(0.16)	1.931(0.68)	0.355(0.03)	0.942(0.14)	3.159(1.34)	0.463(0.13)	1.378(0.54)	3.286(1.19)	0.480(0.11)	1.510(0.53)

The Chamfer distance (CD), absolute average distance (AAD) and the 90th percentile of Hausdorff distance (HD) are reported for white and pial surface meshes of both left and right hemisphere. The first row highlights the performance of our method with surface reconstruction from signed distance (SD) map, whereas the second row shows the limitation of using binary (BW) segmentation map for surface reconstruction. The third and fourth rows report the reconstruction error on OASIS dataset. Likewise, the last two rows show the results on ABIDE dataset. Reconstruction metrics are reported on unseen OASIS and ABIDE data containing both healthy and AD/ASD subjects. The standard deviation is reported inside the parenthesis.

channels produce the signed distance map $\hat{d}_v^w$ and $\hat{d}_v^p$ for each voxel v. The λ_1 and λ_2 are set to 0.05 and 1 respectively. The network parameters, θ, are optimized using a stochastic gradient descent with the Adam optimizer (Kingma and Ba, 2014). During training, we pick the maximum distance of surface voxels in Eq. (3) to be $\epsilon = 2.5$ mm, which corresponds to the overall average thickness reported in Fischl and Dale (2000). We employ the (Lewiner et al., 2003) implementation of the marching cubes in *scikit-image* on the 0-levelset of its predicted signed distance map, smoothed with a Gaussian kernel of sigma = 0.5 mm to ensure that the reconstructed surfaces are one-connected meshes with no cracks or tears. We use an Intel i7 desktop machine with 16Gb RAM and Nvidia RTX 2080 GPU for our work.

3. Experiments and results

To benchmark the performance of our method, we use one of the largest publicly-available dataset containing manual surface parcellation, MindBoggle (Klein et al., 2017). This dataset contains 101 subjects with MRI volumes, FreeSurfer processed meshes, and 32 manually-labeled cortical parcels. We split the dataset randomly into training, validation, and testing using a ratio of 70-10%–20%. We also use the ABIDE-I (Di Martino et al., 2014) and OASIS (Marcus et al., 2007) databases as independent test sets to measure the surface reconstruction error of our method with FreeSurfer-generated cortical and pial surfaces. The ABIDE dataset contains brain surfaces for 1035 subjects with 530 healthy and 505 autism spectrum disorder (ASD) subjects. Likewise, the OASIS dataset comprises a total of 226 brain surfaces from 93 healthy subjects and 133 subjects with Alzheimer's disease (AD). These two datasets are used to validate the robustness of the method to various factors, including image acquisition and processing parameters, age (the majority of ABIDE subjects are children or adolescents, while most OASIS subjects are elders) and cortical surface alterations (AD subjects in OASIS have a thinner cortex, on average, resulting from the neurodegenerative disease).

In a first experiment, we validate the benefit of predicting a signed distance map for surface reconstruction compared to predicting a binary mask. For this experiment, we train the model using only data from MindBoggle, and measure the reconstruction error on subjects from the ABIDE and OASIS datasets. The qualitative results of the reconstructed surface are shown in Figs. 3 and 4. In the next experiment, we evaluate the impact of varying the reference atlas template for predicting parcellation labels, and show that a robust parcellation can be achieved by combining the predictions from multiple atlases. Finally, we highlight the advantages of our joint reconstruction and parcellation model against state-of-the-art methods.

3.1. Surface reconstruction

To assess the quality of reconstructed surfaces quantitatively, we use the Chamfer distance (CD) (Park et al., 2019), absolute average distance (AAD) (Cruz et al., 2021) and Hausdorff distance (HD) (Cruz et al., 2021). Chamfer distance is a widely-used evaluation metric defined as the sum between the average squared-distance from predicted surface points to their nearest point on the ground truth surface and the average squared-distance from ground truth surface points to their nearest point on the predicted surface. Similarly, AAD measures the mean *absolute* nearest-neighbor distance between the predicted and ground-truth surface points, averaging values computed in both directions (predicted to ground-truth and the opposite). Finally, HD computes the maximum distance between a point on a surface and its nearest point on the other surface. As in Cruz et al. (2021), to minimize the impact of outliers, we consider the 90th percentile of nearest-neighbor distances for one direction and the other. For all metrics, a lower value in mm indicates a better surface reconstruction.

We first evaluate the benefit of using a signed distance (SD) map, when reconstructing the white and pial cortical surfaces, by comparing it against using a binary mask (BW). To predict the binary mask, we use an architecture similar to the one in Fig. 2 where the last two output maps (corresponding to white matter and pial surfaces) are generated with sigmoid activations. As reported in Table 1, an improvement in CD from 5.0 mm to 1.7 mm is obtained when a signed distance map is used for white matter surface reconstruction. A similar improvement over the binary mask approach is also observed in terms of AAD and HD. Qualitative results, presented in Fig. 3, show that the meshes reconstructed using signed distance maps are more regular and closer to FreeSurfer-generated meshes, compared to those obtained with binary masks.

Our surface reconstruction method was also tested on the OASIS and ABIDE datasets, not used for training, to evaluate its robustness. As found in Table 1, our method obtained a mean AAD below 0.35 mm and mean HD less than 0.93 mm for the white matter surface, in both datasets. Similarly, reconstructed pial surfaces in both datasets have a mean AAD no greater than 0.58 mm and mean HD less than 1.53 mm for pial surfaces in both datasets. These results, obtained for subjects of very different ages and with cortical alterations, are comparable to those obtained for the MindBoggle test set. The qualitative results in Fig. 4 validate the visual similarity in surface reconstruction of our method, across datasets. Furthermore, we present a visualization of the surface reconstruction geometry overlaid with the curvatures of the white surface on inflated surfaces (Fig. 5). Additionally, the measured Euler characteristic of 2 and the presence of one large connected component in the reconstructed surface validate the topological accuracy of our method.

3.2. Effect of reference atlas on parcellation

Instead of predicting class probabilities for each voxel, as in standard 3D segmentation networks, the proposed network predicts spherical atlas coordinates (i.e., angles $\hat{\phi}_v$ and $\hat{\gamma}_v$). This has two important advantages: *(i)* considerably reducing the number of outputs for the

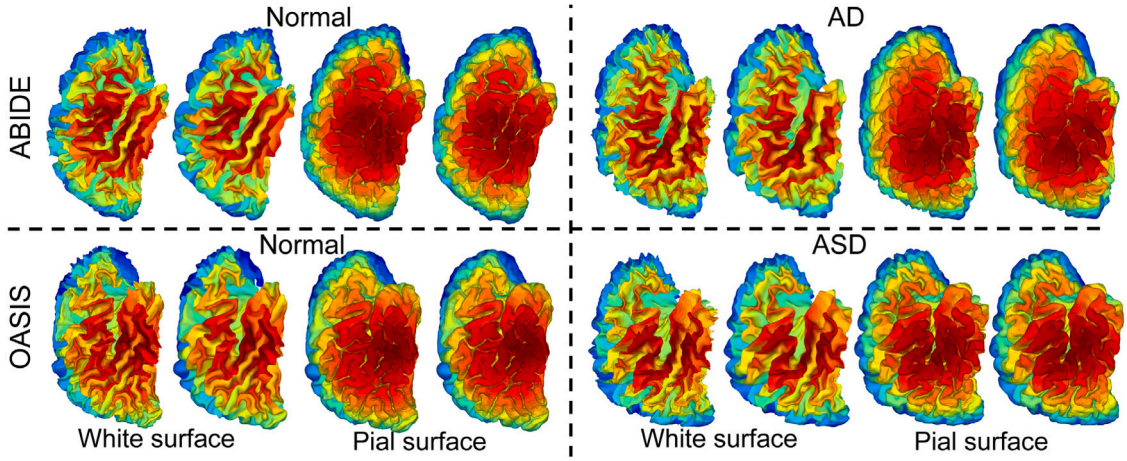

Fig. 4. Reconstruction on unseen ABIDE and OASIS datasets: Comparison of a cortical surface predicted by our joint reconstruction and segmentation method and FreeSurfer (Fischl et al., 1999). Our SegRecon method yields visually similar results while being orders of magnitude faster. Reconstruction error on the unseen ABIDE and OASIS dataset with both healthy and ASD/AD subjects are identical to the MindBoggle dataset the model is trained on, indicating the robustness of the proposed method. Only right hemisphere is shown here.

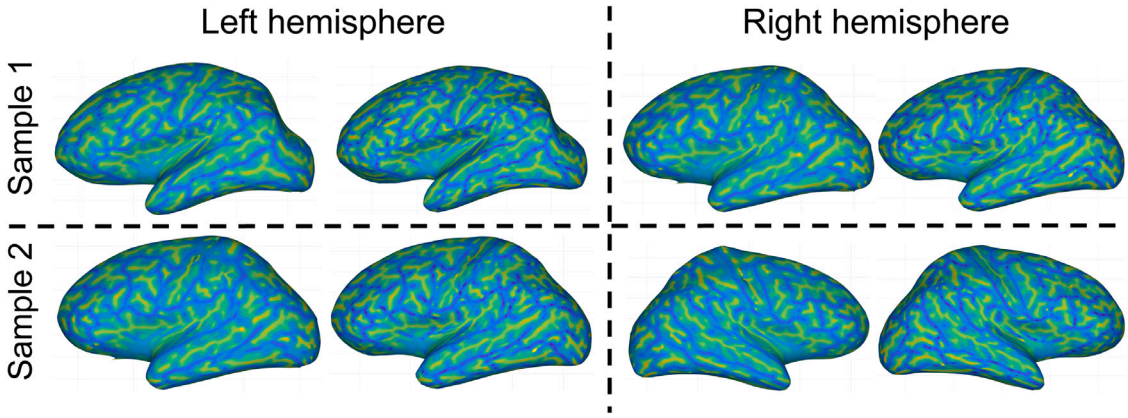

Fig. 5. Qualitative visualization of curvature: The curvature in the range [−1 mm, 1 mm] for the white surface overlaid on the inflated surface. The odd columns illustrates the curvatures overlaid on surfaces reconstructed and inflated by our method and the even columns indicate curvature for the ground-truth FreeSurfer meshes. The We can see that both methods produce a similar geometry for the two samples.

number of classes to only two, and *(ii)* providing information on the precise location of a voxel inside a parcel instead of simply measuring if a voxel is inside a parcel or not. As we will show in the next section, this continuous prediction strategy leads to a higher accuracy compared to a standard segmentation approach. However, the final predicted labels depend on the reference atlas.

For assessing the impact of the reference atlas on segmentation performance, we randomly select five subjects from the training set and use the spherical coordinates and parcellation labels of their surface mesh nodes as different atlases $\text{Ref}_1, \ldots, \text{Ref}_5$. Table 2 reports the mean Dice score obtained for test subjects using each of the five atlases. While a high accuracy is obtained in all cases, the performance also varies significantly from 84.60% to 87.33%.

To provide greater robustness to the choice of an atlas, we apply a simple multi-atlas strategy in which a separate prediction is obtained for each atlas, and individual predictions are combined using majority voting. The number of atlas or references picked for combining labels can impact the average Dice overlap of parcellation. Table 3 reports majority voting computed across multiple atlases. The subscripts indicate the number of atlases used for majority voting. The accuracy obtained by combining a different number of atlases referenced with subscripts in Table 3 shows an increase in Dice overlap from 85.29% to 89.40%. However, the computation time also increases from 1.7 s to 89.9 s, with an increase in the number of atlases referenced for majority voting. Having a trade-off between performance and computation time, Table 2 (last column) shows the majority voting strategy with 5 reference atlas obtaining a Dice score of 88.69%.

3.3. Comparison with the state-of-the-art

We now compare our joint reconstruction and parcellation method SegRecon against several baselines and recent approaches for these tasks. Table 4 reports the performance of tested methods in terms of average Dice scores over both hemispheres, mean Hausdorff distances computed on the surface manifold over both hemispheres, as well as runtimes. To evaluate the benefit of predicting cortical parcels using spherical atlas coordinates, we first train a 3D-UNet to predict the parcellation label probabilities directly at the voxel level as in standard 3D segmentation networks. This baseline, called DirectSeg in Table 4, gives a low Dice score of 79.95%. As mentioned above, this is due to the greater number of network outputs (i.e., one output per class) compared to simply predicting the two spherical atlas coordinates.

We also evaluate the FreeSurfer parcellation against the manual labels provided in the MindBoggle dataset. FreeSurfer considerably improves parcellation accuracy compared to DirectSeg with a Dice score of 84.39%. However, this comes at the price of a significant increase in computation times, from 300 ms per volume for DirectSeg to a few hours for FreeSurfer.

Table 2
Effect of reference atlas.

Ref_1	Ref_2	Ref_3	Ref_4	Ref_5	Voting
84.60(1.90)	85.85(1.79)	85.29(1.93)	85.08(1.54)	87.33(1.90)	**88.69(1.84)**

Column 1–5: The average Dice overlap (in %) obtained after using five different references as an atlas for label propagation. The last column shows the results when we vote across five different atlas references. The standard deviation is reported inside the parenthesis.

Table 3
Effect of number of reference atlases on parcellation.

Metric	Ref_1	Ref_3	Ref_5	Ref_{10}	Ref_{35}	Ref_{70}
Dice overlap	85.29(1.76)	87.62(1.97)	88.69(1.84)	88.55(1.90)	89.33(1.99)	89.4(2.04)
Time	1.7 s	4.2 s	6.5 s	12.9 s	45.0 s	89.9 s

Column 1-6: The average Dice overlap (in %) obtained after using majority voting with 1, 3, 5, 35, 70 different references as an atlas for label propagation. The standard deviation is reported inside the parenthesis.

Table 4
Evaluation of SegRecon: Comparison against state-of-the-art approaches in terms of Dice scores (in %), Hausdorff Distances (in mm), and computational time averaged over both hemispheres.

Methods	Dice overlap (%)	Hausdorff (mm)	Time
DirectSeg	79.95 ± 2.58	–	300 ms
FreeSurfer	84.39 ± 1.91	2.11 ± 0.29	4 h
FS + SRF	79.89 ± 2.62	1.97 ± 0.40	2 h + 18 s
FS + DBPN	84.60 ± 3.40	–	2 h + 1 s
FS + GCN	86.61 ± 2.45	1.66 ± 0.44	2 h + 3 s
FS + ASEGAT	89.00 ± 2.29	–	2 h + 1 s
FS + SPHARM-Net	88.48 ± 1.88	1.36 ± 0.27	2 h + 1 s
Recon+Seg	44.50 ± 0	–	4 s
w/o hemisphere	59.28 ± 12.20	3.94 ± 3.14	8 s
SegRecon (Ours)	**88.69 ± 1.84**	**1.20 ± 1.36**	8 s

The first row shows the performance of a DirectSeg a 3D-CNN network on surface parcellation. The second row illustrates the results of the traditional FreeSurfer algorithm for parcellation. In the third and fourth row, we show the ability of a Spectral Random Forest (SRF), DBPN, graph convolutional network (GCN) and ASEGAT learning based approach to segment the cortical surface. The eighth row shows the importance of learning hemisphere segmentation in our work. Finally, in the last row, we show the performance of our proposed model. The Hausdorff distance is computed directly on the surface mesh.

Third, we show the advantage of predicting cortical surfaces directly from 3D images, as in our method, compared to working with surface meshes computed previously. Toward this goal, we compare our method with four mesh-based surface parcellation methods, named FS + SRF, FS + DBPN (Zhang and Wang, 2019), FS + GCN (Gopinath et al., 2018), and FS + ASEGAT (Li et al., 2022). The results reported for these methods are taken from their original manuscripts since their code is not publicly available. The first one, Spectral Random Forest (SRF) (Lombaert et al., 2015), performs a spectral embedding of nodes in the FreeSurfer mesh graph using the main eigen-components of its Laplace matrix. The labels of embedded nodes are then predicted separately using a Random Forest classifier. In the second method, a two-stage spatial graph convolutional network explores features both locally and globally to learn the brain surface data. In the next spectral method, the connectivity of nodes in the mesh graph is exploited in the prediction using a spectral graph convolutional network (GCN) (Gopinath et al., 2019). Next, the surface parcellation using an anatomically constrained squeeze-and-excitation self attention graph neural network (Li et al., 2022) is evaluated. Finally, we evaluate a spherical harmonics based network for cortical parcellation, SPHARM-Net (Ha and Lyu, 2022b), in our training setting. As shown, predicting labels for all nodes simultaneously in FS + DBPN, FS + GCN, FS + ASEGAT, and FS + SPHARM-Net, instead of individually in FS + RF, largely improves the Dice score by a minimum of 4.71% to a maximum of 9.91%. However, these approaches require generating surface meshes in a former step, which can take around 2 h for FreeSurfer, their total run time remains substantial. In comparison, our joint reconstruction and parcellation method achieves a mean Dice score of 88.69% with an average total run time of only 8 s per volume. That is a 4.30% improvement over the Dice score of FreeSurfer, at a fraction of its computational cost.

Next, we evaluate the performance of our SegRecon method in two different settings. First, we show the importance of the hemisphere prediction loss of Eq. (5) on performance. To do so, we reduce the weight of the loss term ℓ_{hemi} in Eq. (2) to $\lambda = 0.0001$ during training. This ablation baseline is denoted as *w/o hemisphere* in Table 4. As observed, the lack of accurate hemisphere prediction results in ambiguous label prediction for surface voxels in both hemispheres which leads to a low Dice score of 59.28%. Finally, we present the setting of our model for predicting a distance $\hat{d}_v$ for each voxel compared with the hemisphere prediction. In this way, our model predicts the iso-surface for surface reconstruction. The accurate prediction of polar and azimuth angles ($\hat{\phi}_v$ and $\hat{\gamma}_v$) for obtaining parcel labels from the atlas yields an average Dice score of 88.69%. Similar improvements of our method compared to other approaches are also found for the Hausdorff distance metric. Qualitative results obtained by our surface segmentation method are shown in Fig. 6, where we illustrate the differences between the predicted and manual label boundaries for four different parcels or regions.

4. Discussion and conclusion

We presented SegRecon, a novel deep learning end-to-end model for the joint reconstruction and segmentation of nested surfaces, directly from MRI volumes. Our model learns multiple signed distance functions that represent surfaces implicitly as iso-levels. An inter-surface distance loss, computed from the distance maps during training, ensures that surfaces do not cross and that the predicted cortical thickness is anatomically possible. After applying a topological correction method (Bazin and Pham, 2007), a mesh is generated for each surface from their signed distance map using the (Lewiner et al., 2003) implementation of Marching Cubes algorithm (Lorensen and Cline, 1987). Jointly, the model also learns to predict the spherical coordinates of each voxel in a registered atlas space. The propagation of labels from the atlas space effectively segments the cortical white matter surface.

Our experiments used the largest publicly available dataset of manually-labeled brain surfaces (Klein et al., 2017), as well as the ABIDE-I (Di Martino et al., 2014) and OASIS (Marcus et al., 2007) datasets, to evaluate the surface reconstruction and segmentation accuracy of our method. We first showed the advantage of employing a signed distance map over a binary surface mask for reconstructing cortical surfaces. When comparing surfaces reconstructed by our method to those produced by FreeSurfer, using a continuous signed distance map significantly reduces the Hausdorff distance from 2.8 mm to 1 mm. Fig. 3 shows the irregularities and artifacts in the reconstructed surface

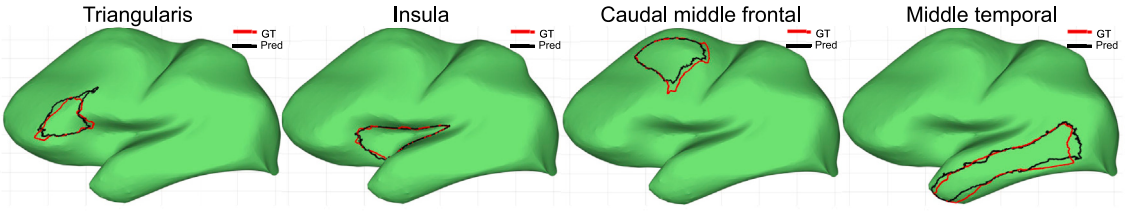

Fig. 6. Parcellation performance: The manual parcellation boundaries are shown in red, with our predicted parcellation boundaries in black. Our model segments 32 parcels in total on the brain surface. We show four parcels, namely, Triangularis, Insula, Caudal middle frontal and middle temporal of the left hemisphere for qualitative analysis. The cortical mesh is inflated here for visualization.

due to the use of binary map. We then validated the robustness of our reconstruction method on the ABIDE and OASIS datasets which were not used in training. The method yields a Hausdorff distance of less than 1.5 mm on samples from these datasets, obtained with varying acquisition protocols and corresponding to subjects with very different ages and cortical alterations. The surfaces reconstructed with our method, presented in Fig. 3, 4, are visually similar to FreeSurfer meshes which require few hours of runtime to generate. Similarly, the curvature of the white surfaces, as seen in Fig. 5, highlights the quality of our surface reconstructions.

We analyzed the impact on performance of the reference atlas selected for transferring cortical parcellation labels to the surface. While Dice scores ranging from 84.60% to 87.33% were obtained with 5 different atlases, an improved Dice of 88.69% was achieved via a multi-atlas strategy combining the predictions for different atlases with majority voting. We also compared our method against several baselines and state-of-the-art approaches for cortical parcellation. Our approach has higher Dice score than directly predicting cortical labels with 3D-UNet (79.9%) which, unlike our method, cannot be used to reconstruct cortical surfaces. Moreover, it achieved a significantly higher mean Dice score than FreeSurfer (84.3%) with substantially reduced computation times over compared to this method (hours vs. seconds). Likewise, it improved by over 2% Dice a state-of-art parcellation method based on GCN that requires pre-computed surfaces as input.

While the potential of our method has been demonstrated in our results, our method has also limitations in both reconstruction and parcellation, identified next.

Reconstruction: The reconstruction of surfaces from explicit distance functions creates a dependency on the input resolution of the MRI. Our results have evaluated the robustness of our method on a $1mm^3$ MRI volume. Our approach could however potentially fail when a surface reconstruction is made on coarser MRI resolutions. Furthermore, our proposed surface reconstruction has only been evaluated on structural T1 MRI. While an extension to other modalities and resolutions should follow the same methodology with a re-training, fine-tuning and additional data augmentation, its validation remains to be performed. On an additional note, our method adds a penalty term on thickness during training to ensure non-overlapping white and pial surfaces implicitly, while also using an external topological correction technique. To this matter, despite the added dependency on these external techniques, our results indicates that the reconstructed cortical surfaces have no measured and observed topological errors. Assessing alternative topological correction methods remains a potential for future research. As an extension of our work, we further anticipate that adding constraints to the CNN network that guarantee a topological correctness could potentially improve the accuracy of the surface reconstruction. Likewise, the presence of abnormalities such as tumors and lesions in the brain could hinder the estimation of the surface distance functions and, hence, their reconstructions.

Segmentation/Parcellation: Our approach consists of performing a cortical surface parcellation by predicting the spherical coordinates in a registered atlas space. Predicting such spherical coordinates also enables the use of different atlases at inference time with no retraining. As shown in Table 2, the results demonstrate a sensitivity to the choice of the reference atlas picked for label propagation. To overcome this limitation, we have proposed to aggregate the labels from multiple reference atlases in a majority voting scheme. The computation time for our parcellation is, therefore, dependent on the number of references picked for a label look-up. Furthermore, our method also assumes a necessary registration between the reference atlases. Our majority voting could therefore potentially fail if these reference atlases are not adequately aligned to the same template space.

Future work: Our method has currently been evaluated on cerebral cortices, while it could also be in practice applied to various other surface data, such as cardio-vascular surfaces. Moreover, although our model includes a loss to control the distance between reconstructed surfaces and prevent them from crossing one another, incorporating more powerful topological constraints during training could possibly remove the need for our current use of external topological correction.

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 data we used for this manuscript is publically available.

Acknowledgments

This work was supported financially by the Canada Research Chair on Shape Analysis in Medical Imaging, the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Fonds de recherche du Québec – Nature et technologies (FRQNT), Canada. We also gratefully acknowledge the support of the NVIDIA Corporation with the donation of a GPU used for this research.

References

Bazin, P.-L., Pham, D.L., 2007. Topology correction of segmented medical images using a fast marching algorithm. Comput. Methods Programs Biomed.

Billot, B., Magdamo, C., Cheng, Y., Arnold, S.E., Das, S., Iglesias, J.E., 2023. Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets. Proc. Natl. Acad. Sci.

Bongratz, F., Rickmann, A.-M., Pölsterl, S., Wachinger, C., 2022. Vox2Cortex: Fast explicit reconstruction of cortical surfaces from 3D MRI scans with geometric deep neural networks. In: Proceedings of the Institute of Electrical and Electronics Engineers /CVF Conference on Computer Vision and Pattern Recognition.

Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention.

Cruz, R.S., Lebrat, L., Bourgeat, P., Fookes, C., Fripp, J., Salvado, O., 2021. DeepCSR: A 3D deep learning approach for cortical surface reconstruction. In: Proceedings of the Institute of Electrical and Electronics Engineers /CVF Winter Conference on Applications of Computer Vision.

Dahnke, R., Yotter, R.A., Gaser, C., 2013. Cortical thickness and central surface estimation. Neuroimage.

Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., et al., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. Neuroimage.

Di Martino, A., Yan, C.-G., Li, Q., Denio, E., Castellanos, F.X., Alaerts, K., Anderson, J.S., Assaf, M., Bookheimer, S.Y., Dapretto, M., et al., 2014. The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. Mol. Psychiatry.

Fischl, B., Dale, A.M., 2000. Measuring the thickness of the human cerebral cortex from magnetic resonance images. Proc. Natl. Acad. Sci.

Fischl, B., Sereno, M.I., Dale, A.M., 1999. Cortical surface-based analysis: II. inflation, flattening, and a surface-based coordinate system. Neuroimage.

Fischl, B., et al., 2004. Automatically parcellating the cortex. Cereb. Cortex.

Germanaud, D., Lefèvre, J., Toro, R., Fischer, C., Dubois, J., Hertz-Pannier, L., Mangin, J.F., 2012. Larger is twistier: Spectral analysis of gyrification (SPANGY) applied to adult brain size polymorphism. NeuroImage 63.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., et al., 2016. A multi-modal parcellation of human cerebral cortex. Nature.

Gopinath, K., Desrosiers, C., Lombaert, H., 2018. Graph convolutions on spectral embeddings: Learning of cortical surface data. ArXiv Preprint 1803.10336. arXiv:1803.10336.

Gopinath, K., Desrosiers, C., Lombaert, H., 2019. Graph convolutions on spectral embeddings for cortical surface parcellation. Med. Image Anal.

Gopinath, K., Desrosiers, C., Lombaert, H., 2020. Graph domain adaptation for alignment-invariant brain surface segmentation. In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Graphs in Biomedical Image Analysis. Springer.

Gopinath, K., Desrosiers, C., Lombaert, H., 2021. SegRecon: Learning joint brain surface reconstruction and segmentation from images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer.

Ha, S., Lyu, I., 2022a. SPHARM-Net: Spherical harmonics-based convolution for cortical parcellation. Inst. Electr. Electron. Eng. Trans. Med. Imag. 41 (10), 2739–2751. http://dx.doi.org/10.1109/TMI.2022.3168670.

Ha, S., Lyu, I., 2022b. SPHARM-Net: Spherical harmonics-based convolution for cortical parcellation. Inst. Electr. Electron. Eng. Trans. Med. Imag.

He, R., Gopinath, K., Desrosiers, C., Lombaert, H., 2020. Spectral graph transformer networks for brain surface parcellation. In: International Symposium on Biomedical Imaging.

Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M., 2020. FastSurfer - A fast and accurate deep learning based neuroimaging pipeline. NeuroImage.

Hoopes, A., Iglesias, J.E., Fischl, B., Greve, D., Dalca, A.V., 2022. TopoFit: Rapid reconstruction of topologically-correct cortical surfaces. In: Proceedings of Machine Learning Research. NIH Public Access.

Kim, J.S., Singh, V., Lee, J.K., Lerch, J., Ad-Dab'bagh, Y., MacDonald, D., Lee, J.M., Kim, S.I., Evans, A.C., 2005. Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a Laplacian map and partial volume effect classification. Neuroimage.

Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. In: International Conference on Learning Representations.

Klein, A., Ghosh, S.S., Bao, F.S., Giard, J., Häme, Y., Stavsky, E., Lee, N., Rossa, B., Reuter, M., Chaibub Neto, E., Keshavan, A., 2017. Mindboggling morphometry of human brains. PLoS Comput. Biol.

Klein, A., Tourville, J., 2012. 101 Labeled brain images and a consistent human cortical labeling protocol. Front. Neurosci.

Kriegeskorte, N., Goebel, R., 2001. An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical MR volumes. NeuroImage.

Lewiner, T., Lopes, H., Vieira, A.W., Tavares, G., 2003. Efficient implementation of marching cubes' cases with topological guarantees. J. Graph. Tools.

Li, X., Tan, J., Wang, P., Liu, H., Li, Z., Wang, W., 2022. Anatomically constrained squeeze-and-excitation graph attention network for cortical surface parcellation. Comput. Biol. Med.

Lombaert, H., Criminisi, A., Ayache, N., 2015. Spectral forests: Learning of surface data, application to cortical parcellation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer.

López-López, N., Vázquez, A., Poupon, C., Mangin, J.-F., Ladra, S., Guevara, P., 2020. GeoSP: A parallel method for a cortical surface parcellation based on geodesic distance. In: Engineering in Medicine & Biology Society.

Lorensen, W.E., Cline, H.E., 1987. Marching cubes: A high resolution 3D surface construction algorithm. In: ACM Siggraph Computer Graphics.

Ma, Q., Robinson, E.C., Kainz, B., Rueckert, D., Alansary, A., 2021. PialNN: A fast deep learning framework for cortical pial surface reconstruction. In: Machine Learning in Clinical Neuroimaging: 4th International Workshop, MLCN 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 4. Springer.

Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open access series of imaging studies (OASIS): Cross-sectional MRI data in Young, middle aged, nondemented, and demented older adults. J. Cogn. Neurosci. 19 (9), 1498–1507.

Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S., 2019. DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the Institute of Electrical and Electronics Engineers /CVF Conference on Computer Vision and Pattern Recognition.

Parvathaneni, P., Bao, S., Nath, V., Woodward, N.D., Claassen, D.O., Cascio, C.J., Zald, D.H., Huo, Y., Landman, B.A., Lyu, I., 2019. Cortical surface parcellation using spherical convolutional neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer.

Puonti, O., Iglesias, J.E., Van Leemput, K., 2016. Fast and sequence-adaptive whole-brain segmentation using parametric Bayesian modeling. NeuroImage.

Querbes, O., Aubry, F., Pariente, J., Lotterie, J.-A., Démonet, J.-F., Duret, V., Puel, M., Berry, I., Fort, J.-C., Celsis, P., et al., 2009. Early diagnosis of Alzheimer's disease using cortical thickness: Impact of cognitive reserve. Brain.

Shattuck, D.W., Leahy, R.M., 2002. BrainSuite: An automated cortical surface identification tool. Med. Image Anal.

Styner, M., et al., 2006. Framework for the statistical shape analysis of brain structures using SPHARM-PDM. Insight J.

Wu, Z., Zhao, F., Xia, J., Wang, L., Lin, W., Gilmore, J.H., Li, G., Shen, D., 2019. Intrinsic patch-based cortical anatomical parcellation using graph convolutional neural network on surface manifold. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer.

Zhang, W., Wang, Y., 2019. Geometric brain surface network for brain cortical parcellation. In: International Workshop on Graph Learning in Medical Imaging. Springer.

Zhao, F., Wu, Z., Wang, L., Lin, W., Gilmore, J.H., Xia, S., Shen, D., Li, G., 2021. Spherical deformable U-Net: Application to cortical surface parcellation and development prediction. Trans. Med. Imag.

Zhao, F., Xia, S., Wu, Z., Duan, D., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G., 2019. Spherical U-Net on cortical surfaces: Methods and applications. In: International Conference on Information Processing in Medical Imaging. Springer.