

This WACV 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# DeepCSR: A 3D Deep Learning Approach for Cortical Surface Reconstruction

Rodrigo Santa Cruz<sup>1,2</sup>, Leo Lebrat<sup>1,2</sup>, Pierrick Bourgeat<sup>1</sup>, Clinton Fookes<sup>2</sup>, Jurgen Fripp<sup>1</sup>, and Olivier Salvado<sup>1,3</sup>

<sup>1</sup>CSIRO Health & Biosecurity, The Australian eHealth Research Centre, Brisbane, QLD, Australia <sup>2</sup>Image and Video Laboratory, Queensland University of Technology (QUT), Brisbane, QLD, Australia <sup>3</sup>CSIRO Data61, Brisbane, Australia

rodrigo.santacruz@csiro.au

### Abstract

The study of neurodegenerative diseases relies on the reconstruction and analysis of the brain cortex from magnetic resonance imaging (MRI). Traditional frameworks for this task like FreeSurfer demand lengthy runtimes, while its accelerated variant FastSurfer still relies on a voxelwise segmentation which is limited by its resolution to capture narrow continuous objects as cortical surfaces. Having these limitations in mind, we propose DeepCSR, a 3D deep learning framework for cortical surface reconstruction from MRI. Towards this end, we train a neural network model with hypercolumn features to predict implicit surface representations for points in a brain template space. After training, the cortical surface at a desired level of detail is obtained by evaluating surface representations at specific coordinates, and subsequently applying a topology correction algorithm and an isosurface extraction method. Thanks to the continuous nature of this approach and the efficacy of its hypercolumn features scheme, DeepCSR efficiently reconstructs cortical surfaces at high resolution capturing fine details in the cortical folding. Moreover, DeepCSR is as accurate, more precise, and faster than the widely used FreeSurfer toolbox and its deep learning powered variant FastSurfer on reconstructing cortical surfaces from MRI which should facilitate large-scale medical studies and new healthcare applications.

### **1. Introduction**

The study of many neurodegenerative diseases and psychological disorders, rely on the analysis of the cerebral cortex using magnetic resonance imaging (MRI) [1, 6, 32]. As shown in Figure 1a, the cortex can be visually described as the inner volume between two cortical surfaces: the inner surface which is the interface between the white matter (WM) and the gray matter (GM) tissues, and the outer surface which is the interface between the gray matter tissue and the cerebrospinal fluid (CSF). Therefore, the goal of cortical surface reconstruction is to estimate reliable, accurate, and topologically correct [30] triangular meshes, for the inner and outer cortical surfaces from a given MRI.

Reconstructing cortical surfaces from MRI is a challenging problem. First, cortical surfaces vary significantly across individuals as exemplified in Figure 1b by overlaying axial slices of co-registered MRIs from 3 different patients. Second, voxels near the cortex tend to hold more than one tissue type due to the discrete nature of MRI and the folding patterns of the cortical surfaces which is known as partial volume effect (PVE) [2] in medical imaging. Consequently, sub-voxel variations of these surfaces can not be captured by pure voxel-wise segmentation approaches. This limitation leads to oversmoothed reconstructed cortical surfaces even when manual expert segmentation is employed as shown in Figure 1c. Finally, we need to enforce the spherical topology (i.e., homeomorphic to a sphere) of the reconstructed surfaces to allow surface-based analysis [34, 39] and visualizations [7]. Figure 1d emphasizes that arbitrarily small defects, known as *holes*, can drastically change the topology of the reconstructed surface.

Despite these challenges, there exist frameworks for cortical reconstruction from MRI [4, 5, 11, 16, 18, 21, 40]. Traditionally, they consist of extensive pipelines of handcrafted image processing algorithms subject to careful hyper-parameter tuning (*e.g.*, thresholds, iteration numbers, and convergence criterion) and very long runtimes which prevent them from being employed in healthcare applications where immediate results are critical. For instance, the well-known FreeSurfer's V6.0 [5] takes around six hours per scan depending on the quality of the input MRI. Concur-

Thanks to Maxwell plus: http://maxwellplus.com/



Figure 1: Cerebral cortex and its reconstruction challenges. (a) Depicts the brain cortex from an axial view. It also provides schematics for the partial volume effect (PVE) problem and the inaccurate boundary estimation using partial volume mapping. (b) Overlays the cortical surfaces of co-registered MRIs from 3 different individuals exemplifying that the visual variability across individuals is like a fingerprint. (c) Shows that meshes obtained from manual segmentation provided by experts, e.g. Neuromorphometrics, can not capture high-frequency details in the cortical surfaces. (d) Present examples of small surface mesh defects that change the topology of the reconstructed surface.

rently to our work, Henschel et al. [11] propose FastSurfer which alleviates this burden by integrating a fast and accurate deep learning based brain segmentation model to the FreeSurfer cortical surface reconstruction pipeline. However, these frameworks still rely on voxel-wise segmentation of the input MRI which has a limited resolution or employ partial volume mapping which also fails to define the tissue boundaries accurately. More specifically, multiple configurations of a tissue boundary have the same partial volume assignment as illustrated in Figure 1a.

In this paper, we propose a 3D deep learning framework for cortical surface reconstruction from MR images named DeepCSR. More specifically, we first reformulate this problem as the prediction of an implicit surface representation for points in a continuous coordinate system. Then, the cortical surfaces are extracted using this implicit surface representation, a lightweight topological correction method, and an isosurface mesh extraction technique. Since our framework predicts implicit surface representation for real-valued points instead of voxels, providing a continuous approximation to surfaces, we can reconstruct cortical surfaces at a higher resolution than the input MR image. This continuous formalism allows us to overcome the PVE problem without relying on inaccurate estimations like partial volume mapping. We also develop a neural network architecture with hypercolumn features able to capture local information from the input MR image allowing the reconstruction of fine details describing the cortical folding.

In order to validate our approach, we first provide ablative studies of its main components. Then, we compare *DeepCSR* to the FreeSurfer V6.0 cross-sectional pipeline and FastSurfer for cortical reconstruction on the Test-Retest [22] and MALC [19, 23] datasets. We conclude that the proposed *DeepCSR* is able to reconstruct cortical surfaces faster and with higher reliability than the competitors, thus facilitating new healthcare applications.

# 2. Related Work

Cortical surface reconstruction. Traditionally, cortical surface reconstruction frameworks involve a lengthy sequence of image processing techniques [4, 5, 16, 18, 21, 40]. For instance, the widely used FreeSurfer [5] performs multiple image normalization routines, linear and non-linear registration techniques, brain tissue labeling, surface fitting models, and topology correction between other techniques to reconstruct the cortical surfaces. The other traditional frameworks follow a similar approach mostly differing in the mechanism used to fit the desired surfaces onto the segmented volumes. As examples, Kim et al. [16] leverage a Laplacian field with stretch and self-proximity regularization terms to prevent mesh self-intersection and Han et al. [9] introduce a topology-preserving geometric deformable surface model. Therefore, these frameworks require a considerable processing time and eventually manual intervention of experts to fine-tune the parameters of some of these image processing techniques [14] which limit their range of applications. Differently, we propose a swift framework able to reconstruct cortical surfaces efficiently by leveraging modern 3D deep learning models.

Concurrently to our work, Henschel et al. [11] propose a faster variant for the FreeSurfer pipeline named FastSurfer. In that work, expensive computations of FreeSurfer are replaced with modern and lightweight alternatives as a deep learning model for whole-brain segmentation and a spectral mesh processing algorithm for spherical mapping. However, this approach still relies on voxel-wise brain segmentation which does not approximate continuous surfaces well due to its discrete nature as exemplified in Figure 1c. In contrast, we propose a 3D deep learning model that directly predicts implicit surface representations for real coordinates in the brain MRI providing a continuous approximation to the cortical surfaces. Nonetheless, it is also important to acknowledge that FreeSurfer and FastSurfer are complete brain morphometry tools [41] providing further volumetric and surface-based analysis which is beyond the scope of this work. Note also that recent approaches tend to indicate that one could bypass the estimation of cortical surfaces directly predicting averaged morphometric measurements from the MR images [33, 37].

Deep learning models for 3D reconstruction. It comprises the set of deep learning models that learn to reconstruct 3D geometries, e.g., objects and scenes, from data [12, 38]. These deep learning models can be broadly categorized according to the 3D shape representation they operate on as either mesh-based, voxel-based, or implicit surfaces models. Mesh-based methods exploit explicit representations like template meshes [8, 31], graphs [45, 47], and parametric 3D models [17, 27] to predict surfaces and object shapes. While these approaches are very convenient, they are also prone to produce noisy meshes with a large number of self-intersections which is very undesirable for cortical surface reconstruction. On the other hand, voxel-based methods reconstruct 3D shapes by predicting 3D voxel-grids of surface representations like discretized signed distances [28] and level-sets [26]. The main limitation of these approaches is to treat the output 3D space as a discretized grid that may not capture the fine details of the cortical surface folding. On the other hand, implicit surface models explore continuous representations of the 3D space where the objects are defined as 3D scalar fields [25, 28]. These approaches are memory efficient and can explore the finer details of the target geometries. Therefore, we follow this approach and develop a novel deep learning model to reconstruct cortical surfaces from MR images.

### 3. Method

In this section, we introduce *DeepCSR*, a deep learning framework for cortical surface reconstruction from MR images. First, we describe how to implicitly define surfaces by functions on the 3D Euclidean space, and how to learn these functions using a deep learning model. Then, we show how to perform cortical surface reconstruction using this model.

### 3.1. Learning Implicit Surface Representation

In order to ease the presentation, we focus on the learning of a single target surface, but this framework can be generalized to multiple surfaces by introducing mild modifications. Without loss of generality, one can implicitly define a Lipschitz surface  $S \subset \mathbb{R}^3$  by the  $\ell$ -level set of a continuous function  $f_S : \mathbb{R}^3 \mapsto \mathbb{R}$  as,

$$S = \left\{ p \in \mathbb{R}^3 \mid f_S(p) = \ell \right\},\tag{1}$$

where the function  $f_S$  maps to a scalar for every point  $p \in \mathbb{R}^3$  in the 3D Euclidean space. Furthermore, given  $f_S$ ,

the mesh representation of S can be obtained by iso-surface extraction algorithms like marching cubes [20] or ray casting [44]. Therefore, reconstructing a specific surface can be thought as modeling a function such that its  $\ell$ -level set defines the desired geometry.

For a given surface S, the function  $f_S$  can be modeled, for instance, using occupancy field or signed distance function. The former is a binary valued function defined as,

$$r_S^{occ}(p) = \mathbb{1}_{p \in S_{int}},\tag{2}$$

where  $\mathbb{1}_{pred}$  is the indicator function evaluating to one, if its predicate *pred* is true, and zero otherwise. In this case, the level set of interest to retrieve the surface is  $\ell = \frac{1}{2}$ . The occupancy field  $r_S^{occ}$  divides the 3D space in the interior  $S_{int}$  and exterior of the oriented surface S. On the other hand, the signed distance function consists of the Euclidean distance between every point in the 3D space and its projection,  $\operatorname{proj}_S(p) = \operatorname{argmin}_{q \in S} ||p - q||_2$ , on the surface S. Mathematically,

$$r_{S}^{sdf}(p) = (2 r_{S}^{occ}(p) - 1) \|p - \operatorname{proj}_{S}(p)\|_{2}.$$
 (3)

Therefore,  $r_S^{sdf}$  assigns positive distances to points in the interior of the surface  $S_{int}$ , zero to points on the surface S (*i.e.*,  $\ell = 0$ ), and negative distances otherwise. These tools are widely used in surface modeling, then we evaluate both in the context of cortical surface reconstruction.

Our final goal is to reconstruct surfaces according to an observed MR image. Rather than conditioning the function  $f_S$  on a specific surface, one should condition it on an input 3D MR image. While this formulation seems logical, we notice that MR images may be defined in different coordinate systems producing an unbounded input space for learning. We overcome this difficulty by co-registering the input MR images to a brain template constraining the space where the target surfaces exists to the coordinates of the bounding box containing the template brain in the template coordinate system. Therefore, our problem boils down to model the function,

$$f: \Omega \times \mathcal{I} \mapsto \mathbb{R},\tag{4}$$

where  $\Omega \subset \mathbb{R}^3$  is a closed and bounded subspace where the template brain is defined and  $\mathcal{I}$  denotes the space of MR images *I* represented by a 3D grid of voxels.

Using a machine learning approach to model such a function, we first define a dataset  $\mathcal{D} = (I_i, S_i)_{i=1,...,N}$  of N MR images  $I_i$  and their corresponding surfaces  $S_i$  which can be obtained using traditional pipelines for automatic cortical surface reconstruction [4, 5, 40]. Then, we parametrize the function  $f_{\theta}$  in terms of learnable parameters  $\theta$  and minimize the regularized empirical risk on the dataset  $\mathcal{D}$ . This learning problem can be stated as,

$$\underset{\theta}{\operatorname{argmin}} \ \frac{1}{N} \sum_{i=1}^{N} \int_{\Omega} \Delta\left(r_{S_{i}}^{*}(p), f_{\theta}(p, I_{i})\right) \mathrm{d}p + \mathcal{R}(\theta), \quad (5)$$



Figure 2: *DeepCSR*'s neural network architecture. It receives as input an image I and the coordinates of a point p in the template coordinate system  $\Omega$ . Then, it predicts the implicit surface representation  $f_{\theta}(p, I)$  for the inner and outer cortical surfaces divided into left and right hemispheres.

where  $\mathcal{R}$  is some regularization function on the parameters  $\theta$  and  $\Delta$  is an appropriate cost function measuring discrepancies between the predicted implicit surface representation  $r_{S_i}^*(p)$  for the given point  $p \in \Omega$ , MR image  $I_i$ , and its corresponding surface  $S_i$ . In the case of occupancy field, we minimize the binary cross-entropy classification loss, while for signed distance function we opt for the  $L^1$ -loss.

In order to use this approach for cortical surface reconstruction, we extend the scalar-valued functions  $f_{\theta}$  to a vector-valued function  $f_{\theta} \in \mathbb{R}^4$  where the predicted values are the surfaces representations for the inner and outer cortical surfaces further divided into left and right brain hemispheres. Consequently, the learning problem descried by Equation 5 is solved jointly for these four cortical surfaces across all of the MRIs in the dataset  $\mathcal{D}$ . This formulation is an instance of multi-task learning which often produces models with good generalization ability [36].

In practice, we approximate the integral in Equation 5 by a finite sum of sampled locations p in the reference space  $\Omega$ . We observed that the sampling scheme employed is critical to the quality of the reconstructed surface. More specifically, as shown later in Section 4.1, the proposed framework reconstructs more accurate cortical surfaces when we sample more points near the surface, in addition to uniformly on the template space  $\Omega$ . In order to sample points near the target surface, we first sample faces proportionally to their area, then we sample points uniformly in these faces using the triangle point picking method [46]. Finally, we perturb these sampled points on the surface by adding a Gaussian  $\mathcal{N}(0,1)$  centered at zero with 1 mm variance getting points near the surface. Such a sampling scheme provides global information to correctly align the predicted surface to the brain and a lot of local information to capture the highly folded geometry of the brain cortex.

# 3.2. Neural Network Architecture and Hypercolumn Features

We implement the parameterized function  $f_{\theta}$  as an encoder-decoder neural network. The encoder takes as input a point p represented by its coordinates in the MNI105 space and a co-registered MR image as a 3D voxel grid I. It first processes the input MR image by a sequence of 3D convolutional layers (Conv3D), ReLu activation functions, max pooling layers (Max3D), and fully connected layers (FC) producing multiple feature maps as indicated in the top of Figure 2. The convolutional layers use  $3^3$  kernels, stride equal to two, and padding equal to one, with the exception of the first convolutional layer where the stride is equal to one. Our goal is to produce a rich hierarchy of visual features with different level of details. Then, inspired by Hariharan et al. [10], we project the input point coordinates into each feature map, linearly interpolate a feature value at these projected locations, and concatenate these interpolated values to form a 512-dimensional hypercolumn feature vector c(I, p). This hypercolumn vector holds global



Figure 3: *DeepCSR* framework for cortical surface reconstruction from MRI.

and local visual cues for predicting the surface representation at the input point p since it collects features from different levels of our hierarchy of features maps. As shown later in Section 4.1, this architectural design is critical to capture the high frequency details on the target surfaces.

On the other hand, the decoder receives as input the point coordinates p and its corresponding hypercolumn vector c(I, p) generated by the encoder for the image I as previously described. The decoder processes these inputs through a sequence of fully connected layers, conditional batch normalization layers (CBN) [29] and ReLU activation functions as also indicated in Figure 2, outputting the predicted implicit surfaces representation  $f_{\theta}(p, I) \in \mathbb{R}^4$  for the inner and outer cortical surfaces divided into left and right brain hemispheres. The decoder network uses skip-connections and follows the method proposed by Mescheder et al. [25]. However, our CBN layers compute a non-linear function conditioned on the hypercolumn vector c(I, p) which encodes global and local visual cues for the prediction of the desired surfaces representation.

We train such a model from scratch by optimizing the objective in Equation 5 extended for multiple surfaces using stochastic gradient descent. More specifically, we use Adam optimizer with an initial learning rate equal to  $10^{-4}$  and back-propagate the loss computed on mini-batches of 5 MR images and 1024 sampled points per image from a precomputed pool of 4 million points whereby 10% of them are sampled uniformly in  $\Omega$  and 90% are sampled near the target surfaces as explained in Section 3.1.

#### 3.3. Reconstructing Cortical Surfaces

*DeepCSR* receives as input a MRI scan and outputs mesh representations for the outer and inner cortical surfaces further divided into the left and right brain hemispheres. As illustrated in Figure 3, we first perform an affine registration of the input MRI scan to the MNI105 brain template [24]. This registration aims to unify the coordinate systems across different MR scans easing the learning and prediction of implicit surfaces.

Second, we construct a Cartesian grid of points at the desired resolution by dividing the template bounding box into evenly spaced points. Using the trained neural network  $f_{\theta}$ , we predict implicit surface representations at these points, represented by their continuous coordinates, for those four cortical surfaces according to the input MRI. Throughout this paper we reconstruct cortical surfaces using  $512^3$  points, unless mentioned, providing sub-voxel precision. The model only needs to extract the feature maps once and can process the points in parallel, generating surfaces at high resolution efficiently.

Third, since most of the applications of cortical surface reconstruction require surfaces with spherical topology (*i.e.*, genus zero) as discussed in Section 1, we employ a lightweight topology correction algorithm to fix the topological defects caused by wrong predictions of the implicit surface representation. Specifically, we use the method proposed by Bazin and Pham [3] which consists of a spherical level-set evolution that avoids evolving over critical points [30], outputting a 3D implicit surface volume with guaranteed spherical topology. Each target surfaces' topology is corrected independently, and these tasks are performed in parallel for efficiency. Note that the proposed framework is agnostic to the topology correction method allowing the application of other techniques.

Finally, in order to obtain the mesh representation of the target surfaces, we apply a topology preserving marching cubes algorithm [20]. The predicted surfaces are in the template brain coordinate system, but these can be mapped back to the native space using the inverse of the transformation obtained during the registration of the input MR scan.

### 4. Experiments

We now evaluate the performance of our method on the reconstruction of cortical surfaces from MR images. In Section 4.1, we present an ablative study while, in Section 4.2, we compare it to the FreeSurfer V6 (crosssectional pipeline) and FastSurfer.

#### 4.1. Ablation Studies

We now perform experiments with variations of the proposed *DeepCSR* to measure the importance of its main components. The quantitative results are presented in Table 1, while the qualitative results are shown in Figure 4. See below a brief summary of the dataset and evaluation metrics used, in addition to the detailed description of these experiments and the discussion of their results.

**Dataset.** We use the MRI data provided by the Alzheimer's Disease Neuroimaging Initiative (ADNI) [15] and their corresponding pseudo-ground truth surfaces generated with FreeSurfer V6.0. The ADNI dataset consists of 3876 MRI images from 820 different subjects collected at different time points. We split this dataset by subjects obtaining 2353 MRI scans from 492 subjects for training ( $\approx 60\%$ ), 375 MRI scans from 82 subjects for validation ( $\approx 10\%$ ) and 1148 MRI scans from 246 subjects for testing ( $\approx 30\%$ ). We train the models on the training set until their loss plateau

	Left Outer Surface			Right Outer Surface			Left Inner Surface			Right Inner Surface		
	EMD	AD (mm)	HD (mm)	EMD	AD (mm)	HD (mm)	EMD	AD (mm)	HD (mm)	EMD	AD (mm)	HD (mm)
DeepCSR (SDF.)	7.084	0.298	0.654	7.083	0.294	0.651	6.440	0.267	0.562	6.401	0.260	0.542
	(± <b>0.797</b> )	(± <b>0.184</b> )	(± <b>0.596</b> )	(±0.851)	(± <b>0.155</b> )	(± <b>0.530</b> )	(± <b>0.810</b> )	(± <b>0.216</b> )	(± <b>0.725</b> )	(± <b>0.692</b> )	(± <b>0.162</b> )	(± <b>0.523</b> )
DeepCSR (Occ.)	7.133	0.312	0.711	7.108	0.305	0.672	6.459	0.281	0.587	6.522	0.289	0.582
	$(\pm 0.757)$	$(\pm 0.167)$	$(\pm 0.525)$	$(\pm 0.857)$	$(\pm 0.134)$	$(\pm 0.554)$	$(\pm 0.802)$	$(\pm 0.284)$	$(\pm 0.724)$	$(\pm 0.632)$	$(\pm 0.127)$	$(\pm 0.570)$
Uniform Sampling	7.342	0.510	1.185	7.415	0.538	1.326	6.584	0.501	1.100	6.537	0.465	0.997
	$(\pm 0.856)$	$(\pm 0.208)$	$(\pm 0.665)$	$(\pm 0.868)$	$(\pm 0.169)$	$(\pm 0.524)$	$(\pm 0.804)$	$(\pm 0.236)$	$(\pm 0.728)$	$(\pm 0.691)$	$(\pm 0.184)$	$(\pm 0.521)$
No hypercolumns	10.892	5.145	12.252	10.649	5.173	12.346	15.644	5.629	12.811	15.434	5.698	12.917
	$(\pm 1.375)$	$(\pm 0.489)$	$(\pm 1.267)$	$(\pm 1.201)$	$(\pm 0.437)$	$(\pm 1.106)$	$(\pm 1.914)$	$(\pm 0.603)$	$(\pm 1.254)$	$(\pm 1.814)$	$(\pm 0.591)$	$(\pm 1.113)$
Single Surface Model	7.109	0.314	0.702	7.088	0.323	0.731	6.454	0.280	0.568	6.429	0.278	0.564
	$(\pm 0.872)$	$(\pm 0.229)$	$(\pm 0.885)$	$(\pm 1.042)$	$(\pm 0.249)$	$(\pm 0.958)$	$(\pm 0.807)$	$(\pm 0.199)$	$(\pm 0.708)$	$(\pm 0.709)$	$(\pm 0.186)$	$(\pm 0.606)$
Resolution 64 <sup>3</sup>	9.164	3.227	8.039	8.992	3.268	8.123	8.318	2.889	6.189	8.299	2.996	6.429
	$(\pm 1.114)$	$(\pm 0.369)$	$(\pm 1.103)$	$(\pm 1.068)$	$(\pm 0.344)$	$(\pm 1.047)$	$(\pm 0.962)$	$(\pm 0.365)$	$(\pm 0.950)$	$(\pm 0.893)$	$(\pm 0.325)$	$(\pm 0.828)$
Resolution 128 <sup>3</sup>	8.340	1.400	3.879	8.277	1.408	3.868	6.680	0.872	1.929	6.681	0.893	1.982
	$(\pm 1.117)$	$(\pm 0.314)$	$(\pm 1.023)$	$(\pm 1.077)$	$(\pm 0.305)$	$(\pm 1.006)$	$(\pm 0.776)$	$(\pm 0.261)$	$(\pm 0.779)$	$(\pm 0.669)$	$(\pm 0.207)$	$(\pm 0.618)$
Resolution $256^3$	7.349	0.453	1.112	7.315	0.447	1.096	6.522	0.327	0.697	6.525	0.324	0.691
	$(\pm 0.882)$	$(\pm 0.198)$	$(\pm 0.649)$	$(\pm 0.906)$	$(\pm 0.167)$	$(\pm 0.571)$	$(\pm 0.818)$	$(\pm 0.226)$	$(\pm 0.740)$	$(\pm 0.652)$	$(\pm 0.168)$	$(\pm 0.549)$

Table 1: Results of the ablation study on the proposed *DeepCSR* framework for cortical surface reconstruction from MRI using the ADNI study data [15].

on the validation set and report their performance on the test set. We emphasize that these splits do not have MRIs or subjects in common for an unbiased evaluation.

**Evaluation metrics.** We report the mean and standard deviation of well known surface comparison metrics like earth mover's distance (EMD), average absolute distance (AD), and Hausdorff distance (HD) [35, 42, 43]. EMD is the minimum amount of "work" to morph one surface to the other (lower is better), while AD and HD is the mean and maximum distance between closest points in two meshes (lower is better), respectively. Since HD is very sensitive to outliers, we use the  $90^{th}$  percentile instead of the maximum as suggested by Huttenlocher et al. [13]. It is also important to mention that AD and HD are computed in a bidirectional way for symmetry and over a sub-sample of 100k points.

**Signed distance vs. Occupancy Field.** We compare the performance of the *DeepCSR* framework when using signed distance function (first row in Table 1) and occupancy field (second row in Table 1) as implicit surface representation. The signed distance function provides better results than occupancy field since such a representation offers more information than simple binary labels to the learning process and mesh extraction method. More specifically, it also indicates the remoteness of the target surface at every point in the 3D space whereas the occupancy field just says whether a given point is inside or outside the target surface.

**Surface-based Sampling.** Despite uniform sampling has presented better results on generic object surface reconstruction as in Mescheder et al. [25], it is not the best choice for cortical surfaces due to the high frequency details localized on the surface. To support such a claim, we train the *DeepCSR* framework to regress signed distance function of points sampled uniformly in the 3D template space and compare with the proposed sampling scheme described in Section 3.1. The results for the uniform sampling approach is reported in the third row of Table 1, while our surface-

based sampling scheme is in the first row of the same table. The surface-based schema outperforms uniform sampling in all of the evaluation metrics.

**Hypercolumn features.** It is the most important component in our model, since we are not able to recover the high frequency details in the cortical surfaces without it. In order to support this claim, we remove the hypercolumn formation steps from our encoder depicted in Figure 2, reshape the encoder's last layer to output a 512-dimensional vector, use this vector as the decoder input, and train this model as before. These changes remove the hypercolumn schema but keeps the model with similar capacity since no layer or weights are removed. The performance of such a model is reported in the fourth row of Table 1 and it is much worse than the proposed model for all the evaluation metrics. Furthermore, Figure 4 depicts oversmoothed cortical surfaces that were reconstructed without hypercolumn features.

**Single Surface Model.** As previously observed by the machine learning community [36], multi-task learning tends to provide better generalization. Similarly, we observe that learning the surface representations jointly produces better results than learn multiple models independently. We demonstrate such a result by training and evaluating a network for each target surface independently and comparing them to a network trained jointly for all the target surfaces. The results for the single-surface approach is presented in the fifth row of Table 1, while the multi-surface is presented in the first row of the same table.

**Output Resolution.** We reconstruct cortical surfaces at different resolutions by predicting implicit surface representations for 3D evenly spaced grids of points of different sizes as explained in Section 3.3. Figure 4 depicts cortical surfaces generated at  $64^3$ ,  $128^3$ ,  $256^3$ , and  $512^3$  3D evenly spaced grids of points. As the resolution grows, our method reconstructs more details of the cortical surfaces. The quantitative results are presented in the last three rows of Table 1.



Figure 4: Reconstructed outer cortical surfaces at different resolution or without hypercolumns ("No hc.").

#### 4.2. Comparison to FreeSurfer and FastSurfer

We now compare the precision, accuracy, and runtime of the proposed model to the FreeSurfer V6 (cross-sectional pipeline) and FastSurfer. More specifically, we use the *DeepCSR* (SDF) trained on the ADNI dataset and reconstruct surfaces at  $512^3$  resolution as explained Section 3.3. Table 2 summarizes the results and Figure 5 shows examples of reconstructed surfaces.

Precision Analysis. We measure the precision, i.e., repeatability, of the evaluated frameworks using the Test-Retest dataset (TRT) [22] and the experimental protocol described by Tosun et al. [43]. More specifically, the TRT dataset provides 120 T1-weighted MRI scans from 3 subjects which are scanned twice in 20 sessions spanning 31 days. Since no morphological changes should occur between successive scans of the same subject at the same session, the reconstructed surfaces should be identical up to the variations in the image acquisition and minor physiological changes. Therefore, following the aforementioned protocol, we reconstruct the surfaces from every MRI scan using the evaluated frameworks, align pairs of surfaces from the same subject and session using the ICP algorithm, compute discrepancy measures between these pairs of aligned surfaces, and report group statistics of these metrics. As discrepancy measures, we use the already discussed average absolute distance (AD), in millimeters (mm), between meshes and the percentage of distances greater than one (% > 1 mm) and two ( $\% > 2 \ mm$ ) millimeters as group statistics. The "Precision" column in Table 2 summarizes the results for this experiment. We observe that the proposed *DeepCSR* has better reproducibility than both FreeSurfer and Fast-Surfer which is critical for medical studies.

Accuracy Analysis. This experiment aims to measure how close the reconstructed cortical surfaces are to the true cortical surfaces. Since there is no manually annotated surface data for such a task, we decide to compare the evaluated methods on the segmentation of the brain cortex using the Multi-Atlas Labelling Challenge (MALC) dataset [19, 23].

This dataset consists of 30 brain volumes manually segmented by experts using the NeuroMorphometric labelling schema for the whole brain including cortical and subcortical structures. We first select all cortical labels to form the brain cortex ground-truth segmentation. Then, for a given MR image, we reconstruct the inner and outer surfaces using the evaluated methods, discretize these surfaces to a grid of voxels of  $1 mm^3$ , and resample these generated voxelgrids to the input MRI resolution and dimensions. Next, we remove the resampled voxel-grid of the inner surface from the interior of the resampled voxel-grid of the outer surface and perform morphological opening and dilatation to generate a surface-based brain cortex segmentation. Finally, we report the Dice score (Dice) and volume similarity (VS) [42] between the surface-based generated segmentation and the cortex ground-truth segmentation as an accuracy measure of the evaluated frameworks. The "Accuracy" column in Table 2 summarizes the results for this experiment. We observe that *DeepCSR* provides brain cortex segmentation with slightly greater overlap and more similar volume to the manually annotated data than the competitors.

Runtime Analysis. In order to compare the processing time required by the evaluated frameworks, we report the average elapsed time, in minutes, for these frameworks to reconstruct the cortical surfaces of the MRI scans in the MALC dataset [19, 23]. The "Runtime" column in Table 2 presents the results for this experiment. It also important to note that the FreeSufer and FastSurfer runtimes reported just takes into account the processing steps necessary to reconstruct the cortical surfaces and ignores any other computations for a fair comparison. In summary, using a NVIDIA P100 GPU and Intel Xeon (E5-2690) CPU, our model is at least thirteen times faster than the FreeSurfer using the same hardware and input MRI scans. When compared to the Fast-Surfer, the speed-up is much smaller but the variance in runtime is drastically reduced. These improvements facilitate large medical studies and new healthcare applications.

Moreover, profiling *DeepCSR*, we observe that the initial registration takes  $2.334 (\pm 0.022)$  minutes, the implicit sur-

Table 2: Precision, accuracy and runtime comparison between FreeSurfer, FastSurfer, and *DeepCSR* on cortical reconstruction from MRI using Test-Retest (TRT) and MALC datasets.

	P	Precision on TR	T	Accuracy	Runtime	
Method	AD(mm)	$\% > 1 \ mm$	% > 2mm	Dice	VS	(minutes)
FreeSurfer	0.241	9 479	0.083	0.841	0.953	373.86
	$(\pm 0.291)$	2.412	0.985	$(\pm 0.020)$	$(\pm 0.027)$	$(\pm 47.64)$
FastSurfer	0.204	1 402	0.274	0.834	0.942	28.943
	$(\pm 0.028)$	1.492	0.374	$(\pm 0.021)$	$(\pm 0.029)$	$(\pm 13.281)$
DeepCSR	0.193	1 966	0.962	0.846	0.958	27.824
	$(\pm 0.051)$	1.200	0.205	$(\pm 0.019)$	$(\pm 0.024)$	$(\pm 1.393)$



Figure 5: Example of outer and inner cortical surfaces reconstructed with *DeepCSR*. The surfaces are color coded with the absolute distance to the pseudo ground-truth surface.

face prediction takes 8.672 ( $\pm 0.655$ ) minutes, the topology correction takes 16.493 ( $\pm 1.064$ ) minutes, and the marching cubes takes 0.325 ( $\pm 0.025$ ) minutes. Therefore, further speed-up can be achieved by employing a faster topology correction method.

# 5. Conclusion

In this paper, we tackle the problem of directly reconstructing the brain cortex from MRI which is a critical task in clinical studies of neurodegenerative diseases. We formulate this problem as the prediction of implicit surfaces for points in a continuous brain template coordinate system. We also develop an encoder network architecture with hypercolumn features that is able to extract local and global image features from brain MR Images. Due to the continuous nature of this formulation and the efficient design of our network, we are able to accurately reconstruct cortical surfaces at high resolution capturing the brain surface geometry. Compared to the widely used FreeSurfer toolbox and its deep learning powered variant FastSurfer in two standard datasets, the proposed *DeepCSR* was found to be as accurate, more precise, and faster, which should facilitate large-scale medical studies and new healthcare applications.

# References

- [1] Liana G Apostolova, Shannon L Risacher, Tugce Duran, Eddie C Stage, Naira Goukasian, John D West, Triet M Do, Jonathan Grotts, Holly Wilhalme, Kwangsik Nho, et al. Associations of the top 20 alzheimer disease risk variants with brain amyloidosis. *JAMA Neurology*, 75(3):328–341, 2018.
- [2] Miguel Angel Gonzalez Ballester, Andrew P Zisserman, and Michael Brady. Estimation of the partial volume effect in mri. *Medical Image Analysis*, 6(4):389–405, 2002.
- [3] Pierre-Louis Bazin and Dzung L Pham. Topology correction of segmented medical images using a fast marching algorithm. *Computer methods and programs in biomedicine*, 88(2):182–190, 2007.
- [4] Robert Dahnke, Rachel Aine Yotter, and Christian Gaser. Cortical thickness and central surface estimation. *NeuroIm-age*, 65:336–348, 2013.
- [5] Anders M Dale, Bruce Fischl, and Martin I Sereno. Cortical surface-based analysis: I. segmentation and surface reconstruction. *NeuroImage*, 9(2):179–194, 1999.
- [6] Vincent Doré, Victor L Villemagne, Pierrick Bourgeat, Jurgen Fripp, Oscar Acosta, Gael Chetélat, Luping Zhou, Ralph Martins, Kathryn A Ellis, Colin L Masters, et al. Crosssectional and longitudinal analysis of the relationship between aβ deposition, cortical thickness, and memory in cognitively unimpaired individuals and in alzheimer disease. JAMA Neurology, 70(7):903–911, 2013.
- [7] Bruce Fischl, Martin I Sereno, and Anders M Dale. Cortical surface-based analysis: Ii: inflation, flattening, and a surfacebased coordinate system. *NeuroImage*, 9(2):195–207, 1999.
- [8] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [9] Xiao Han, Dzung L Pham, Duygu Tosun, Maryam E Rettmann, Chenyang Xu, and Jerry L Prince. Cruise: cortical reconstruction using implicit surface evolution. *NeuroImage*, 23(3):997–1012, 2004.
- [10] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 447–456, 2015.
- [11] Leonie Henschel, Sailesh Conjeti, Santiago Estrada, Kersten Diers, Bruce Fischl, and Martin Reuter. Fastsurfer-a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage*, page 117012, 2020.
- [12] Derek Hoiem, Alexei A Efros, and Martial Hebert. Recovering surface layout from an image. *International Journal of Computer Vision (IJCV)*, 75(1):151–172, 2007.
- [13] Daniel P Huttenlocher, Gregory A Klanderman, and William J Rucklidge. Comparing images using the hausdorff distance. *The IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 15(9):850–863, 1993.
- [14] Zafer Iscan, Tony B Jin, Alexandria Kendrick, Bryan Szeglin, Hanzhang Lu, Madhukar Trivedi, Maurizio Fava, Patrick J McGrath, Myrna Weissman, Benji T Kurian, et al. Test–retest reliability of freesurfer measurements within and

between sites: effects of visual approval process. *Human brain mapping*, 36(9):3472–3485, 2015.

- [15] Clifford R Jack Jr, Matt A Bernstein, Nick C Fox, Paul Thompson, Gene Alexander, Danielle Harvey, Bret Borowski, Paula J Britson, Jennifer L. Whitwell, Chadwick Ward, et al. The alzheimer's disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging*, 27(4):685–691, 2008.
- [16] June Sic Kim, Vivek Singh, Jun Ki Lee, Jason Lerch, Yasser Ad-Dab'bagh, David MacDonald, Jong Min Lee, Sun I Kim, and Alan C Evans. 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.
- [17] Chen Kong, Chen-Hsuan Lin, and Simon Lucey. Using locally corresponding cad models for dense 3d reconstructions from a single image. In *The Conference on Computer Vision* and Pattern Recognition (CVPR), pages 4857–4865, 2017.
- [18] Nikolaus Kriegeskorte and Rainer Goebel. An efficient algorithm for topologically correct segmentation of the cortical sheet in anatomical mr volumes. *NeuroImage*, 14(2):329– 346, 2001.
- [19] B Landman and S Warfield. Miccai 2012 workshop on multiatlas labeling. In *Medical image computing and computer* assisted intervention conference, 2012.
- [20] Thomas Lewiner, Hélio Lopes, Antônio Wilson Vieira, and Geovan Tavares. Efficient implementation of marching cubes' cases with topological guarantees. *Journal of Graphics Tools*, 8(2):1–15, 2003.
- [21] David MacDonald, Noor Kabani, David Avis, and Alan C Evans. Automated 3-d extraction of inner and outer surfaces of cerebral cortex from mri. *NeuroImage*, 12(3):340–356, 2000.
- [22] Julian Maclaren, Zhaoying Han, Sjoerd B Vos, Nancy Fischbein, and Roland Bammer. Reliability of brain volume measurements: A test-retest dataset. *Scientific Data*, 1:140037, 2014.
- [23] Daniel S Marcus, Tracy H Wang, Jamie Parker, John G Csernansky, John C Morris, and Randy L Buckner. Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults. *Journal of Cognitive Neuroscience*, 19(9):1498– 1507, 2007.
- [24] John C Mazziotta, Arthur W Toga, Alan Evans, Peter Fox, Jack Lancaster, et al. A probabilistic atlas of the human brain: theory and rationale for its development. *NeuroIm*age, 2(2):89–101, 1995.
- [25] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4460–4470, 2019.
- [26] Mateusz Michalkiewicz, Jhony K. Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Implicit surface representations as layers in neural networks. In *The International Conference on Computer Vision (ICCV)*, October 2019.
- [27] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying

deep learning and model based human pose and shape estimation. In *International Conference on 3D Vision (3DV)*, pages 484–494. IEEE, 2018.

- [28] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [29] Ethan Perez, Harm de Vries, Florian Strub, Vincent Dumoulin, and Aaron Courville. Learning visual reasoning without strong priors. arXiv preprint arXiv:1707.03017, 2017.
- [30] Dzung L Pham, Pierre-Louis Bazin, and Jerry L Prince. Digital topology in brain imaging. *IEEE Signal Processing Magazine*, 27(4):51–59, 2010.
- [31] Jhony K Pontes, Chen Kong, Sridha Sridharan, Simon Lucey, Anders Eriksson, and Clinton Fookes. Image2mesh: A learning framework for single image 3d reconstruction. In *The Asian Conference on Computer Vision (ACCV)*, pages 365–381. Springer, 2018.
- [32] Olivier Querbes, Florent Aubry, Jérémie Pariente, Jean-Albert Lotterie, Jean-François Démonet, Véronique Duret, Michèle Puel, Isabelle Berry, Jean-Claude Fort, Pierre Celsis, and The Alzheimer's Disease Neuroimaging Initiative. Early diagnosis of Alzheimer's disease using cortical thickness: impact of cognitive reserve. *Brain*, 132(8):2036–2047, 05 2009.
- [33] Michael Rebsamen, Yannick Suter, Roland Wiest, Mauricio Reyes, and Christian Rummel. Brain morphometry estimation: From hours to seconds using deep learning. *Frontiers in neurology*, 11:244, 2020.
- [34] Claudia E Rodriguez-Carranza, Pratik Mukherjee, D Vigneron, J Barkovich, and Colin Studholme. A framework for in vivo quantification of regional brain folding in premature neonates. *NeuroImage*, 41(2):462–478, 2008.
- [35] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. A metric for distributions with applications to image databases. In *The International Conference on Computer Vision (ICCV)*, pages 59–66. IEEE, 1998.
- [36] Sebastian Ruder. An overview of multi-task learning in deep neural networks. ArXiv, abs/1706.05098, 2017.
- [37] Rodrigo Santa Cruz, Léo Lebrat, Pierrick Bourgeat, Vincent Doré, Jason Dowling, Jurgen Fripp, Clinton Fookes, and Olivier Salvado. Going deeper with brain morphometry using neural networks. arXiv e-prints, pages arXiv–2009, 2020.
- [38] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Make3d: Learning 3d scene structure from a single still image. *The IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 31(5):824–840, 2008.
- [39] Marie Schaer, Meritxell Bach Cuadra, Lucas Tamarit, François Lazeyras, Stephan Eliez, and Jean-Philippe Thiran. A surface-based approach to quantify local cortical gyrification. *IEEE Transactions on Medical Imaging*, 27(2):161– 170, 2008.
- [40] David W Shattuck and Richard M Leahy. Brainsuite: an automated cortical surface identification tool. *Medical Image Analysis*, 6(2):129–142, 2002.
- [41] Gianfranco Spalletta, Fabrizio Piras, and Tommaso Gili.

Brain morphometry. Springer, 2018.

- [42] Abdel Aziz Taha and Allan Hanbury. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC medical imaging*, 15(1):29, 2015.
- [43] Duygu Tosun, Maryam E Rettmann, Daniel Q Naiman, Susan M Resnick, Michael A Kraut, and Jerry L Prince. Cortical reconstruction using implicit surface evolution: accuracy and precision analysis. *NeuroImage*, 29(3):838–852, 2006.
- [44] Ingo Wald, Heiko Friedrich, Gerd Marmitt, Philipp Slusallek, and H-P Seidel. Faster isosurface ray tracing using implicit kd-trees. *IEEE Transactions on Visualization* and Computer Graphics, 11(5):562–572, 2005.
- [45] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *The European Conference* on Computer Vision (ECCV), pages 52–67, 2018.
- [46] Eric W Weisstein. Triangle point picking. Wolfram Research, Inc., 1999.
- [47] Chao Wen, Yinda Zhang, Zhuwen Li, and Yanwei Fu. Pixel2mesh++: Multi-view 3d mesh generation via deformation. In *The Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1042–1051, 2019.