



# Improved Segmentation of Deep Sulci in Cortical Gray Matter Using a Deep Learning Framework Incorporating Laplace's Equation

Sadhana Ravikumar<sup>1</sup>(✉), Ranjit Ittyerah<sup>1</sup>, Sydney Lim<sup>1</sup>, Long Xie<sup>1</sup>, Sandhitsu Das<sup>1</sup>, Pulkit Khandelwal<sup>1</sup>, Laura E. M. Wisse<sup>2</sup>, Madigan L. Bedard<sup>1</sup>, John L. Robinson<sup>1</sup>, Terry Schuck<sup>1</sup>, Murray Grossman<sup>1</sup>, John Q. Trojanowski<sup>1</sup>, Edward B. Lee<sup>1</sup>, M. Dylan Tisdall<sup>1</sup>, Karthik Prabhakaran<sup>1</sup>, John A. Detre<sup>1</sup>, David J. Irwin<sup>1</sup>, Winifred Trotman<sup>1</sup>, Gabor Mizsei<sup>1</sup>, Emilio Artacho-Pérula<sup>3</sup>, Maria Mercedes Iñiguez de Onzono Martin<sup>3</sup>, Maria del Mar Arroyo Jiménez<sup>3</sup>, Monica Muñoz<sup>3</sup>, Francisco Javier Molina Romero<sup>3</sup>, Maria del Pilar Marcos Rabal<sup>3</sup>, Sandra Cebada-Sánchez<sup>3</sup>, José Carlos Delgado González<sup>3</sup>, Carlos de la Rosa-Prieto<sup>3</sup>, Marta Córcoles Parada<sup>3</sup>, David A. Wolk<sup>1</sup>, Ricardo Insausti<sup>3</sup>, and Paul A. Yushkevich<sup>1</sup>(✉)

<sup>1</sup> University of Pennsylvania, Philadelphia, USA  
{ravikums, pauly2}@penmedicine.upenn.edu

<sup>2</sup> Lund University, Lund, Sweden

<sup>3</sup> University of Castilla-La Mancha, Ciudad Real, Spain

**Abstract.** When developing tools for automated cortical segmentation, the ability to produce topologically correct segmentations is important in order to compute geometrically valid morphometry measures. In practice, accurate cortical segmentation is challenged by image artifacts and the highly convoluted anatomy of the cortex itself. To address this, we propose a novel deep learning-based cortical segmentation method in which prior knowledge about the geometry of the cortex is incorporated into the network during the training process. We design a loss function which uses the theory of Laplace's equation applied to the cortex to locally penalize unresolved boundaries between tightly folded sulci. Using an *ex vivo* MRI dataset of human medial temporal lobe specimens, we demonstrate that our approach outperforms baseline segmentation networks, both quantitatively and qualitatively.

**Keywords:** Cortical segmentation · topology correction · ex vivo MRI

## 1 Introduction

Segmentation of the cerebral cortex from MRI is an important first step in many neuroimaging pipelines such as quantitative morphometry analyses aimed at understanding the pathophysiology of neurological disorders. Automated

segmentation methods applied to the cortex are challenged by various artifacts such as image noise, partial volume effects and intensity inhomogeneities which make accurate identification of the tissue boundaries difficult and result in geometrically inaccurate cortical reconstructions. The cerebral cortex or gray matter (GM) can be defined as the space between two cortical surfaces; the pial surface which separates the GM from the surrounding cerebrospinal fluid (CSF), and the white matter (WM) surface which separates the GM from WM. The cortex has a complex geometry and is often modelled as a highly folded 2D sheet, with spatially varying curvature and thickness [14]. Geometrically accurate segmentation of the cortex requires accurate reconstruction of both the WM and pial cortical surfaces, complete with all cortical folds and narrow sulci. A commonly used simplification when solving cortical surface reconstruction problems is to view the cortical surfaces as having the topology of a 3D sphere [4]. However, unless explicitly corrected for, imaging artifacts often introduce topological defects in the resulting surface reconstructions. Defects due to partial volume effects are particularly apparent in tightly folded sulci and result in opposing banks of sulci appearing fused together. This creates either ‘bridged’ or ‘unresolved’ sulci in the resulting cortical reconstruction, which cause errors in downstream quantitative brain morphometry measures such as cortical thickness. While topological defects can be corrected by manual editing, these checks can be time-consuming.

Topology-corrected reconstruction of cortical surfaces is a well studied topic in *in vivo* neuroimaging literature, and several state-of-the-art methods have been developed to address this problem [4, 5, 11]. The widely used FreeSurfer framework employs a mesh-based approach to topology-correction which consists of two main steps [4]. First, the inner WM surface is generated by applying mesh tessellation to a volumetric WM segmentation that has been corrected for topological defects [4]. Second, this WM surface is expanded using a deformable surface model to reconstruct the outer, pial surface, while ensuring that the topology of the initial surface is preserved [4, 5, 11]. In recent years, elements of the FreeSurfer pipeline have been implemented as deep learning networks, resulting in significant speedups [1, 7, 8, 12]. However, these frameworks still either require the time-consuming post-processing step of topology correction [1, 7] or rely on a predefined initial mesh with the correct, spherical topology to reconstruct the cortical surface [8, 12].

In this work, we are specifically interested in developing an automated cortical segmentation method that can be applied to *ex vivo* brain MRI datasets to generate geometrically valid models of the cortex. Instead of a mesh deformation-based approach, we propose a novel volumetric deep image segmentation method that learns to segment the cortex while explicitly modeling the ‘sheet-like’ geometry of the cortex. In *ex vivo* studies, it is common to image only a portion of the brain hemisphere, thus violating the assumption of spherical WM topology made by mesh-based approaches. Furthermore, many of the *in vivo* cortical segmentation methods contain algorithms that are optimized for data with a standard 1 mm voxel size [4], and would result in unrealistic computational times if applied to high-resolution *ex vivo* MRI datasets. As a result, existing methods are not easily applicable to *ex vivo* MRI scans. As far as we know, no prior work has focused on topology correction of *ex vivo* cortical segmentations.

Previous studies have used the concept of Laplace’s equation as a tool for modelling the cortex [10,11,14]. By setting different boundary conditions at the WM/GM and GM/CSF interfaces, Laplace’s equation can be solved within the GM volume to generate a laminar ‘potential’ field that smoothly varies in value depending on its distance between the two cortical surfaces. The gradient of the Laplacian field can be used to compute cortical thickness [10] and defines an expansion path which guarantees topology-preserving deformation between the WM and pial cortical surfaces [11,14]. Building on this idea and the success of deep convolutional neural networks (CNN) in medical image segmentation tasks, here we design a differentiable numerical solver for Laplace’s equation and incorporate it within a deep segmentation framework to locally impose a Laplacian mapping between the predicted WM and pial surfaces. We train the segmentation network by comparing the predicted tissue segmentations and corresponding Laplacian field maps with the equivalent ground truth images, thus penalizing self-intersections in the predicted segmentations. Our results show that when compared to a baseline network trained without Laplacian constraints, our method is able to better reconstruct the intrinsic, layered geometry of the cortex. To our knowledge, this is the first time that an iterative numerical solver has been incorporated within a cortical segmentation network to directly compute Laplacian fields in an end-to-end setting.

## 2 Methods

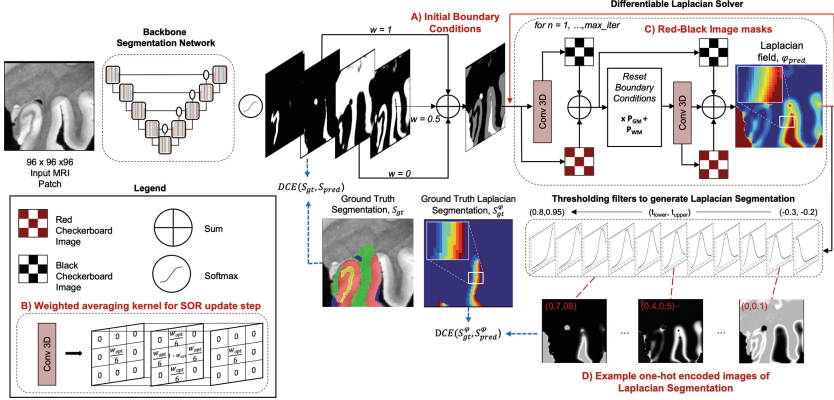
As illustrated in Fig. 1, our proposed framework builds upon any given backbone segmentation network (Sect. 2.1) by appending a numerical solver for Laplace’s equation to the output of the network. We reformulate the solver to be differentiable with respect to the input image to allow for gradient-based learning, used within standard CNN training (Sect. 2.2). In addition to the standard tissue segmentation loss, we design a loss function which compares the predicted Laplacian field to the solution of Laplace’s equation applied to the ground truth cortical segmentation, which is assumed to have correct topology (Sect. 2.3).

### 2.1 Backbone Segmentation Network

The proposed Laplacian solver is compatible with any semantic segmentation network since it only relies on the segmentation map output by the backbone network. We conducted experiments using two backbone networks for 3D image segmentation: the state-of-art nnU-Net framework [9] based on the U-net architecture, and nnFormer [18], a variant framework based on the recently popular transformer architecture. Both frameworks use image patches, deep supervision, and a variety of data-augmentation techniques to train the network [9].

### 2.2 Differentiable Laplacian Solver

To compute the Laplacian field corresponding to the cortical segmentation predicted based on a given input patch, the iterative solver for Laplace’s equation



**Fig. 1.** Schematic illustration of the proposed framework. A differentiable numerical solver for Laplace’s equation, based on the successive over relaxation (SOR) algorithm, is incorporated within end-to-end training of a segmentation network.

is appended after the final layer of the backbone network. Laplace’s equation,  $\Delta\varphi = 0$ , is a second-order partial differential equation, where  $\Delta$  is the Laplacian operator ( $\partial_{xx}^2 + \partial_{yy}^2 + \partial_{zz}^2$ ) and  $\varphi$  is a twice-differentiable, real-valued function. To solve Laplace’s equation within a domain (in our case, the GM volume), specific conditions need to be set that the Laplacian field  $\varphi$  must satisfy at the boundaries of the domain. In our case, we set  $\varphi_{x,y,z} = 0$  at the GM/WM boundary and  $\varphi_{x,y,z} = 1$  at the GM/pial boundary. Voxels within the GM domain are initialized with  $\varphi_{x,y,z} = 0.5$ .

Given these boundary conditions, the Laplacian field can be approximated using the finite-difference method, which is solved using an iterative numerical solver. Here, we use the Successive Over Relaxation (SOR) algorithm, a variant of the Gauss-Seidel method, to solve for the Laplacian field [6]. In the Gauss-Seidel method, given initial values for all the voxels in an image, at each iteration, the new value for a particular voxel within the GM volume,  $\tilde{\varphi}_{xyz}$ , is computed by taking the weighted sum of the most recently updated values of its six neighboring voxels (Eq. 1). The superscript  $n$  refers to the iteration of the algorithm and the subscripts are the voxel indices.

$$\tilde{\varphi}_{xyz}^n = \frac{1}{6}(\varphi_{x-1,y,z}^n + \varphi_{x+1,y,z}^{n-1} + \varphi_{x,y-1,z}^n + \varphi_{x,y+1,z}^{n-1} + \varphi_{x,y,z-1}^n + \varphi_{x,y,z+1}^{n-1}) \quad (1)$$

The SOR algorithm accelerates this approach by taking, at each iteration  $n$ , the weighted sum of the current solution and the solution from the previous iteration (Eq. 2). The over-relaxation parameter,  $\omega$  accelerates the rate of convergence of the Gauss-Seidel method when  $1 < \omega < 2$  [17]. It has been shown that the optimum value,  $\omega_{opt}$  is given by  $\omega_{opt} = \frac{2}{1 + \sin(\frac{\pi}{N+1})}$  where  $N$  is the minimum dimension of the input grid [17].

$$\begin{aligned} \varphi_{x,y,z}^n = & (1 - \omega_{opt})\varphi_{x,y,z}^{n-1} + \frac{\omega_{opt}}{6}(\varphi_{x-1,y,z}^n + \varphi_{x+1,y,z}^{n-1} \\ & + \varphi_{x,y-1,z}^n + \varphi_{x,y+1,z}^{n-1} + \varphi_{x,y,z-1}^n + \varphi_{x,y,z+1}^{n-1}) \end{aligned} \quad (2)$$

Instead of updating the value of each voxel in the image serially, computation of Eq. 2 can be parallelized using the Red-Black SOR approach, wherein the voxels in an image are divided into ‘red’ and ‘black’ following a checkerboard pattern [3]. During the update step, the ‘red’ voxels only depend on the values of the ‘black’ voxels and vice versa. Therefore, at each iteration, the Laplacian solution is updated in two steps; first, the update equation is applied to all of the black voxels in parallel and second, the update equation is applied to all of the red voxels in parallel, using the updated values computed at the black voxels. After convergence, the resulting Laplacian field contains values within the GM volume increasing smoothly from 0 at the WM surface to 1 at the pial surface.

To incorporate this numerical solver within a CNN, an important consideration is that the computations used to generate the Laplacian field must be differentiable with respect to the predicted tissue class probabilities to allow for back-propagation of the final loss through the network. To this end, we initialized the boundary conditions for the Laplacian solver by taking a weighted sum of the GM, WM and background probability maps, with weights of 0.5, 0 and 1 respectively (Fig. 1A). Additionally, the SOR update step (Eq. 2) was reformulated as a  $1 \times 1$  convolutional layer with fixed neighborhood weights specified in a  $3 \times 3 \times 3$  kernel (Fig. 1B). The voxels in the image were divided into a red and black grid by generating 3D ‘red’ ( $\text{mod}(x + y + z, 2) = 0$ ), and ‘black’ ( $\text{mod}(x + y + z, 2) = 1$ ) binary checkerboard images which were applied as image masks to retain values of interest after each convolutional layer (Fig. 1C). Lastly, the maximum number of iterations for the Laplacian solver was empirically set to 60, as a trade-off between computational time and convergence of the Laplacian solution. Since the solver numerically computes the solution to Laplace’s equation, it does not introduce any additional training parameters within the network.

### 2.3 Loss Function

To train the model, the backbone networks compare the predicted tissue segmentation,  $S_{pred}$  with the ground truth cortical segmentation,  $S_{gt}$  using a combination of Dice and cross-entropy loss,  $DCE(S_{gt}, S_{pred})$  [9]. We introduce an additional loss term which compares the Laplacian field computed from the predicted tissue segmentation,  $\varphi_{pred}$ , with the solution of Laplace’s equation applied to the ground truth segmentation,  $\varphi_{gt}$ . To simplify comparison between the predicted Laplacian field and the ground truth solution, we convert the Laplacian field computed by the solver to a multi-label segmentation,  $S_{pred}^\varphi$  using a series of thresholding functions. The advantage of this approach is that it enables the use of the same combined Dice and cross-entropy loss used by the backbone network on the outputs of the Laplacian solver, instead of the mean square error loss, thereby allowing us to equally weight the two loss terms:  $\mathcal{L} = DCE(S_{gt}, S_{pred}) + DCE(S_{gt}^\varphi, S_{pred}^\varphi)$ . To threshold the Laplacian field and create a multi-label segmentation in a differentiable way, the computed Laplacian field is passed through a product of two sigmoid functions,  $(1 + e^{-\beta(x-t_{lower})})^{-1} \times (1 + e^{\beta(x-t_{upper})})^{-1}$ , which together create a ‘band-pass’

thresholding filter.  $\beta$  controls the steepness of the filter, and  $t_{lower}$  and  $t_{upper}$  control the domain of the filter. Each filter creates an image that has values close to 1 for voxels in the Laplacian field lying within the domain of the filter, and values close to 0 otherwise. Therefore, by varying the lower and upper threshold values, a one-hot encoded image can be created, where each channel corresponds to a different label along the laminar axis of the GM. A multi-label Laplacian segmentation is then generated by applying the argmax operation to the computed one-hot encoded image (Fig. 1D).

### 3 Experiments

#### 3.1 Dataset

*MRI Image Acquisition.* To train and evaluate the proposed framework, we used *ex vivo* images of intact temporal lobe specimens, obtained from 27 brain donors from either the brain bank operated by the National Disease Research Interchange, or autopsies performed at the University of Pennsylvania Center for Neurodegenerative Disease Research (CNDR) and the University of Castilla-La Mancha (UCLM) Human Neuroanatomy Laboratory (HNL) in Spain. Brain specimens were obtained in accordance with the University of Pennsylvania Institutional Review Board guidelines, and the Ethical Committee of UCLM. Where possible, pre-consent during life and, in all cases, next-of-kin consent at death was given. Following 4+ weeks of fixation, the tissue specimens were scanned overnight on a 9.4 T 31 cm bore MRI scanner using a T2-weighted, multi-slice spin echo sequence (TE = 9330 ms, TR = 23 ms), with a resolution of  $0.2 \times 0.2 \times 0.2 \text{ mm}^3$ . Following image acquisition, the images were corrected for bias field non-uniformity and normalized to a common intensity range of [0, 1000]. In our work, we are specifically interested in segmenting the medial temporal lobe (MTL), a region affected early in Alzheimer’s Disease. Therefore, to facilitate semi-automated MTL segmentation, each scan was re-oriented so that the long axis of the MTL aligned with the anterior-poster direction.

*Ground Truth Tissue Segmentation.* To generate 3D segmentations of the MTL cortex, we adopted a semi-automatic interpolation technique [15]. The boundary of the MTL was manually traced approximately every 3 mm (i.e. 12–15 slices per dataset). Given the subset of labeled slices, the interpolation method uses contour and intensity information to compute the intermediate segmentations. This algorithm was applied iteratively, allowing the interpolated result to be reviewed and manually edited at each step to refine the segmentation. When editing, we ensured that in narrow and bridged sulci, the full extent of each sulcus was correctly labeled as background. Additionally, in a small region surrounding the MTL, the white matter and background voxels were semi-automatically labeled using a combination of intensity-based thresholding and morphological operations. The ground truth segmentations also contain a separate label for the stratum radiatum lacunosum moleculare (SRLM), which is the thin, hypo-intense layer within the hippocampus. We note that the ground truth segmentations

only cover the region in the image encompassing the MTL, and not the entirety of the *ex vivo* MRI scan.

*Ground Truth Laplacian Maps.* The proposed framework requires the Laplacian field maps corresponding to the ground truth segmentations to train the model. To solve Laplace’s equation within the ground truth GM volume, we used the iterative finite-differences approach, as implemented in [2]. This implementation employs a 26-neighbour average to compute the updated potential field and terminates the numerical solver when the Laplacian field change is below a specified threshold (sum of changes  $<0.001\%$  of total volume). To initialize the solver, source and sink boundary conditions were semi-manually labeled as the WM and pial surfaces of the MTL respectively. We note that the hippocampus voxels were not included in the GM domain of Laplace’s equation.

### 3.2 Implementation Details

We used Pytorch 1.9.1 and Nvidia Quadro RTX 500 GPUs to train the models. We implemented the differentiable Laplacian solver within the standardized training framework presented in [9] that is employed by both backbone networks. The framework includes pre-processing, automated hyper-parameter selection and fixed techniques for data augmentation. In our experiments, we made a few modifications to the default training parameters. First, since the ground truth segmentations only cover a portion of the input *ex vivo* MRI scans, we set the *ignore\_label* parameter in the loss function to 0 to exclude the unlabeled background voxels from the training process. Additionally, we increased the *oversample\_foreground* parameter such that only foreground patches are sampled during training. Lastly, we used an input patch size of  $96 \times 96 \times 96$ , to encourage the network to learn more local image features instead of larger contextual information, like the anatomical boundaries of the MTL. The networks were trained with a batch size of 2 (nnFormer) and 4 (nnU-Net), for 250 epochs in a five-fold cross validation setting. We used the results of the first fold to tune the network parameters. Consistent with the evaluation scheme used within nnU-Net, we aggregated the results across the remaining four folds for reporting test accuracy. We tested the performance of the network when using either 5 or 10 class labels (i.e. laminar layers) for generating the Laplacian segmentation. We found that increasing the number of class labels and training the network using Laplacian segmentations with a denser number of laminar layers improved the network’s ability to detect obscured sulci. To convert the Laplacian fields to segmentations, we used  $\beta = 10$  (softmax scaling parameter) and selected evenly spaced thresholds spanning the  $[0, 1]$  range of the Laplacian field. More specifically, the following lower ( $t_{lower}$ ) and upper ( $t_{upper}$ ) threshold values were used:  $[(-0.3, -0.2), (0, 0.1), (0.1, 0.2), (0.2, 0.3), (0.3, 0.4), (0.4, 0.5), (0.5, 0.6), (0.6, 0.7), (0.7, 0.8), (0.8, 0.95), (0.95, 1.05)]$ . Additionally, we tested the effect of increasing the weight given to the Laplacian segmentation loss relative to the tissue segmentation loss and found that it had minimal effect on cortical segmentation accuracy.

### 3.3 Evaluation

We compared the performance of our approach with the performance of the corresponding backbone segmentation networks, trained only with the tissue segmentation loss. We measured segmentation accuracy by computing the DSC between the predicted and ground truth tissue segmentations, and Laplacian field segmentations within the MTL region of interest. We also report the symmetric Hausdorff Distance (HD) 95<sup>th</sup> percentile between the predicted and ground truth segmentation of the MTL cortex. Since the numerical solvers used to generate the ground truth Laplacian fields and embedded in the network leverage different finite-difference approximation methods, during evaluation, we re-computed the Laplacian field for both the ground truth and predicted cortical segmentations using 120 iterations of the SOR solver used by the network, and computed the corresponding Laplacian segmentation with 5 laminar layers.

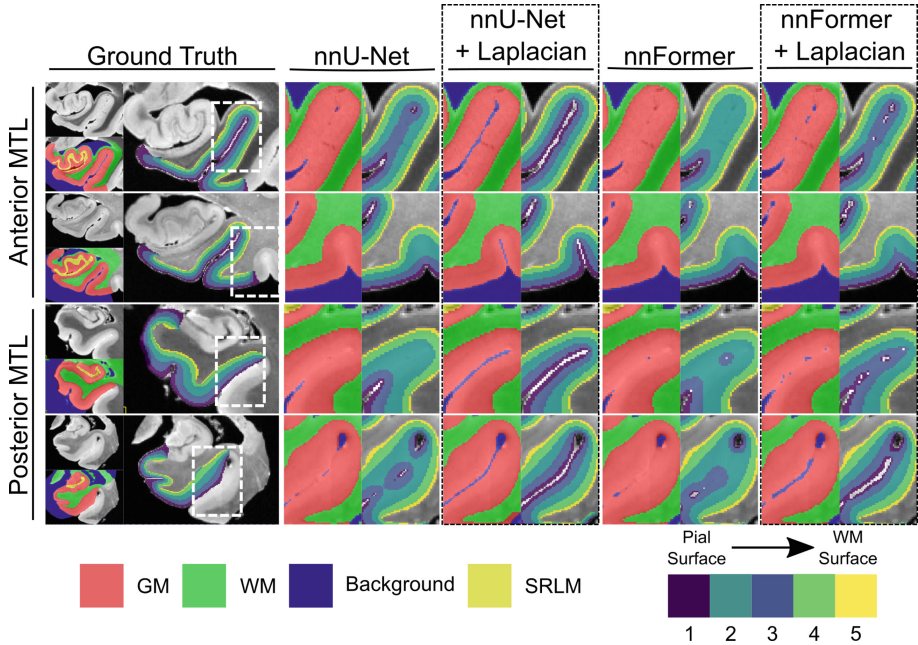
In a secondary analysis we evaluated the effect of introducing the Laplacian constraint on downstream cortical thickness measures. We applied the nnU-Net models to a dataset of 36 temporal lobe specimens obtained from individuals not included in the training dataset. For each specimen, we quantified MTL thickness at 6 manually identified landmarks corresponding to the anterior and posterior locations of MTL subregions Brodmann Area (BA) 35, BA36 and the parahippocampal cortex (PHC). We chose these subregions since they typically lie along the banks of the collateral sulcus (CS), and are therefore mostly likely to be affected by topological errors in the segmentation. For each location, we extracted the GM segmentation surrounding the landmark and measured cortical thickness using the pipeline described in [16]. In brief, given the GM segmentation surrounding a landmark, a maximally inscribed sphere is computed using Voronoi skeletonization [13], and the diameter of the sphere gives the thickness at that landmark. We compared thickness measurements obtained when using automatic GM segmentations generated by the baseline nnU-Net and the proposed model (*nnU-Net+Laplacian*), and reference thickness measurements computed using semi-automatic segmentations of the GM in terms of Pearson’s correlation and the average fixed-raters Intra-Class Correlation Coefficient (ICC).

## 4 Results and Discussion

### 4.1 Segmentation Accuracy

Table 1 presents the quantitative results, averaged across four cross-validation folds, evaluating the proposed framework using the corresponding backbone networks as baseline. Since sulci can be very thin structures, any improvements to mislabeled sulci would contribute minimally to the tissue segmentation DSC. Therefore, it is not surprising that we do not see statistically significant differences in the GM DSC measures. However, when looking at the Laplacian segmentation DSC, the results show that the Laplacian solver significantly improves upon the baseline performance of both nnU-Net and nnFormer in terms of better





**Fig. 2.** Qualitative comparison of the cortical segmentations generated by the proposed method and nnU-Net. Cross-sectional views are provided through the anterior and posterior MTL, using four different specimens. The white dashed boxes are used to indicate cortical folds demonstrating improved geometric accuracy. GM: Gray Matter; WM: White Matter; SRLM: Stratum Radiatum Lacunosum Moleculare (Color figure online)

preserving the layered structure of the cortex. This is visualized in Fig. 2 which provides a qualitative comparison of the predicted segmentations generated by each method. Labels 1–3 correspond to layers closest to the pial surface, and are therefore the class labels mostly likely to reflect errors such as bridged or unresolved sulci.

We observe that in the anterior portion of the MTL, the baseline models are often able to distinguish the sulcus, even without the Laplacian constraint (Fig. 2, row 2). This is likely because in the ground truth segmentation protocol, the labeled GM extends to include both banks of the CS in the anterior MTL, but only the medial bank of the CS in the posterior MTL. As a result, in the ground truth segmentation, the sulcus is clearly labeled in the anterior MTL and therefore included in the tissue segmentation loss. Conversely, in the posterior MTL, the ground truth tissue segmentation does not explicitly enforce the presence of the sulcus. However, in this region, the Laplacian segmentation term implicitly includes information about the location of the sulcus and therefore drives the network to learn the correct pial boundary of the cortex. To further investigate the contribution of the proposed loss function in the anterior and

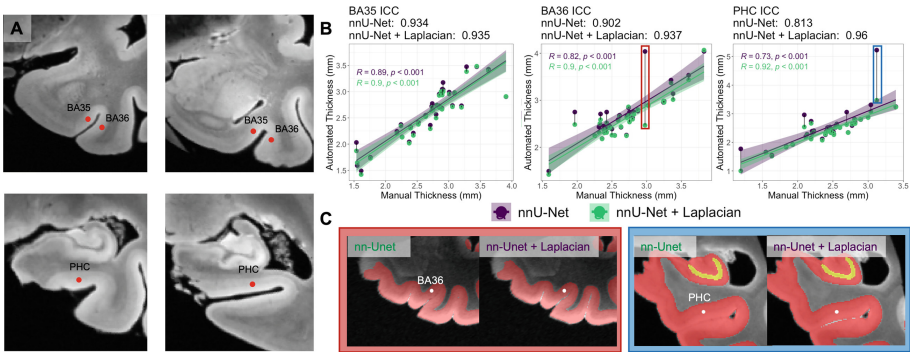
**Table 1.** Quantitative metrics comparing cortical segmentation accuracy of the proposed network and the baseline networks. We compute the Laplacian Segmentation accuracy, which reflects the networks ability to capture the layered nature of the cortex, across the whole MTL, and also separately for the anterior and posterior MTL. Metrics: Dice Score Coefficient (DSC) per label; Hausdorff Distance (HD). Standard deviations are reported in parentheses. (Statistical significance was assessed using paired t-tests; \*\*\* < 0.001, \*\*:  $p < 0.01$ , \*:  $p < 0.05$ )

	Method	DSC Laplacian Segmentation (%)					DSC Tissue Segmentation (%)				HD 95 (mm)
		Pial Surface → WM Surface					GM	WM	BG	SRLM	
Whole MTL	nnU-Net	78.3 (3.3)	78.6 (3.1)	77.4 (3.1)	77.7 (5.0)	58.7 (12.3)	94.5 (1.5)	96.0 (1.3)	95.3 (6.1)	85.8 (3.8)	0.341 (0.134)
	nnU-Net + Laplacian	80.7*** (2.6)	81.8*** (2.4)	80.4*** (3.5)	78.9* (5.8)	58.2 (13.4)	94.5 (1.7)	95.9 (1.3)	95.4 (5.9)	85.6 (3.8)	0.344 (0.205)
	nnFormer	75.5 (4.1)	75.5 (3.7)	74.6 (3.4)	75.9 (4.7)	57.0* (10.6)	93.8** (1.9)	95.1* (2.1)	95.1 (6.1)	83.8** (4.2)	0.559 (0.885)
	nnFormer + Laplacian	78.1*** (3.4)	79.8*** (3.0)	78.4*** (3.6)	77.3*** (5.2)	55.7 (11.6)	93.5 (2.2)	94.9 (2.3)	95.1 (6.0)	83.2 (4.5)	0.589 (0.980)
Anterior MTL	nnU-Net	80.0 (3.7)	80.6 (3.4)	79.9 (3.3)	79.6 (5.0)	59.3 (12.8)	94.5 (1.6)	95.0** (1.6)	95.7 (5.1)	85.9 (4.4)	0.208 (0.025)
	nnU-Net + Laplacian	81.6** (3.0)	82.8*** (2.7)	81.8*** (3.7)	80.3 (5.9)	58.1 (13.7)	94.4 (1.9)	94.7 (1.7)	95.9 (5.0)	85.7 (4.5)	0.217 (0.049)
	nnFormer	76.8 (4.5)	76.9 (4.3)	76.5 (4.4)	78.0 (5.1)	58.4** (11.4)	93.8* (2.5)	93.9 (3.6)	95.7 (5.2)	83.7** (5.0)	0.208 (0.0205)
	nnFormer + Laplacian	79.3*** (3.4)	80.9*** (3.0)	80.0*** (3.8)	79.1** (5.2)	56.8 (12.1)	93.6 (2.7)	93.7 (3.9)	95.8 (5.1)	83.1 (5.4)	0.212 (0.030)
Posterior MTL	nnU-Net	74.8 (6.2)	74.7 (5.8)	73.2 (5.4)	74.6 (6.1)	57.2 (12.9)	94.7 (1.6)	96.9 (1.4)	94.3* (9.0)	85.3 (3.4)	0.228 (0.096)
	nnU-Net + Laplacian	78.9*** (4.7)	79.8*** (3.9)	78.1*** (4.1)	76.8*** (6.5)	58.0 (13.5)	94.6 (1.7)	97.0 (1.1)	94.1 (9.0)	85.2 (3.4)	0.223 (0.062)
	nnFormer	72.7 (8.1)	72.8 (7.0)	71.1 (6.6)	72.7 (6.9)	55.0 (11.9)	93.9*** (2.5)	96.2* (2.0)	93.8* (8.9)	84.3* (4.0)	0.348 (0.430)
	nnFormer + Laplacian	75.4*** (8.3)	77.5*** (6.9)	75.4*** (6.6)	74.3** (7.9)	54.0 (13.1)	93.2 (3.2)	96.0 (2.2)	93.6 (9.0)	83.6 (4.0)	0.436 (0.692)

posterior MTL, we computed the DSC metrics of the Laplacian segmentations separately for the anterior and posterior MTL (Table 1). We observe that even when considering the anterior MTL on its own, the proposed framework improves Laplacian segmentation accuracy compared to the baseline networks, confirming that the addition of the Laplacian term is in fact contributing towards the network better learning the layered organization of the cortex. This is further seen in the *nnU-Net+Laplacian* result in Fig. 2, row 2, where the proposed network is able to detect a buried sulcus in a cortical fold not included within the ground truth region of interest. We note that Laplacian segmentation label 5, which corresponds to the innermost cortical surface at the GM/WM boundary, forms a very thin layer and therefore has greater variation in DSC compared to the other labels.

## 4.2 Downstream Thickness Measurements

Figure 3 shows the results of the morphometry analysis correlating automated and manual measurements of cortical thickness in BA35, BA36 and the PHC, when using cortical segmentations generated with and without the Laplacian-based loss term. Since we found that nnU-Net achieves better segmentation performance than nnFormer, we only conducted experiments using nnU-Net in



**Fig. 3.** A) Example scan showing the 6 landmarks where cortical thickness is measured. For each subregion, the thickness measurement is averaged across two landmark locations. B) Scatter plots showing the correlation between automated segmentation-based cortical thickness measurements and reference measurements based on semi-automatic segmentations for three MTL subregions, with (green) and without (purple) the Laplacian constraint. C) Segmentations produced by *nnUnet* and *nnU-Net+Laplacian* for BA36 (red) and PHC (blue) landmarks where thickness measures derived from the two networks differed the most. BA: Brodmann Area; PHC: parahippocampal cortex (Color figure online)

the secondary analysis. Compared to the baseline *nnU-Net*, we observe that the thickness measurements of BA36 and the PHC computed using the proposed network are more strongly correlated with manual measurements, in terms of both correlation coefficient and ICC. Both models achieve similar correlations in BA35. BA36 is located in the anterior MTL, whereas PHC corresponds to the posterior MTL. The strengthened correlations in both these regions further demonstrate that the proposed method is able to improve the accuracy of the predicted segmentations across the whole length of MTL.

## 5 Conclusions

We present a novel deep learning-based solution for cortical segmentation, applied to *ex vivo* MRI, that is able to learn the layered geometry of the cortex by locally imposing Laplacian mappings between the predicted WM and pial cortical surfaces. A limitation of this approach is the long run-time of the iterative solver during training ( $\sim 9x$  slower/epoch relative to the backbone). However, at inference time, the input image is only passed through the backbone segmentation network which typically takes 3–5 min per scan. Another limitation is the need for the sulci to be well delineated in the training data. In the future, we will explore ways to relax this requirement, perhaps using additional geometric priors. While in this work we demonstrate the utility of our approach in the context of MTL cortical segmentation, this approach can be extended to other high-resolution neuroimaging datasets such as *ex vivo* whole hemisphere scans

or other *in vivo* image segmentation tasks which involve similar sheet-like structures. Future work will focus on applying this method to *in vivo* brain MRI, thus allowing for evaluation of our approach against existing cortical surface reconstruction methods.

**Acknowledgements.** We gratefully acknowledge the tissue donors and their families. This work was supported by the NIH (Grants RF1 AG069474, P30 AG072979 and R01 AG056014), a UCLM travel and research grant (to R.I), and an Alzheimer’s Association grant (AARF-19-615258) (to L.E.M.W).

## References

1. Cruz, R.S., 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)
2. DeKraker, J., Ferko, K.M., Lau, J.C., Köhler, S., Khan, A.R.: Unfolding the hippocampus: an intrinsic coordinate system for subfield segmentations and quantitative mapping. *Neuroimage* **167**, 408–418 (2018)
3. Epicoco, I., Mocavero, S.: The performance model of an enhanced parallel algorithm for the SOR method. In: Murgante, B., et al. (eds.) ICCSA 2012. LNCS, vol. 7333, pp. 44–56. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-31125-3\\_4](https://doi.org/10.1007/978-3-642-31125-3_4)
4. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
5. Han, X., Pham, D.L., Tosun, D., Rettmann, M.E., Xu, C., Prince, J.L.: CRUISE: cortical reconstruction using implicit surface evolution. *Neuroimage* **23**(3), 997–1012 (2004)
6. Hansen, P.B.: Numerical solution of Laplace’s equation (1992)
7. 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)
8. Hoopes, A., Iglesias, J.E., Fischl, B., Greve, D., Dalca, A.V.: TopoFit: rapid reconstruction of topologically-correct cortical surfaces. In: Medical Imaging with Deep Learning (2021)
9. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**(2), 203–211 (2021)
10. Jones, S.E., Buchbinder, B.R., Aharon, I.: Three-dimensional mapping of cortical thickness using Laplace’s equation. *Hum. Brain Mapp.* **11**(1), 12–32 (2000)
11. Kim, J.S., et al.: 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)
12. 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 (2021)
13. Ogniewicz, R., Kübler, O.: Hierarchic Voronoi skeletons. *Pattern Recognit.* **28**(3), 343–359 (1995)
14. Osechinskiy, S., Kruggel, F.: Cortical surface reconstruction from high-resolution MR brain images. *Int. J. Biomed. Imaging* **2012** (2012)

15. Ravikumar, S., Wisse, L., Gao, Y., Gerig, G., Yushkevich, P.: Facilitating manual segmentation of 3D datasets using contour and intensity guided interpolation. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019), pp. 714–718 (2019)
16. Wisse, L.E., et al.: Downstream effects of polypathology on neurodegeneration of medial temporal lobe subregions. *Acta Neuropathol. Commun.* **9**(1), 1–11 (2021)
17. Yang, S., Matthias, K.G.: The optimal relaxation parameter for the SOR method applied to a classical model problem. Technical report, Technical Report TR2007-6, University of Maryland, Baltimore County (2007)
18. Zhou, H.Y., Guo, J., Zhang, Y., Yu, L., Wang, L., Yu, Y.: nnFormer: interleaved transformer for volumetric segmentation. arXiv preprint [arXiv:2109.03201](https://arxiv.org/abs/2109.03201) (2021)