CorticalFlow^{++}: Boosting Cortical Surface Reconstruction Accuracy, Regularity, and Interoperability

Rodrigo Santa Cruz[† 1,2][0000−0002−5273−7296], Léo Lebrat[† 1,2], Darren Fu[3], Pierrick Bourgeat[1], Jurgen Fripp[1], Clinton Fookes[2], and Olivier Salvado[1,2]

[1] CSIRO
[2] Queensland University of Technology
[3] University of Queensland, Australia.

Abstract. The problem of Cortical Surface Reconstruction from magnetic resonance imaging has been traditionally addressed using lengthy pipelines of image processing techniques like FreeSurfer, CAT, or CIVET. These frameworks require very long runtimes deemed unfeasible for real-time applications and unpractical for large-scale studies. Recently, supervised deep learning approaches have been introduced to speed up this task cutting down the reconstruction time from hours to seconds. Using the state-of-the-art CorticalFlow model as a blueprint, this paper proposes three modifications to improve its accuracy and interoperability with existing surface analysis tools, while not sacrificing its fast inference time and low GPU memory consumption. First, we employ a more accurate ODE solver to reduce the diffeomorphic mapping approximation error. Second, we devise a routine to produce smoother template meshes avoiding mesh artifacts caused by sharp edges in CorticalFlow's convex-hull based template. Last, we recast pial surface prediction as the deformation of the predicted white surface leading to a one-to-one mapping between white and pial surface vertices. This mapping is essential to many existing surface analysis tools for cortical morphometry. We name the resulting method CorticalFlow^{++}. Using large-scale datasets, we demonstrate the proposed changes provide more geometric accuracy and surface regularity while keeping the reconstruction time and GPU memory requirements almost unchanged.

Keywords: Cortical Surface Reconstruction · CorticalFlow · 3D Deep Learning.

1 Introduction

The problem of cortical surface reconstruction (CSR) consists of estimating triangular meshes for the inner and outer cortical surfaces from a magnetic res-

---

onance image (MRI). It is a pivotal problem in Neuroimaging and a fundamental task in clinical studies of neurodegenerative diseases [4] and psychological disorders [21]. Traditionally, this problem is tackled by extensive pipelines of handcrafted image processing algorithms [6,25,13,18] which are subject to careful hyperparameter tuning (e.g., thresholds, iteration numbers, and convergence criterion) and very long runtimes.

To overcome these issues, deep learning (DL) based approaches have been proposed recently [23,9,17,15]. These methods can directly predict cortical surfaces geometrically close to those produced by traditional methods while reducing their processing time from hours to seconds. More specifically, the current state-of-the-art DL model for CSR, named CorticalFlow [15], consists of a sequence of diffeomorphic deformation modules that learn to deform a template mesh towards the target cortical surface from an input MRI. This method takes only a few seconds to produce cortical surfaces with sub-voxel accuracy.

In this paper, we propose three modifications to improve the accuracy of CorticalFlow and its interoperability to other surface processing tools without increasing its reconstruction time and GPU memory consumption. First, we upgrade the Euler method used by CorticalFlow to solve the flow ODE responsible for computing the diffeomorphic mapping to the fourth-order Runge-Kutta method [20]. This tool provides more accurate ODE solutions leading the model to better approximate the target surfaces and reducing the number of self-intersecting faces on the reconstructed meshes. Second, instead of using the convex hull of the training surfaces as a template, we propose a simple routine to generate smooth templates that tightly wrap the training surfaces. This new template eases the approximation problem leading to more accurate surfaces while suppressing mesh artifacts in highly curved regions. Finally, inspired by Ma et al. [17], we leverage the estimated white surfaces as the initial mesh template for learning and predicting the pial surfaces. This approach provides a better "starting point" to the approximation problem as well as a one-to-one mapping between white and pial surface vertices which facilitates the use of the generated surfaces in existing surface-based analysis tools [24,8]. In acknowledgment of the CorticalFlow framework, we name our resulting method CorticalFlow^{++}.

Using a large dataset of MRI images and pseudo-ground-truth surfaces, we compare CorticalFlow and CorticalFlow^{++} performance in the reconstruction of cortical surfaces from MRI on three perspectives: geometric accuracy, mesh regularity, and time and space complexity for the surface reconstruction. We conclude that the proposed CorticalFlow^{++} improves upon CorticalFlow, on average, by 19.11% in terms of Chamfer distance and 56.77% in terms of the percentage of self-intersecting faces. Additionally, it adds only half a second to the final surface reconstruction time while keeping the same GPU memory budget.

2   Related Work

Traditional cortical surface reconstruction frameworks like FreeSurfer[6], Brain-Suite [25], and CIVET [18] involve two major steps usually accomplished by lengthy sequences of image processing techniques. They first voxel-wise segment the input MRI, then fit surfaces enclosing the gray matter tissue delimiting the brain cortex. More specifically, the widely used FreeSurfer V6 framework for cortical surface analysis from MRI uses an atlas-based segmentation [7] and a deformable model [3] for surface fitting on these segmented volumes. Recently, Henschel et al. [9] accelerated this framework with a modern DL brain segmentation model and a fast spectral mesh processing algorithm for spherical mapping, cutting down FreeSurfer's processing time to one hour.

In contrast, supervised deep learning approaches leverage large datasets of MRIs and pseudo-ground-truth produced with traditional methods to train high-capacity neural networks to predict surfaces directly from the MRI. DeepCSR [23] trains an implicit surface predictor and CorticalFlow [15] learns a diffeomorphic deformable model. Similarly, PialNN [17] also learns a deformable model but focuses only on pial surface reconstruction, receiving as input the white matter surface generated with traditional methods.

Building upon the success and the powerful framework of CorticalFlow, this paper proposes to address its main limitations, aiming to improve its accuracy and interoperability with existent surface analysis tools, but without severely degrading its inference time and GPU memory consumption.

3   Method

In this Section, we present the proposed method CorticalFlow^{++}. We start by reviewing the original CorticalFlow framework and its main components, then we introduce our proposed changes.

3.1   CorticalFlow Framework

Originally proposed by Lebrat et al. [15], CorticalFlow consists of a chain of deformation blocks that receives as input a 3-dimensional MRI $\mathbf{I} \in \mathbb{R}^{H \times W \times D}$ and deforms an initial template mesh $\mathcal{T}$ producing the desired cortical surface mesh. As shown in Figure 1a, a CorticalFlow's deformation block comprises a U-Net [22] and a Diffeomorphic Mesh Deformation module (DMD) [15]. The U-Net in the i-th deformation block, denoted as $\mathrm{UNet}_{\theta_i}^i$ and parametrized by θ_i, outputs a stationary flow field $\mathbf{U}_i \in \mathbb{R}^{H \times W \times D \times 3}$, while it receives as input the channel-wise concatenation of the input MRI and all the previous blocks' U-Nets outputs $\{\mathbf{U}_1, \ldots, \mathbf{U}_{i-1}\}$. The predicted flow field encodes how the template mesh should be deformed to approximate the target cortical surface. The DMD module receives as input the block's UNet predicted flow field $\mathbf{U}_i$ and computes a diffeomorphic mapping $\Phi : [0,1] \times \mathbb{R}^3 \to \mathbb{R}^3$ for each vertex position $\mathbf{x} \in \mathbb{R}^3$ in
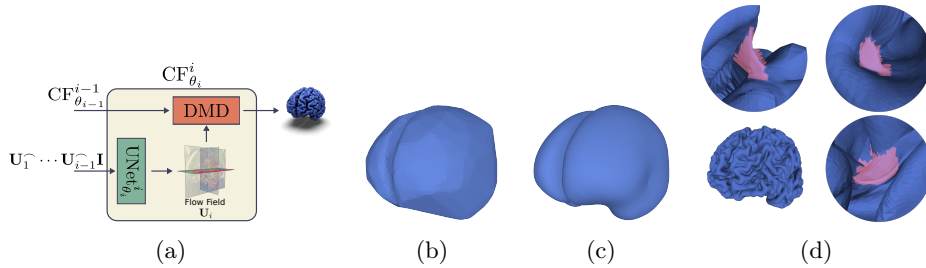
Fig. 1: (a) CorticalFlow's Deformation block. (b) Convex-hull based Corti-calFlow's mesh template. (c) CorticalFlow^{++}'s proposed template. (d) Examples of mesh artifacts caused by sharp edges in CorticalFlow's template.

the resulting mesh computed by the previous deformation block. Formally, the DMD module solves the flow ODE,

$$\frac{d\Phi(s; \mathbf{x})}{ds} = \mathbf{U}\left(\Phi(s; \mathbf{x})\right), \text{ with } \Phi(0; \mathbf{x}) = \mathbf{x}, \tag{1}$$

using the forward Euler method [5]. This flow ODE formulation preserves the topology of the initial template mesh without producing self-intersecting faces. Hence, CorticalFlow with k deformations (CF^k) can be written using the following recurrence,

$$\mathrm{CF}^1_{\theta_1}(\mathbf{I}, \mathcal{T}) = \mathrm{DMD}(\mathrm{UNet}^1_{\theta_1}(\mathbf{I}), \mathcal{T}))$$

$$\mathrm{CF}^{i+1}_{\theta_{i+1}}(\mathbf{I}, \mathcal{T}) = \mathrm{DMD}(\mathrm{UNet}^{i+1}_{\theta_{i+1}}(\mathbf{U}_1 \frown \cdots \mathbf{U}_i \frown \mathbf{I}), \mathrm{CF}^i_{\theta_i}(\mathbf{I}, \mathcal{T})) \quad \text{for } i \geq 1, \tag{2}$$

with $\mathbf{A} \frown \mathbf{B}$ denoting the channel-wise concatenation of the tensors $\mathbf{A}$ and $\mathbf{B}$. As in [15], we focus on CorticalFlow with three deformation blocks (CF^3).

In order to train CF^3, the authors propose a multi-stage approach where the deformation blocks are trained successively keeping previous blocks' weights frozen and using template meshes of increasing resolution $\mathcal{T}_i$. Mathematically, at each stage i, they optimize the following objective,

$$\arg\min_{\theta_i} \sum_{(\mathbf{I},\mathbf{S})\in\mathcal{D}} \mathcal{L}\left(\mathrm{CF}^i_{\theta_i}(\mathbf{I}, \mathcal{T}_i), S\right), \tag{3}$$

where $\mathcal{D}$ is a dataset composed of pairs of MRIs and their respective triangle meshes S representing some cortical surface. $\mathcal{L}(\cdot, \cdot)$ is the training loss composed of mesh edge loss [27] and Chamfer Distance loss [27]. Note that for each cortical surface and brain hemisphere, a separate CorticalFlow is trained independently.

3.2   Higher Order ODE Solver

DMD modules in CorticalFlow solve the flow ODE (1) using the forward Euler method which consists of an iterative method defined by the following integration step:

$$\hat{\Phi}(h, \mathbf{x}) = \mathbf{x} + h\mathbf{U}(\mathbf{x}), \tag{4}$$

where h is the algorithm step-size and $\mathbf{U}(\mathbf{x})$ is the linear interpolation of a predicted flow field $\mathbf{U}$ at mesh vertex position $\mathbf{x}$. $\hat{\Phi}$ is a numerical approximation of the true diffeomorphic mapping Φ due to the interpolation and discretization errors inherent to this application.

Since the Euler method only provides an approximation of the continuous problem with an error that decreases linearly as h decreases, it may require a large number of integration steps to approximate accurately the continuous solution. Otherwise, the resulting mapping may cease to be invertible. For these reasons, we propose to use the Runge-Kutta [20] explicit fourth-order approximation method also known as $RK4$. The integration step of this method consists of the weighted average of four slopes estimated at the beginning, two different midpoints, and end of the step size interval. For our stationary vector field, the $RK4$ integration step is defined as,

$$\hat{\Phi}(h, \mathbf{x}) = \mathbf{x} + \frac{1}{6}\left[k_1 + 2k_2 + 2k_3 + k_4\right], \tag{5}$$

where $k_1 = \mathbf{U}\left(\mathbf{x}\right)$, $k_2 = \mathbf{U}\left(\mathbf{x} + h\frac{k_1}{2}\right)$, $k_3 = \mathbf{U}\left(\mathbf{x} + h\frac{k_2}{2}\right)$, and $k_4 = \mathbf{U}\left(\mathbf{x} + hk_3\right)$ are the averaged slopes.

3.3   Smooth Templates

CorticalFlow's template mesh consists of the convex-hull of all cortical surface meshes in the training set. This approach has two main shortcomings:

1. Even target surfaces with small differences between them can lead to a "loose" template. Consequently, the model has to learn "large" deformations making the smooth approximation problem harder. For pial surfaces, this problem is even worse because some template regions may lay outside the image bounds where the predicted flow field is undefined.
2. The convex-hull is defined by a set of intersecting planes which leads to sharp edges as shown in Figure 1b. These edges are very hard to be unfolded by a smooth deformable model like CorticalFlow, because it requires non-smooth deformations with drastic local changes of direction in its flow field representation. Hence, these edges may remain in the predicted mesh as undesirable artifacts (see Figure 1d).

To overcome these issues, we develop genus zero smooth mesh templates that tightly wrap all training meshes. We first compute a signed distance map for every target surface mesh in the training set by computing the largest 3D bounding box that contains these meshes, create 512^3 voxel-grids into this bounding box, and populate these voxel-grids with the signed distance to each target mesh. These signed distance maps are implicit representations of the target meshes where voxels with positives values are outside of the mesh and voxels with negative values are inside of the mesh. Then, by thresholding the binary union of these maps and running the standard marching cubes algorithm [16], we obtain a template mesh that is very tight around all training surfaces. However, this template mesh looks "blocky" with many small sharp edges and undesired topological

defects. The template mesh is thus smoothed using the Laplacian smoothing algorithm [10,26] and re-meshed with Delaunay triangulation [1]. The result is a smooth template mesh with spherical topology tightly wrapping all training set surfaces (see Figure 1c). Finally, we apply a topology preserving mesh subdivision algorithm [2] to generate template meshes at different resolutions which are required to train CorticalFlow. Implementations of the used algorithms are available in the Libigl [12] and MeshLab [2] toolboxes.

3.4   White To Pial Surface Morphing

In Lebrat et al. [15], a separate CorticalFlow model is trained for each cortical surface (i.e., pial and white) and each brain hemisphere (i.e., left and right). This approach leads to reconstructed surface meshes without a one-to-one mapping between the vertices in the white and pial surfaces on the same brain hemisphere. In the absence of this mapping, many existent surface analysis tools can not process the generated surfaces. Additionally, as also observed in Ma et al. [17], the pial surface may only differ from the white surface by a "small" deformation thanks to the natural anatomical agreement between these surfaces. Therefore, we propose to predict pial surfaces by learning to deform the predicted white surfaces instead of using a pial surface template mesh.

   Formally, the resulting model for pial surfaces with k deformation blocks (CFP^k) can be restated by the following recurrence,

$$\text{CFP}^1_{\theta_1}(\mathbf{I}, \mathcal{T}^w) = \text{DMD}(\text{UNet}^1_{\theta_1}(\mathbf{I}), \text{CFW}^{k'}(\mathbf{I}, \mathcal{T}^w))$$
$$\text{CFP}^{i+1}_{\theta_{i+1}}(\mathbf{I}, \mathcal{T}^w) = \text{DMD}(\text{UNet}^{i+1}_{\theta_{i+1}}(\mathbf{U}_1^\frown \cdots \mathbf{U}_i^\frown \mathbf{I}), \text{CFP}_i(\mathbf{I}, \mathcal{T}^w)) \quad \text{for } i \geq 1, \quad (6)$$

where $\text{CFW}^{k'}$ is a CorticalFlow model with k' deformation blocks pretrained to reconstruct the same hemisphere white surface as described in Section 3.1 and $\mathcal{T}^w$ is its respective template mesh for white surfaces generated as described in Section 3.3. Note that the formulation of $\text{CFW}^{k'}$ remains the one stated in Equation 2 and we only use template meshes for the white surfaces since the pial surfaces are obtained by deforming the predicted white surface.

4   Experiments

We now evaluate CorticalFlow^{++} in the cortical surface reconstruction problem. First, using the CSR benchmark introduced by Santa Cruz et al. [23], we quantify the performance impact of each proposed modification separately. This benchmark consists of 3,876 MRI images extracted from the Alzheimer's Disease Neuroimaging Initiative[4] (ADNI)[11] and their respective pseudo-ground-truth

---

[4] Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: `http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf`

	Left Pial Surface				Right Pial Surface			
	CH(mm)↓	HD(mm)↓	CHN↑	% SIF↓	CH(mm)↓	HD(mm)↓	CHN↑	% SIF↓
CorticalFlow	0.681	0.802	0.932	0.686	0.693	0.815	0.929	1.239
3.22 sec / 2.82 GB	(±0.098)	(±0.049)	(±0.006)	(±0.469)	(±0.091)	(±0.046)	(±0.006)	(±0.629)
CorticalFlow + RK4	0.629	0.761	0.937	0.502	0.580	0.751	0.943	0.280
3.55 sec / 2.82 GB	(±0.100)	(±0.042)	(±0.006)	(±0.196)	(±0.082)	(±0.038)	(±0.006)	(±0.133)
CorticalFlow + W2P	0.545	0.730	0.943	0.188	0.540	0.729	0.945	0.176
3.63 sec / 2.82 GB	(±0.082)	(±0.037)	(±0.006)	(±0.116)	(±0.075)	(±0.033)	(±0.006)	(±0.134)
PialNN†	5.500	2.793	0.792	4.730	5.948	2.898	0.789	4.537
white gen. + 0.880 secs / 1.92 GB	(±0.786)	(±0.220)	(±0.009)	(±0.841)	(±0.811)	(±0.212)	(±0.011)	(±0.815)
PialNN*	1.388	1.251	0.864	10.507	1.374	1.236	0.863	11.159
white gen. + 0.880 secs / 1.92 GB	(±0.223)	(±0.120)	(±0.011)	(±1.908)	(±0.217)	(±0.115)	(±0.011)	(±1.930)
CorticalFlow^{++}	0.529	0.721	0.946	0.069	0.528	0.723	0.946	0.099
3.76 sec / 2.82 GB	(±0.088)	(±0.036)	(±0.006)	(±0.060)	(±0.074)	(±0.031)	(±0.005)	(±0.093)
	Left White Surface				Right White Surface			
	CH(mm)↓	HD(mm)↓	CHN↑	% SIF↓	CH(mm)↓	HD(mm)↓	CHN↑	% SIF↓
CorticalFlow	0.608	0.785	0.941	0.033	0.599	0.783	0.942	0.030
3.22 sec / 2.82 GB	(±0.098)	(±0.060)	(±0.007)	(±0.030)	(±0.093)	(±0.059)	(±0.007)	(±0.029)
CorticalFlow + RK4	0.540	0.733	0.948	0.042	0.517	0.716	0.951	0.010
3.55 sec / 2.82 GB	(±0.107)	(±0.055)	(±0.006)	(±0.039)	(±0.089)	(±0.044)	(±0.006)	(±0.023)
CorticalFlow + NEWTPL	0.598	0.780	0.942	0.030	0.558	0.747	0.945	0.104
3.04 sec / 2.82 GB	(±0.101)	(±0.062)	(±0.007)	(±0.027)	(±0.091)	(±0.053)	(±0.006)	(±0.079)
CorticalFlow^{++}	0.514	0.712	0.952	0.017	0.510	0.711	0.952	0.031
3.35 sec / 2.82 GB	(±0.090)	(±0.044)	(±0.006)	(±0.023)	(±0.083)	(±0.040)	(±0.006)	(±0.040)

Table 1: Cortical Surface Reconstruction Benchmark [23]. ↓ indicates smaller metric value is better, while ↑ indicates greater metric value is better.

surfaces generated with the FreeSurfer V6.0 cross-sectional pipeline. We strictly follow the benchmark's data splits and evaluation protocol. We refer the reader to [23,15] for full details on this dataset. As evaluation metrics for measuring geometric accuracy, we use Chamfer distance (CH), Hausdorff distance (HD), and Chamfer normals (CHN). These distances are computed for point clouds of 200k points uniformly sampled from the predicted and target surfaces. As a measure of surface regularity, we compute the percentage of self-intersecting faces (%SIF) using PyMeshLab [19]. Finally, we also report the average time (in seconds) and the maximum GPU memory (in GB) required by the evaluated methods to reconstruct a cortical surface. Table 1 presents these results.

We observe that the adoption of the $RK4$ ODE solver (CorticalFlow + RK4) and the white to pial surface morphing formulation (CorticalFlow + W2P) significantly improves the geometric accuracy and surface regularity of the CorticalFlow baseline. For instance, we noticed an average decrease of 12.2% and 21.02% in chamfer distance, respectively. Likewise, the percentage of self-intersecting faces reduces on average by 35.90% and 79.19%, respectively. On the other hand, the proposed new templates (CorticalFlow + NEWTPL) present a modest improvement in those criteria, but greatly succeeds in suppressing mesh artifacts as qualitatively shown in Figure 2. Additionally, none of these changes incur a significant increase in reconstruction time or memory consumption due to their GPU friendly nature. All together, these modifications in CorticalFlow^{++} establish a new state-of-the-art method for cortical surface reconstruction.

Since the white to pial morphing approach has been previously introduced in [17], we compare CorticalFlow^{++} and PialNN [17] on the pial surface re-
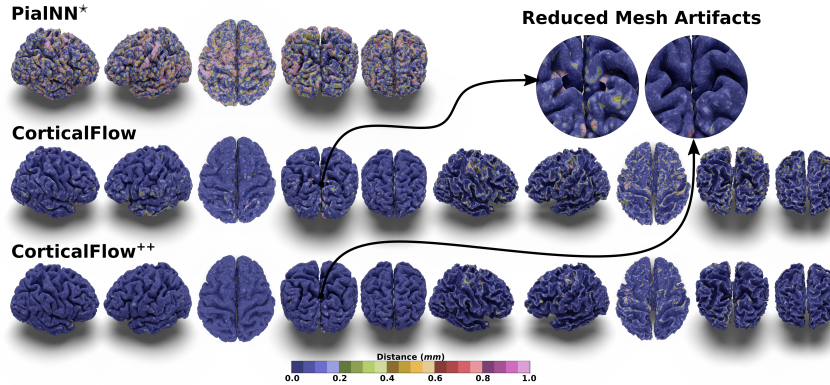
Fig. 2: Predicted cortical surfaces color-coded with the distance to the pseudo-ground-truth surfaces. See our supplementary materials for more examples.

	Left Pial Surface				Right Pial Surface			
	CH$(mm)\downarrow$	HD$(mm)\downarrow$	CHN $\uparrow$	% SIF $\downarrow$	CH$(mm)\downarrow$	HD$(mm)\downarrow$	CHN $\uparrow$	% SIF $\downarrow$
CorticalFlow	0.677	0.803	0.923	0.594	0.724	0.845	0.918	1.467
3.22 sec / 2.82 GB	($\pm$0.099)	($\pm$0.056)	($\pm$0.007)	($\pm$0.319)	($\pm$0.106)	($\pm$0.063)	($\pm$0.007)	($\pm$0.519)
PialNN†	5.426	2.763	0.781	4.281	5.944	2.873	0.777	4.033
white gen. + 0.880 secs / 1.92 GB	($\pm$0.486)	($\pm$0.144)	($\pm$0.009)	($\pm$0.709)	($\pm$0.503)	($\pm$0.139)	($\pm$0.009)	($\pm$0.671)
PialNN*	1.307	1.243	0.860	9.661	1.264	1.211	0.857	10.482
white gen. + 0.880 secs / 1.92 GB	($\pm$0.202)	($\pm$0.119)	($\pm$0.013)	($\pm$1.604)	($\pm$0.194)	($\pm$0.117)	($\pm$0.013)	($\pm$1.678)
CorticalFlow^{++}	0.520	0.711	0.935	0.136	0.528	0.727	0.935	0.245
3.76 sec / 2.82 GB	($\pm$0.082)	($\pm$0.044)	($\pm$0.006)	($\pm$0.096)	($\pm$0.079)	($\pm$0.047)	($\pm$0.006)	($\pm$0.167)

Table 2: Out-of-train-distribution evaluation on OASIS3 [14]. $\downarrow$ indicates smaller metric value is better, while $\uparrow$ indicates greater metric value is better.

construction. For this comparison, we report the performance of the publicly available pretrained PialNN model[5] (PialNN†) as well as training it by ourselves in the CSR benchmark training dataset (PialNN*). CorticalFlow^{++} compares favourably to both PialNN variants and importantly, it does not need traditional methods to generate white surfaces.

Finally, we extend our evaluation to an out-of-training-distribution dataset. More specifically, we use the trained models from the previous experiment to reconstruct cortical surfaces for a subset of MRIs extracted from the OASIS3 dataset [14]. These generated surfaces are also compared to FreeSurfer V6.0 pseudo-ground-truth surfaces using the same evaluation metrics described above. As shown in Table 2, CorticalFlow^{++} significantly outperforms CorticalFlow and PialNN, while presenting comparable surface reconstruction runtime and GPU memory consumption.

---

[5] https://github.com/m-qiang/PialNN

5   Conclusion

This paper tackles some limitations of CorticalFlow, the current state-of-the-art model for Cortical surface reconstruction from MRI, in order to improve its accuracy, regularity, and interoperability without sacrificing its computational requirements for inference (reconstruction time and maximum GPU memory consumption). The resulting method, CorticalFlow^{++}, achieves state-of-the-art performance on geometric accuracy and surface regularity while keeping the GPU memory consumption constant and adding less than a second to the entire surface reconstruction process.

6   Compliance with Ethical Standards

This research was approved by CSIRO ethics 2020 068 LR.

7   Acknowledgements
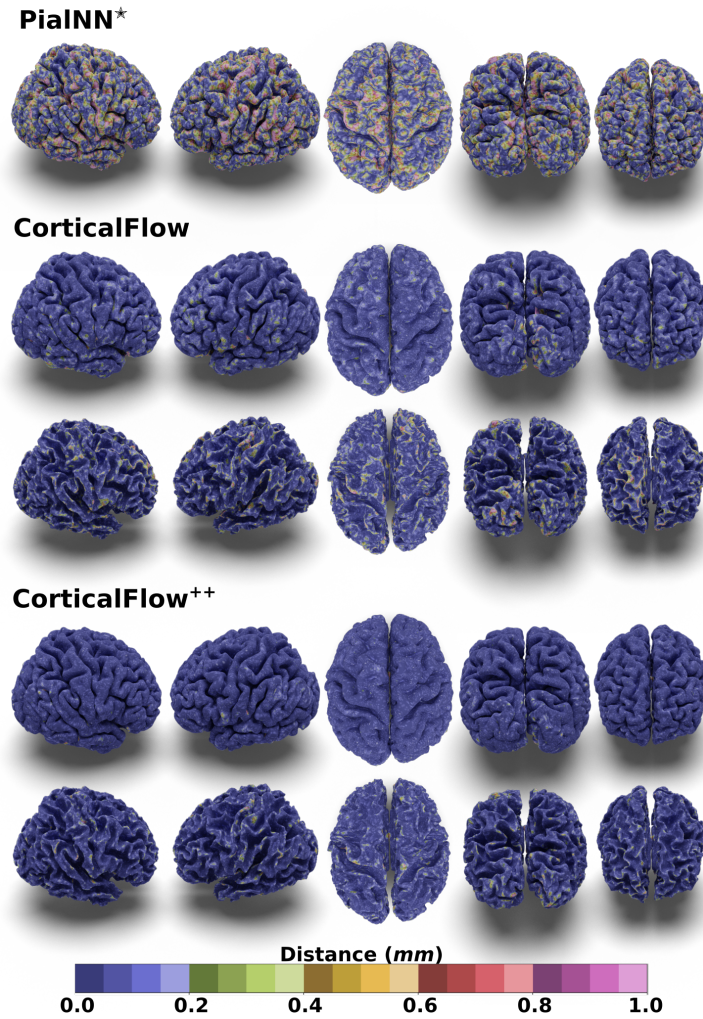
Supplementary Material

Fig. 3: Predicted cortical surfaces for subject **099_S_0551_m36** in **ADNI** dataset. The surfaces are color-coded with the distance to the pseudo-ground-truth surfaces.
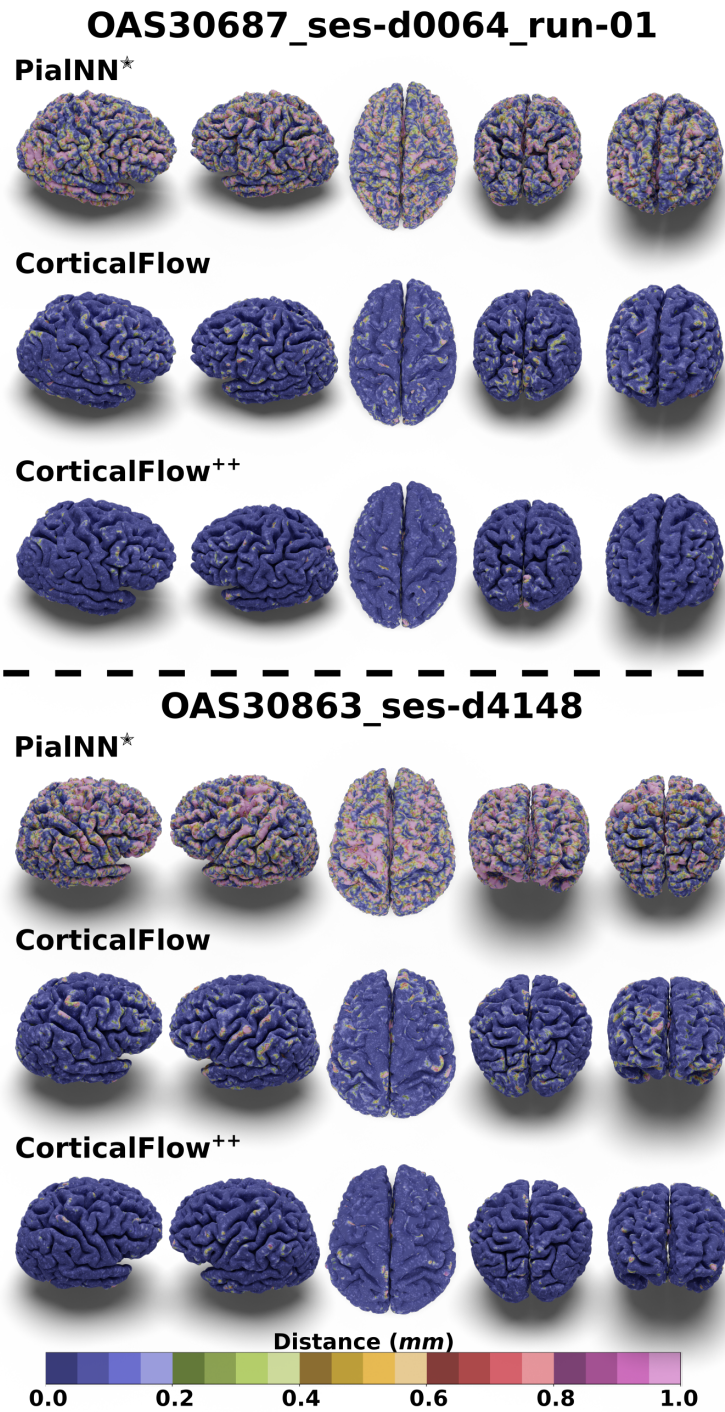
Fig. 4: Predicted cortical surfaces for **OASIS3** subjects. The surfaces are color-coded with the distance to the pseudo-ground-truth surfaces.

References

1. Bobenko, A.I., Springborn, B.A.: A discrete laplace–beltrami operator for simplicial surfaces. Discrete & Computational Geometry **38**(4), 740–756 (2007)
2. Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., Ranzuglia, G.: MeshLab: an Open-Source Mesh Processing Tool. In: Scarano, V., Chiara, R.D., Erra, U. (eds.) Eurographics Italian Chapter Conference. The Eurographics Association (2008). `https://doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136`
3. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. segmentation and surface reconstruction. Neuroimage **9**(2), 179–194 (1999)
4. Du, A.T., Schuff, N., Kramer, J.H., Rosen, H.J., Gorno-Tempini, M.L., Rankin, K., Miller, B.L., Weiner, M.W.: Different regional patterns of cortical thinning in alzheimer's disease and frontotemporal dementia. Brain **130**(4), 1159–1166 (2007)
5. Euler, L.: Institutiones calculi integralis, vol. 4. Academia Imperialis Scientiarum (1794)
6. Fischl, B.: Freesurfer. Neuroimage **62**(2), 774–781 (2012)
7. Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al.: Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. Neuron **33**(3), 341–355 (2002)
8. Fischl, B., Sereno, M.I., Dale, A.M.: Cortical surface-based analysis: Ii: inflation, flattening, and a surface-based coordinate system. Neuroimage **9**(2), 195–207 (1999)
9. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: Fastsurfer - a fast and accurate deep learning based neuroimaging pipeline. NeuroImage **219**, 117012 (2020)
10. Herrmann, L.R.: Laplacian-isoparametric grid generation scheme. Journal of the Engineering Mechanics Division **102**(5), 749–756 (1976)
11. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., et al.: The alzheimer's disease neuroimaging initiative (adni): Mri methods. Journal of Magnetic Resonance Imaging **27**(4), 685–691 (2008)
12. Jacobson, A., Panozzo, D., et al.: libigl: A simple C++ geometry processing library (2018), https://libigl.github.io/
13. 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.: Automated 3-d extraction and evaluation of the inner and outer cortical surfaces using a laplacian map and partial volume effect classification. Neuroimage **27**(1), 210–221 (2005)
14. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., et al.: Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. MedRxiv (2019)
15. Lebrat, L., Santa Cruz, R., de Gournay, F., Fu, D., Bourgeat, P., Fripp, J., Fookes, C., Salvado, O.: Corticalflow: A diffeomorphic mesh transformer network for cortical surface reconstruction. Advances in Neural Information Processing Systems **34** (2021)
16. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. ACM siggraph computer graphics **21**(4), 163–169 (1987)

17. Ma, Q., Robinson, E.C., Kainz, B., Rueckert, D., Alansary, A.: Pialnn: A fast deep learning framework for cortical pial surface reconstruction. In: International Workshop on Machine Learning in Clinical Neuroimaging. pp. 73–81. Springer (2021)

18. MacDonald, D., Kabani, N., Avis, D., Evans, A.C.: Automated 3-d extraction of inner and outer surfaces of cerebral cortex from mri. NeuroImage **12**(3), 340–356 (2000)

19. Muntoni, A., Cignoni, P.: PyMeshLab (Jan 2021). `https://doi.org/10.5281/zenodo.4438750`

20. Press, W., Flannery, B., Teukolsky, S.A., Vetterling, W.: Runge-kutta method. Numerical recipes in Fortran: The art of scientific computing pp. 704–716 (1992)

21. Rimol, L.M., Nesvåg, R., Hagler Jr, D.J., Bergmann, Ø., Fennema-Notestine, C., Hartberg, C.B., Haukvik, U.K., Lange, E., Pung, C.J., Server, A., et al.: Cortical volume, surface area, and thickness in schizophrenia and bipolar disorder. Biological psychiatry **71**(6), 552–560 (2012)

22. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)

23. Santa Cruz, R., Lebrat, L., Bourgeat, P., Fookes, C., Fripp, J., Salvado, O.: Deepcsr: A 3d deep learning approach for cortical surface reconstruction. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 806–815 (2021)

24. Schaer, M., Cuadra, M.B., Tamarit, L., Lazeyras, F., Eliez, S., Thiran, J.P.: A surface-based approach to quantify local cortical gyrification. IEEE transactions on medical imaging **27**(2), 161–170 (2008)

25. Shattuck, D.W., Leahy, R.M.: Brainsuite: an automated cortical surface identification tool. Medical image analysis **6**(2), 129–142 (2002)

26. Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., Seidel, H.P.: Laplacian surface editing. In: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing. pp. 175–184 (2004)

27. Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.: Pixel2mesh: Generating 3d mesh models from single rgb images. In: Proceedings of the European conference on computer vision (ECCV). pp. 52–67 (2018)