

Vox2Cortex: Fast Explicit Reconstruction of Cortical Surfaces from 3D MRI Scans with Geometric Deep Neural Networks

Fabian Bongratz^{1†*}, Anne-Marie Rickmann^{1,2*}, Sebastian Pölsterl², Christian Wachinger^{1,2}
¹Technical University of Munich, ²Ludwig-Maximilians-University Munich

Abstract

The reconstruction of cortical surfaces from brain magnetic resonance imaging (MRI) scans is essential for quantitative analyses of cortical thickness and sulcal morphology. Although traditional and deep learning-based algorithmic pipelines exist for this purpose, they have two major drawbacks: lengthy runtimes of multiple hours (traditional) or intricate post-processing, such as mesh extraction and topology correction (deep learning-based). In this work, we address both of these issues and propose Vox2Cortex, a deep learning-based algorithm that directly yields topologically correct, three-dimensional meshes of the boundaries of the cortex. Vox2Cortex leverages convolutional and graph convolutional neural networks to deform an initial template to the densely folded geometry of the cortex represented by an input MRI scan. We show in extensive experiments on three brain MRI datasets that our meshes are as accurate as the ones reconstructed by state-of-the-art methods in the field, without the need for time- and resource-intensive post-processing. To accurately reconstruct the tightly folded cortex, we work with meshes containing about 168,000 vertices at test time, scaling deep explicit reconstruction methods to a new level.

1. Introduction

Over the last years, fully-convolutional neural networks (F-CNNs) have been introduced for the automatic whole-brain segmentation of MRI scans. Architecture choices range from multi-view 2D F-CNNs [17, 43] and patch-based 3D F-CNNs [18] to full-size 3D F-CNNs [41]. In contrast to time-consuming traditional methods for brain analysis, such networks perform the segmentation in seconds. The resulting voxel-based segmentation can be adequately used for volume measurements of the whole brain or subcortical structures. However, the cerebral cortex, a thin and highly-folded sheet of neural tissue of the brain,




Figure 1. Coronal slices of a brain MRI scan with overlays of mesh- and voxel-based segmentation, and 3D renderings of corresponding pial and white matter (WM) surfaces. The top row shows meshes generated by our method, the bottom row shows meshes generated by applying marching cubes on the voxel segmentation. Marching cubes meshes clearly show stair-case artifacts.

is not well captured by voxel-based segmentations since it is just a few voxels thick. Instead, cortical surfaces have been reconstructed as triangular meshes [9] to capture the intrinsic two-dimensional structure, which enables accurate measurements like thickness, volume, and gyrification. Such *in-vivo* measurements have not only been important for studying higher-level processes like consciousness, thought, emotion, reasoning, language, and memory [39, 44], but also for detecting cortical atrophy in numerous brain disorders like Alzheimer’s disease [27] and schizophrenia [25].

The current standard for cortical surface reconstruction is FreeSurfer [5, 9], which produces smooth, accurate, and topologically correct surface meshes but runs for several hours per scan. On the other hand, voxel-based segmentation algorithms are quick and their outputs can be used to obtain mesh-based surface representations using algorithms like marching cubes [28, 29]. However, the extracted meshes suffer from typical stair-case artifacts and their resolution is limited by the resolution of the segmentation, respectively the resolution of the initial MRI scan, see Fig-

*: Equal contribution, †: Corr. author (fabo.bongratz@tum.de)

Figure 1. Further, the predicted segmentation might have topological defects and a topology-correction algorithm needs to be applied before or after marching cubes mesh extraction, see Figure 2. This need for post-processing steps is not only cumbersome; while topology correction algorithms generate meshes with the desired topology, they do not always produce anatomically correct meshes [9], see also supplementary Figure 4.

Recently, seminal research in the field of single-view or multi-view reconstruction [50, 52] introduced network architectures that can produce explicit surface representations from 2D images by transforming a template mesh (e.g., a sphere) using graph convolutions. These methods are promising for our task, as they can enforce the desired topology of the template mesh and directly output a mesh without the need for post-processing. So far, these methods have not been applied to reconstruct shapes with such complex folding patterns as the cerebral cortex.

In this work, we introduce Vox2Cortex, a deep learning-based algorithm that allows for *direct* reconstruction of *explicit* meshes of cortical surfaces from brain MRI. Vox2Cortex simultaneously predicts white matter (WM) and pial surfaces of both hemispheres, resulting in an output of four meshes. Vox2Cortex enforces a spherical topology to prevent the formation of holes or handles in the mesh and does therefore not need post-processing like [4]. We use up to 168,000 vertices per mesh to capture the highly complex folding patterns of the cortex with high accuracy. This number is substantially higher than in prior work on single-view reconstruction that typically uses less than 10,000 vertices per surface. Our main contributions are:

- Vox2Cortex is the first network specifically designed to extract explicit surface representations of the cortex from MRI scans using a combination of a convolutional neural network and a graph neural network.
- We ensure spherical topology of the generated meshes by learning to deform a template mesh with fixed topology and arbitrary resolution.
- Vox2Cortex models the interdependency between WM and pial surfaces by exchanging information between them.
- We propose a new loss function for training explicit reconstruction methods that leverages local curvature to weight Chamfer distances. It turns out that this loss is crucial for learning to reconstruct the tightly folded surfaces of the human cortex.
- We demonstrate on multiple brain datasets that Vox2Cortex performs on par with or better than existing implicit and voxel-based reconstruction methods in terms of accuracy and consistency while being 25 times faster at inference time.

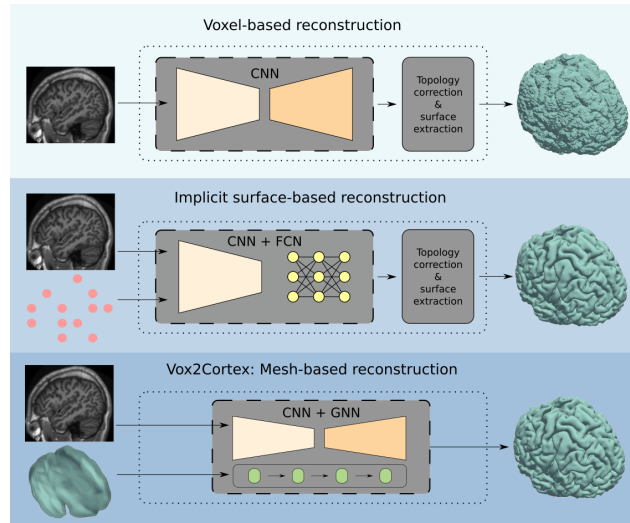


Figure 2. Existing deep learning-based approaches for cortical surface reconstruction from MRI usually rely on an implicit or voxel representation of the cortex. Those methods require intricate post-processing steps, including topological correction and marching cubes [28, 29] to generate surface meshes. These meshes are, however, crucial for downstream applications such as the measurement of cortical thickness. In contrast, our model directly yields highly accurate meshes of the WM and the pial surfaces.

The main objective of our work is to develop a fast and highly accurate reconstruction of the cortex, which has high clinical value. Nevertheless, our contributions to predicting multiple, tightly folded meshes are generic and can also be applied to other applications. Our code is publicly available at <https://github.com/ai-med/Vox2Cortex>.

2. Related Work

Traditional brain MRI processing pipelines [6, 10, 47] consist of many steps including registration, segmentation, and cortical surface extraction. These pipelines are computationally intensive and, therefore, delay the availability of cortex measures after scan acquisition. Current deep learning approaches for cortical surface reconstruction focus on voxel-based or implicit surface reconstruction methods. Instead, we propose a mesh-based approach. We depict an overview of the different approaches in Figure 2 and review related work in the following.

Voxel-based surface reconstruction: FastSurfer [17] is a deep learning approach that speeds up the FreeSurfer pipeline but focuses on voxel-based segmentation and requires the use of marching cubes for surface generation and topology correction. SegRecon [13] proposes a 3D CNN for joint segmentation and surface reconstruction, which learns a 3D signed distance function (SDF) and therefore still requires marching cubes and topology correction.

Deep implicit representations: Deep implicit represen-

tations have become a popular research field in 3D computer vision in recent years [33, 38, 45, 54]. The idea behind this is to learn a function that maps 3D coordinates to a continuous implicit representation of shape, usually an SDF. DeepCSR [4] was explicitly proposed for cortical surface reconstruction. As their shape representation is continuous, meshes of any desired resolution can be obtained. DeepCSR achieves state-of-the-art results in terms of distance metrics but still requires marching cubes and topology correction.

Mesh-based surface reconstruction: Mesh-deformation networks [23, 24, 50, 52, 53] are learning deformable mesh models that take as input a template mesh or sphere initialization and the image and iteratively deform the mesh by learning the deformation field of the vertices. Voxel2Mesh [53] and MeshDeformNet [23, 24] were applied to medical data and showed promising results, albeit mainly applied to relatively simple shapes like the hippocampus, liver, or heart. As the cortex is a significantly more complicated shape, it is still to be evaluated if this approach can work well on cortex reconstruction. PialNN [30] learns to deform an initial WM surface to a pial surface by a series of graph convolutions but requires FreeSurfer to create the WM surface.

3. Method

In this section, we provide a detailed description of Vox2Cortex, a method for fast cortical surface reconstruction from 3D magnetic resonance (MR) images based on geometric deep learning.

3.1. Vox2Cortex Architecture

Vox2Cortex takes a 3D brain MRI scan and a mesh template as input and computes simultaneously four cortical surfaces, namely the white-matter and the pial surfaces of each hemisphere. As a side product, the network predicts a voxel-wise brain segmentation (i.e., of the gray matter and the tissue enclosed by the gray matter). Inspired by previous related methods [23, 24, 49, 53], our architecture consists of two neural sub-networks, a convolutional neural network (CNN) that operates on voxels and a graph neural network (GNN) responsible for mesh deformation. Both networks are connected via feature-sampling modules that map features extracted by the CNN to vertex locations of the meshes. Figure 3 depicts the whole architecture together with exemplary in- and outputs. In the following, we detail the individual building blocks: image encoder, image decoder, mesh deformation network, and information exchange.

Image encoding and decoding (CNN) We use a residual UNet architecture [3, 22, 42, 56] for image-feature extraction. This encoder-decoder F-CNN solely takes the 3D

brain scan as input and segments it into brain and background voxels. The UNet consist of multiple residual convolution blocks with batch-normalization layers [20] and ReLU activations [36]. For up- and downsampling of the feature maps, we use (transposed) convolutional layers with stride 2. Moreover, we add deep-supervision branches [55] to propagate the segmentation loss directly to lower decoder layers. Figure 3 illustrates the network in detail.

Mesh deformation (GNN) The second major building block of our architecture is a GNN that takes a template mesh as input [23, 24, 49, 53]. Its precedence over an MLP [51, 14] is validated by our ablation study (cf. Table 1). The GNN deforms the template in four mesh-deformation steps that build upon each other, see Figure 3. Each deformation step predicts vertex-specific displacement vectors relative to the mesh produced by the previous step. Similar to the image-processing CNN, our GNN sub-architecture consists of multiple residual blocks, allowing for effective residuum-based learning [16, 24]. Each graph-residual block consists of three graph convolutions with subsequent batch-norm and ReLU operations. The input residuum is added before the last ReLU operation and potentially reshaped with nearest-neighbor interpolation. Following [24], the initial graph-residual block has a large number of channels compared to later blocks to weight the extracted vertex features similar to the image-based features in terms of feature-vector length. Ultimately, a single graph convolutional layer is used to output the displacement vectors for each vertex at the current deformation stage.

We leverage spectrum-free graph convolutions, in the flavor of message-passing operations [2], and choose the implementation provided in [40]. Formally, each graph convolutional layer transforms the features from a previous layer $\mathbf{f}_i \in \mathbb{R}^{d_{in}}$ of a vertex $\mathbf{v}_i \in \mathcal{V}$ by aggregating

$$\mathbf{f}'_i = \frac{1}{1 + |\mathcal{N}(i)|} \left[\mathbf{W}_0 \mathbf{f}_i + \mathbf{b}_0 + \sum_{j \in \mathcal{N}(i)} (\mathbf{W}_1 \mathbf{f}_j + \mathbf{b}_1) \right], \quad (1)$$

where $\mathbf{W}_0, \mathbf{W}_1 \in \mathbb{R}^{d_{out} \times d_{in}}$ together with $\mathbf{b}_0, \mathbf{b}_1 \in \mathbb{R}^{d_{out}}$ represent linear transformations and $\mathcal{N}(i)$ is the set of neighbors of \mathbf{v}_i . Besides the mesh-output layers, graph convolutional layers are followed by batch-norm and ReLU layers in our network.

Image features While CNNs and GNNs are established network architectures, so far there have only been few attempts for combining them [49, 53, 24, 23] and the question regarding the best way to exchange information between them has not been answered yet. We combine feature maps from the UNet encoder as well as the UNet decoder at multiple resolutions. This allows the network to identify

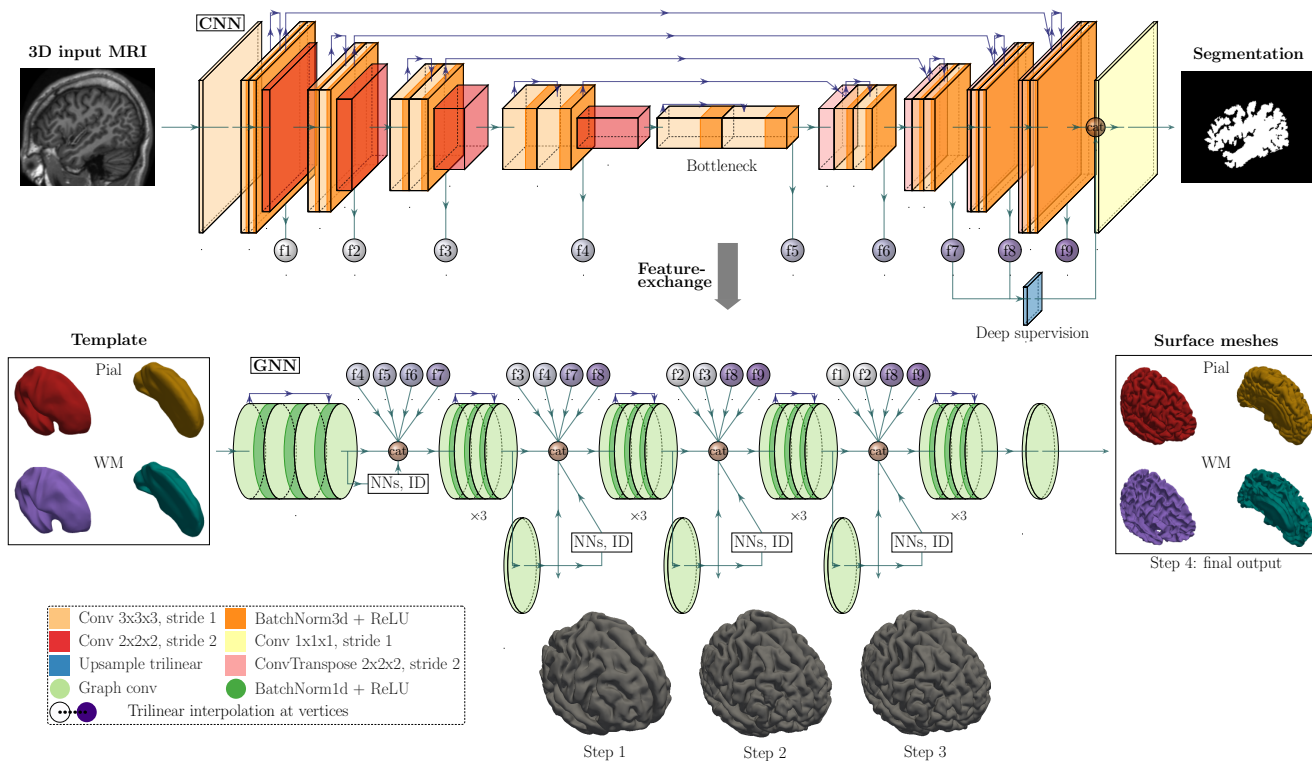


Figure 3. Vox2Cortex pipeline. The architecture takes as input a 3D brain scan and a mesh template and predicts a voxel-wise segmentation map and cortical surface meshes. The principal building blocks are a CNN and a GNN. The GNN deforms the initial template within four deformation steps based on image- and shape-descriptive features to the predicted output meshes. We created the drawing using [21].

relevant information during the training process, which is not possible with a narrow pre-selection of feature maps. In order to get features in continuous 3D space, we interpolate the feature maps from discrete voxels trilinearly.

Inter-mesh neighbors To the best of our knowledge, deep explicit reconstruction methods have been only applied for the reconstruction of *isolated* objects so far. However, we deal with multiple surfaces, among which we have an *interdependence* as the inner and outer brain surfaces are always aligned. To improve the quality of the reconstruction, we model the interdependence of the meshes by exchanging information between them. To this end, we concatenate the coordinates of the five nearest outer vertices and a surface identifier to the features of each inner vertex (and vice versa) before every mesh-deformation block. Traditional methods like FreeSurfer [6] also model this interdependency by extracting an outer surface based on the inner one (and not simultaneously as we do).

Mesh template Existing explicit reconstruction methods usually start from simple mesh templates like spheres or ellipsoids. To improve the cortex reconstruction quality, we propose using a more application-specific template.

We created four genus-0 templates by choosing a random FreeSurfer mesh and applying Laplacian smoothing [48] until the surface did not change anymore. Figure 3 illustrates the template as input to the network.

Importantly, we observed that it is possible to use a template with fewer vertices during training compared to testing and thereby increase the surface accuracy of a trained model. In other words, we can choose the desired mesh resolution independently of the template used during training. In our experiments, we used about 42,000 vertices per surface during training, whereas we increase this number to about 168,000(!) during testing. Across all four surfaces, this yields a total over 672,000 vertices.

3.2. Loss Functions

In this section, we describe the loss function for training the Vox2Cortex network with particular emphasis on our novel *curvature-weighted Chamfer loss*. Definitions of the other loss functions are in the supplementary material.

Let $y^p = \{\{\mathcal{M}_s^p, \Delta_s^p, B_l^p\} | s = 1, \dots, S; l = 1, \dots, L\}$ be the prediction of our model, where $\mathcal{M}_s^p = (\mathcal{V}_s^p, \mathcal{F}_s^p, \mathcal{E}_s^p)$ are the predicted meshes at S different deformation stages, Δ_s^p are the predicted displacement vectors stacked into a tensor, and $B_l^p \in [0, 1]^{HWD}$ are the voxel-wise binary-segmentation maps (final prediction of the CNN and $L - 1$

deep-supervision outputs). In general, each mesh \mathcal{M}_s^p consists again of C different surfaces, i.e., $\mathcal{M}_s^p = \{\mathcal{M}_{s,c}^p | c = 1, \dots, C\}$. In our case, C is equal to four (WM and pial surfaces for each hemisphere). As we train our model in a supervised manner, we assume having corresponding ground-truth meshes and segmentation maps $y^{\text{gt}} = (\mathcal{M}^{\text{gt}}, B^{\text{gt}})$ available for each sample in the training set.

Total loss The loss function of Vox2Cortex consists of a voxel \mathcal{L}_{vox} and a mesh $\mathcal{L}_{\text{mesh}}$ part

$$\mathcal{L}(y^p, y^{\text{gt}}) = \mathcal{L}_{\text{vox}}(y^p, y^{\text{gt}}) + \mathcal{L}_{\text{mesh}}(y^p, y^{\text{gt}}). \quad (2)$$

Voxel loss Let $\mathcal{L}_{\text{BCE}}(B_l^p, B^{\text{gt}})$ be the binary cross-entropy loss between a predicted segmentation map $B_l^p \in [0, 1]^{HWD}$ and a label $B^{\text{gt}} \in \{0, 1\}^{HWD}$. Taking into account that we have L segmentation outputs (the final segmentation of the CNN as well as $L - 1$ deep-supervision outputs), we compute the voxel loss as

$$\mathcal{L}_{\text{vox}}(y^p, y^{\text{gt}}) = \sum_{l=1}^L \mathcal{L}_{\text{BCE}}(B_l^p, B^{\text{gt}}). \quad (3)$$

Mesh loss Compared to the voxel loss, it is more challenging to define an adequate loss function between two meshes. This is partly due to the lack of point correspondences between the predicted meshes and the labels, making it complicated to define a loss function that enforces a “good” cortical surface mesh. Inspired by previous explicit reconstruction methods [24, 49, 53], we use a combination of geometry-consistency and regularization losses

$$\mathcal{L}_{\text{mesh}}(y^p, y^{\text{gt}}) = \mathcal{L}_{\text{mesh, cons}}(y^p, y^{\text{gt}}) + \mathcal{L}_{\text{mesh, reg}}(y^p). \quad (4)$$

For the geometry-consistency loss, we add up a novel curvature-weighted Chamfer loss \mathcal{L}_C , further described in the next section, and an inter-mesh normal consistency (also called normal distance [12]) $\mathcal{L}_{n, \text{inter}}$

$$\begin{aligned} \mathcal{L}_{\text{mesh, cons}}(y^p, y^{\text{gt}}) &= \sum_{s=1}^S \sum_{c=1}^C [\lambda_{1,c} \mathcal{L}_C(\mathcal{M}_{s,c}^p, \mathcal{M}_c^{\text{gt}}) \\ &+ \lambda_{2,c} \mathcal{L}_{n, \text{inter}}(\mathcal{M}_{s,c}^p, \mathcal{M}_c^{\text{gt}})]. \end{aligned} \quad (5)$$

The regularization loss consists of Laplacian smoothing of the displacement fields $\mathcal{L}_{\text{Lap, rel}}$, intra-mesh normal consistency $\mathcal{L}_{n, \text{intra}}$, and edge length $\mathcal{L}_{\text{edge}}$

$$\begin{aligned} \mathcal{L}_{\text{mesh, reg}}(y^p) &= \sum_{s=1}^S \sum_{c=1}^C [\lambda_{3,c} \mathcal{L}_{\text{Lap, rel}}(\mathcal{M}_{s,c}^p, \Delta_{s,c}^p) \\ &+ \lambda_{4,c} \mathcal{L}_{n, \text{intra}}(\mathcal{M}_{s,c}^p) \\ &+ \lambda_{5,c} \mathcal{L}_{\text{edge}}(\mathcal{M}_{s,c}^p)]. \end{aligned} \quad (6)$$




Figure 4. Our curvature-weighted Chamfer loss (green) often leads to a more accurate model of the cortical folds compared to the standard Chamfer distance (blue). FreeSurfer ground truth (white).

More details about the loss hyperparameters can be found in the supplementary material.

Curvature-weighted Chamfer loss The Chamfer distance has become an essential building block for learning deformable-shape models of surfaces [8, 49], but needs to be combined with regularization terms, as it can otherwise lead to edge intersections [15]. In regions with high curvature, as often found in the densely folded geometry of the cortex, the smoothing effect of the regularization terms can lead to lower geometric accuracy. Hence, we propose a curvature-weighted Chamfer loss function that emphasizes high-curvature regions and improves the reconstruction of densely folded areas, cf. Figure 4. A proof showing that this loss “pushes” predicted points in high-curvature regions more towards their correct position compared to points in low-curvature regions, under mild assumptions, is in the supplementary material. Note that we only use curvature weights from ground-truth points as the potentially wrong curvature in a prediction could be misleading.

Assume we have a curvature function $\kappa(\mathbf{p}) \in \mathbb{R}_{\geq 0}$ assigning a curvature to a point \mathbf{p} , then we define the curvature-weighted-Chamfer loss as

$$\begin{aligned} \mathcal{L}_C(\mathcal{M}_{s,c}^p, \mathcal{M}_c^{\text{gt}}) &= \frac{1}{|\mathcal{P}_c^{\text{gt}}|} \sum_{\mathbf{u} \in \mathcal{P}_c^{\text{gt}}} \kappa(\mathbf{u}) \min_{\mathbf{v} \in \mathcal{P}_{s,c}^p} \|\mathbf{u} - \mathbf{v}\|^2 \\ &+ \frac{1}{|\mathcal{P}_{s,c}^p|} \sum_{\mathbf{v} \in \mathcal{P}_{s,c}^p} \kappa(\tilde{\mathbf{u}}) \min_{\mathbf{u} \in \mathcal{P}_c^{\text{gt}}} \|\mathbf{v} - \mathbf{u}\|^2, \end{aligned} \quad (7)$$

where $\tilde{\mathbf{u}} = \arg \min_{\mathbf{r} \in \mathcal{P}_c^{\text{gt}}} \|\mathbf{v} - \mathbf{r}\|^2$. In practice, we found that

$$\kappa(\mathbf{p}) = \min\{1 + \bar{\kappa}(\mathbf{p}), \kappa_{\text{max}}\}, \quad (8)$$

where $\bar{\kappa}(\mathbf{p})$ is the discrete mean curvature [34, 37], is a good choice ($\kappa_{\text{max}} = 5$ in our experiments). During training, every $\mathcal{P}_c^{\text{gt}} |_{c=1, \dots, C}$ contains as many vertices as the smallest ground-truth surface in the training set. On the other hand, we sample the point clouds $\mathcal{P}_c^p |_{c=1, \dots, C}$ in a differentiable manner from the surface of the predicted meshes [12, 46]

such that they contain the same number of points as the reference meshes.

4. Results

In the following, we compare Vox2Cortex to related approaches on multiple datasets. In addition, we reveal the individual contributions of the main building blocks of Vox2Cortex in an extensive ablation study.

4.1. Datasets

Pre-processing All data used in this work was processed using FreeSurfer v5.3 [9]. We used the orig.mgz files and WM and pial surfaces generated by FreeSurfer. We follow the pre-processing pipeline by [4] for experiments with DeepCSR. We registered the MRI scans to the MNI152 space using rigid and subsequent affine registration. When used as supervised training labels, the FreeSurfer meshes were simplified to about 40,000 vertices per surface using quadric edge collapse decimation [11]. We first padded all input images to have shape $192 \times 208 \times 192$ and then resized them to $128 \times 144 \times 128$ voxels. Intensity values were min-max-normalized to the range $[0, 1]$.

ADNI The Alzheimer’s Disease Neuroimaging Initiative (ADNI) (<http://adni.loni.usc.edu>) provides MRI T1 scans for subjects with Alzheimer’s Disease, Mild Cognitive Impairment, and healthy subjects. After removing data with processing artifacts, we split the data into training, validation, and testing set, balanced according to diagnosis, age, and sex. As ADNI is a longitudinal study, we only use the initial (baseline) scan for each subject. In our experiments, we use two different splits of the ADNI data. ADNI_{small} contains 299 subjects for training and 60 for validation and testing, respectively. We use it to conduct the experiments for our architecture ablation study. A second, larger split ADNI_{large} contains 1,155 subjects for training, 169 for validation, and 323 for testing. ADNI_{large} will be used to compare our model with state-of-the-art methods.

OASIS The OASIS-1 dataset [32] contains MRI T1 scans of 416 subjects. 100 subjects have been diagnosed with very mild to moderate Alzheimer’s disease. We split the data balanced on diagnosis, age, and sex, resulting in 292, 44, and 80 subjects for training, validation, and testing. We used a small subset of OASIS, the MALC dataset [26], for hyperparameter tuning as described in the supplement.

Test-retest The test-retest dataset (TRT) [31] contains 120 MRI T1 scans from 3 subjects, where each subject has been scanned twice in 20 days.

4.2. Ablation Study

We evaluate the individual design choices in Vox2Cortex and focus on characteristics of explicit image-based surface

reconstruction methods. Especially the combination of a CNN and a GNN poses the question of how to pass information in the form of features from one subnetwork to the other (in our case from the CNN to the GNN). For an encoder-decoder CNN, for example, it is possible to use extracted feature maps from the encoder [23], from the decoder [53], or from both. The impact of this and several other choices is hard to assess from existing literature and has not yet been studied extensively. We address this question in our ablation study on the ADNI_{small} subset. More precisely, we compare Vox2Cortex with the following modifications and present quantitative measures in terms of average symmetric surface distance (ASSD), called D_{23} in [7], and 90-percentile Hausdorff distance (HD) [19] in Table 1.

Voxel2Mesh: We replace our network architecture with the Voxel2Mesh network [53]. At each mesh deformation step, we sample the CNN features from the corresponding voxel-decoder stage. We double the number of voxel-decoder channels at each stage compared to our model for a fair comparison since those are the only ones passed to the GNN in Voxel2Mesh. Here, we also use learned neighborhood sampling (LNS) instead of trilinear interpolation at vertex locations as in [53]. **Encoder features:** We sample CNN features at each mesh-deformation stage from the corresponding voxel-encoder stage only, see for example [23, 24]. **Classic Chamfer:** We train our architecture with the classic Chamfer loss, which is equal to setting $\kappa(\cdot) \equiv 1$ in Equation (7). **W/o inter-mesh NNs:** The exchange of nearest-neighbor vertex positions between white and pial surfaces is omitted. **Ellipsoidal template:** Start the deformation from an ellipsoidal template instead of our smoothed cortex template. **W/o voxel decoder:** Omit the voxel decoder entirely and sample CNN features from the encoder at the respective stage, similar to [49]. **Lap. on abs. coord.:** Compute the Laplacian loss on absolute vertex coordinates instead of relative displacements, i.e., smooth the mesh instead of the displacement field. **MLP deform:** Every layer of the GNN replaced with a linear layer.

Table 1 shows that the combination of design choices in Vox2Cortex yields the best performance. The curvature-weighted Chamfer loss is the most important ingredient since reconstruction accuracy drops considerably when it is replaced by the standard Chamfer loss. It is worth noticing that whether sampling the image features from the encoder or the decoder has negligible impact on the overall performance. This could be a possible explanation why both approaches have succeeded in the past [23, 53]. Moreover, the voxel decoder can even be omitted entirely without a drastic loss in segmentation accuracy. This might be interesting for practitioners to reduce the memory demand, especially during training.

	Left WM Surface		Right WM Surface		Left Pial Surface		Right Pial Surface	
	ASSD	HD	ASSD	HD	ASSD	HD	ASSD	HD
Vox2Cortex	0.401 ± 0.065	0.894 ± 0.177	0.403 ± 0.057	0.896 ± 0.142	0.375 ± 0.055	0.965 ± 0.210	0.378 ± 0.060	1.012 ± 0.248
Vox2Cortex*	0.455 ± 0.063	1.057 ± 0.195	0.457 ± 0.056	1.055 ± 0.145	0.467 ± 0.057	1.316 ± 0.278	0.470 ± 0.0611	1.371 ± 0.281
Voxel2Mesh* [53]	0.528 ± 0.222	1.209 ± 0.732	0.528 ± 0.197	1.186 ± 0.625	0.486 ± 0.114	1.457 ± 0.398	0.476 ± 0.108	1.440 ± 0.384
Encoder features	0.453 ± 0.072	0.984 ± 0.177	0.456 ± 0.054	1.007 ± 0.144	0.432 ± 0.067	1.057 ± 0.211	0.430 ± 0.059	1.040 ± 0.174
Classic Chamfer	0.852 ± 0.081	2.175 ± 0.340	0.985 ± 0.074	2.282 ± 0.313	0.716 ± 0.063	1.906 ± 0.282	0.913 ± 0.056	2.391 ± 0.160
w/o inter-mesh NNs	0.444 ± 0.063	0.960 ± 0.174	0.438 ± 0.052	0.958 ± 0.142	0.390 ± 0.051	0.892 ± 0.146	0.396 ± 0.049	0.946 ± 0.168
Ellipsoid template	0.459 ± 0.065	0.970 ± 0.145	0.452 ± 0.071	0.954 ± 0.140	0.407 ± 0.044	0.948 ± 0.145	0.412 ± 0.053	0.983 ± 0.201
w/o voxel decoder	0.413 ± 0.069	0.914 ± 0.168	0.424 ± 0.065	0.928 ± 0.150	0.392 ± 0.057	0.916 ± 0.147	0.400 ± 0.059	0.942 ± 0.180
Lap. on abs. coord.	0.467 ± 0.075	0.958 ± 0.150	0.444 ± 0.065	0.952 ± 0.140	0.414 ± 0.050	1.102 ± 0.182	0.425 ± 0.050	1.057 ± 0.178
MLP deform	0.538 ± 0.062	1.237 ± 0.195	0.542 ± 0.057	1.228 ± 0.181	0.533 ± 0.057	1.472 ± 0.227	0.566 ± 0.055	1.447 ± 0.218

Table 1. Ablation study results in terms of average symmetric surface distance (ASSD) and 90-percentile Hausdorff distance (HD). An asterisk (*) indicates that a template with fewer vertices ($\approx 42,000$ instead of $\approx 168,000$) has been used during testing. Best values are **highlighted**. See subsection 4.2 for a detailed description of the individual variations. All values are in mm.

Data	Method	Left WM Surface		Right WM Surface		Left Pial Surface		Right Pial Surface	
		ASSD	HD	ASSD	HD	ASSD	HD	ASSD	HD
ADNI large	Vox2Cortex	0.345 ± 0.056	0.720 ± 0.125	0.347 ± 0.046	0.720 ± 0.087	0.327 ± 0.031	0.755 ± 0.102	0.318 ± 0.029	0.781 ± 0.102
	DeepCSR [4]	0.422 ± 0.058	0.852 ± 0.134	0.420 ± 0.058	0.880 ± 0.156	0.454 ± 0.059	0.927 ± 0.243	0.422 ± 0.053	0.890 ± 0.197
	nnUNet [22]	1.176 ± 0.345	1.801 ± 2.835	1.159 ± 0.242	1.739 ± 1.880	1.310 ± 0.292	3.152 ± 2.374	1.317 ± 0.312	3.295 ± 2.387
OASIS	Vox2Cortex	0.315 ± 0.039	0.680 ± 0.137	0.318 ± 0.048	0.682 ± 0.151	0.362 ± 0.036	0.894 ± 0.141	0.373 ± 0.041	0.916 ± 0.137
	DeepCSR [4]	0.360 ± 0.042	0.731 ± 0.104	0.335 ± 0.050	0.670 ± 0.195	0.458 ± 0.056	1.044 ± 0.290	0.442 ± 0.058	1.037 ± 0.294

Table 2. Comparison of Vox2Cortex with DeepCSR and nnUNet on the ADNI_{large} and OASIS datasets for the four surfaces. All values are in mm.

4.3. Comparison with Related Work

We compare Vox2Cortex with DeepCSR [4] on the ADNI_{large} and OASIS datasets. For DeepCSR, we sample points with 0.5mm distance to obtain high-resolution predictions. We chose 3D nnUNet [22] for a comparison with a voxel-based segmentation method, as it is a state-of-the-art segmentation model that has been shown to work on multiple segmentation tasks. We perform topology correction on the predicted segmentation maps, following the post-processing in [4], and use marching cubes [28] to extract meshes. Table 2 shows that nnUNet achieves less accurate predictions, which is expected as the method is solely voxel-based and the resolution of the generated meshes depends on the image resolution. Vox2Cortex achieves superior results over DeepCSR in almost all metrics on both datasets. In the supplement, we show more results on mesh topology and number of vertices for DeepCSR and Vox2Cortex.

4.4. Consistency

We analyze the consistency of Vox2Cortex, DeepCSR (both trained on ADNI_{large}), and FreeSurfer on the TRT dataset [31]. To this end, we compute cortical surfaces of brain scans of the same subject from the same day and com-

pare them with respect to their ASSD and HD. Since the morphology of the brain does not change within two consecutive scans of the same day, the resulting reconstructions should be identical up to some variations resulting from the imaging process. We follow [4] in this experiment and align pairs of images with the iterative closest-point algorithm (ICP) before comparing them. The results in Table 3 demonstrate that Vox2Cortex has better reproducibility than DeepCSR and FreeSurfer. Further, we compare the inference time of the methods and find that Vox2Cortex is around 25 times faster than DeepCSR.

4.5. Cortical Thickness

The surface reconstruction enables the computation of cortical thickness, which is an important biomarker for assessing cortical atrophy in neurodegenerative disorders. In this experiment, we compute cortical thickness measurements from Vox2Cortex surfaces and compare them against FreeSurfer thickness measures. The thickness is computed by means of the closest-point correspondences between the WM and pial surfaces [35]. The median cortical thickness error per vertex on the OASIS test set compared to closest FreeSurfer vertices is 0.305mm [lower quartile: 0.140mm, upper quartile: 0.564mm]. For comparison, atrophy in re-

Method	ASSD (mm)	HD (mm)	> 1mm	> 2mm	Inference time
Vox2Cortex (ours)	0.228 \pm 0.048	0.478 \pm 0.101	0.80%	0.06%	18.0s
Vox2Cortex* (ours)	0.225 \pm 0.049	0.471 \pm 0.103	0.77%	0.06%	2.1s
DeepCSR	0.357 \pm 0.284	0.739 \pm 0.595	5.82%	2.23%	445.7s
FreeSurfer	0.291 \pm 0.133	0.605 \pm 0.279	2.87%	0.67%	>4h

Table 3. Comparison of reconstruction consistency in terms of ASSD and HD on the TRT dataset. In addition, we provide the inference time of the respective method per 3D scan. An asterisk (*) indicates smaller templates ($\approx 42,000$ instead of $\approx 168,000$ vertices per surface).

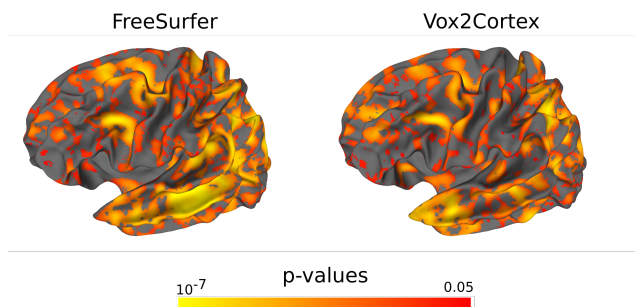


Figure 5. Group comparison of cortical atrophy in the left hemisphere between patients diagnosed with Alzheimer’s disease and healthy controls on the ADNI_{large} test-split.

gions affected by Alzheimer’s disease often lies in the range of one millimeter or more [27]. In the supplement, we also visualize the respective measurements on one subject from the OASIS dataset. On the ADNI_{large} test-split, we performed a group analysis by testing individual vertices for significantly lower thickness in subjects diagnosed with Alzheimer’s disease ($n = 50$) compared to healthy controls ($n = 124$). For this purpose, we leveraged the FreeSurfer FsAverage template to have all thickness measurements in the same domain and perform statistical tests per vertex. We also computed cortical thickness on FreeSurfer meshes with closest-point correspondences instead of using the default thickness measurements to have a fair comparison. Figure 5 visualizes the p-values (t-test, one sided) for the left hemisphere. As there is a high similarity between significance maps created by Vox2Cortex and FreeSurfer, we conclude that Vox2Cortex is well-suited for group analysis for cortical thickness.

5. Limitations and Potential Negative Impact

All presented results were created with the unmodified output of our model. However, we would like to mention that we cannot guarantee that predicted meshes have no self-intersections. This property might be necessary for some tasks, e.g., volume measurements. However, all self-intersections can be easily removed with algorithms like MeshFix [1] (runtime app. 16s for a surface with 168,000 vertices), keeping the remaining mesh unchanged. As there exist no manually generated ground-truth surfaces for our

data, we used FreeSurfer surfaces as pseudo ground-truth as in [4, 30]. Although we identified and removed some faulty labels from our training set, some labels might still be noisy. Our method has the potential to help radiologists to visualize the brain surfaces and to compute measurements like cortical thickness quickly. The predictions should, however, not be used to make clinical decisions. Our model has only been tested on the data discussed in this work, and we cannot guarantee that it will perform as well on unseen data from different domains. Further, the data analyzed in this study only includes healthy subjects and subjects with dementia. If other morphological changes in the brain are present (e.g., tumors), the trained model might provide wrong predictions. We have balanced our data splits with respect to patient’s age and sex to avoid bias, but our model might still lack fairness and might discriminate against groups of people underrepresented in the given datasets.

6. Conclusion

In this work, we have presented a novel method for the simultaneous reconstruction of white matter and pial surfaces from brain MR images. For the first time, we start from a general brain template and deform it in multiple iterations leveraging thoroughly combined convolutional and graph convolutional neural networks. To achieve a high reconstruction accuracy in densely folded regions, we successfully exploit ground-truth curvatures in a novel curvature-weighted Chamfer-loss function. We believe that this loss can also be helpful in other medical and non-medical fields where complex 3D geometries appear. Our experiments show that the proposed combination of loss functions directly yields state-of-the-art cortical surfaces as the output of an end-to-end trainable architecture, while being orders of magnitude faster. Finally, we demonstrated that accurate cortical thickness maps can be derived for studying atrophy in Alzheimer’s disease.

Acknowledgment This research was partially supported by the Bavarian State Ministry of Science and the Arts and coordinated by the bidt, and the BMBF (DeepMentia, 031L0200A). We gratefully acknowledge the computational resources provided by the Leibniz Supercomputing Centre (www.lrz.de).

References

- [1] Marco Attene. A lightweight approach to repairing digitized polygon meshes. *The Visual Computer*, 26(11):1393–1406, Nov. 2010. **8**
- [2] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Velickovic. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *CoRR*, abs/2104.13478, 2021. **3**
- [3] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: Learning dense volumetric segmentation from sparse annotation. In Sebastien Ourselin, Leo Joskowicz, Mert R. Sabuncu, Gozde Unal, and William Wells, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016*, pages 424–432, Cham, 2016. Springer International Publishing. **3**
- [4] Rodrigo Santa Cruz, Léo Lebrat, P. Bourgeat, C. Fookes, J. Fripp, and O. Salvado. Deepcsr: A 3d deep learning approach for cortical surface reconstruction. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 806–815, 2021. **2, 3, 6, 7, 8**
- [5] A. Dale, B. Fischl, and M. Sereno. Cortical surface-based analysis i. segmentation and surface reconstruction. *NeuroImage*, 9:179–194, 1999. **1**
- [6] Anders M Dale, Bruce Fischl, and Martin I Sereno. Cortical surface-based analysis: I. segmentation and surface reconstruction. *Neuroimage*, 9(2):179–194, 1999. **2, 4**
- [7] M.-P. Dubuisson and A.K. Jain. A modified hausdorff distance for object matching. In *Proceedings of 12th International Conference on Pattern Recognition*, volume 1, pages 566–568 vol.1, 1994. **6**
- [8] H. Fan, H. Su, and L. Guibas. A point set generation network for 3d object reconstruction from a single image. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2463–2471, 2017. **5**
- [9] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012. **1, 2, 6**
- [10] Bruce Fischl, Martin I Sereno, and Anders M Dale. Cortical surface-based analysis: Ii: inflation, flattening, and a surface-based coordinate system. *Neuroimage*, 9(2):195–207, 1999. **2**
- [11] Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '97, page 209–216, USA, 1997. ACM Press/Addison-Wesley Publishing Co. **6**
- [12] Georgia Gkioxari, Justin Johnson, and Jitendra Malik. Mesh r-cnn. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9784–9794, 2019. **5**
- [13] Karthik Gopinath, Christian Desrosiers, and Herve Lombaert. Segrecon: Learning joint brain surface reconstruction and segmentation from images. In Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, pages 650–659, Cham, 2021. Springer International Publishing. **2**
- [14] Thibault Groueix, Matthew Fisher, Vladimir G. Kim, Bryan Russell, and Mathieu Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018. **3**
- [15] Kunal Gupta and Manmohan Chandraker. Neural mesh flow: 3d manifold mesh generation via diffeomorphic flows. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1747–1758. Curran Associates, Inc., 2020. **5**
- [16] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. **3**
- [17] Leonie Henschel, Sailesh Conjeti, Santiago Estrada, Kersten Diers, Bruce Fischl, and Martin Reuter. Fastsurfer - a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage*, 219:117012, 2020. **1, 2**
- [18] Yuankai Huo, Zhoubing Xu, Yunxi Xiong, Katherine Aboud, Prasanna Parvathaneni, Shunxing Bao, Camilo Bermudez, Susan M. Resnick, Laurie E. Cutting, and Bennett A. Landman. 3d whole brain segmentation using spatially localized atlas network tiles. *NeuroImage*, 194:105–119, 2019. **1**
- [19] D.P. Huttenlocher, G.A. Klanderman, and W.J. Rucklidge. Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993. **6**
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 448–456, Lille, France, 07–09 Jul 2015. PMLR. **3**
- [21] Haris Iqbal. Plotneuralnet. <https://github.com/HarisIqbal88/PlotNeuralNet>, 2020. **4**
- [22] Fabian Isensee, Paul F. Jaeger, Simon A. A. Kohl, Jens Petersen, and Klaus Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 2020. **3, 7**
- [23] Fanwei Kong and Shawn C. Shadden. Whole heart mesh generation for image-based computational simulations by learning free-from deformations, 2021. **3, 6**
- [24] Fanwei Kong, Nathan Wilson, and Shawn C. Shadden. A deep-learning approach for direct whole-heart mesh reconstruction, 2021. **3, 5, 6**
- [25] Gina R. Kuperberg, Matthew R. Broome, Philip K. McGuire, Anthony S. David, Mohammad Istiadjid Eddy, Fukujiro Ozawa, Donald C. Goff, W. Caroline West, S. C. R. Williams, Andre van der Woburn Kouwe, D. H. Salat, A M Dale, and Bruce R. Fischl. Regionally localized thinning of the cerebral cortex in schizophrenia. *Schizophrenia Research*, 60:199–200, 2003. **1**
- [26] B. Landman and S. Warfield. Miccai 2012 workshop on multi-atlas labeling. In *MICCAI Grand Challenge and Workshop on Multi-Atlas Labeling, CreateSpace Independent Publishing Platform, Nice, France*, 2012. **6**

- [27] Jason P. Lerch, Jens C. Pruessner, Alex P. Zijdenbos, Harald Hampel, Stefan J. Teipel, and Alan C. Evans. Focal decline of cortical thickness in alzheimer's disease identified by computational neuroanatomy. *Cerebral cortex*, 15 7:995–1001, 2005. 1, 8
- [28] 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. 1, 2, 7
- [29] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87*, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery. 1, 2
- [30] Qiang Ma, Emma C. Robinson, Bernhard Kainz, Daniel Rueckert, and Amir Alansary. PIALNN: A Fast Deep Learning Framework for Cortical PIAL Surface Reconstruction. In Ahmed Abdulkadir, Seyed Mostafa Kia, Mohamad Habes, Vinod Kumar, Jane Maryam Rondina, Chantal Tax, and Thomas Wolfers, editors, *Machine Learning in Clinical Neuroimaging*, volume 13001, pages 73–81. Springer International Publishing, Cham, 2021. 3, 8
- [31] Julian R. Maclaren, Zhaoying Han, Sjoerd B. Vos, Nancy J. Fischbein, and Roland Bammer. Reliability of brain volume measurements: A test-retest dataset. *Scientific Data*, 1, 2014. 6, 7
- [32] 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, 09 2007. 6
- [33] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [34] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In Hans-Christian Hege and Konrad Polthier, editors, *Visualization and Mathematics III*, pages 35–57, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. 5
- [35] Michael I. Miller, Allan B. Massie, J. Tilak Ratnanather, Kelly N. Botteron, and John G. Csernansky. Bayesian construction of geometrically based cortical thickness metrics. *NeuroImage*, 12(6):676–687, 2000. 7
- [36] Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. ICML 2010, page 807–814, Madison, WI, USA, 2010. Omnipress. 3
- [37] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. Laplacian mesh optimization. In Y. T. Lee, Siti Mariyam Hj. Shamsuddin, Diego Gutierrez, and Norhaida Mohd. Suaib, editors, *Proceedings of the 4th International Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia 2006, Kuala Lumpur, Malaysia, November 29 - December 2, 2006*, pages 381–389. ACM, 2006. 5
- [38] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 3
- [39] CG Phillips, S Zeki, and HB Barlow. Localization of function in the cerebral cortex: past, present and future. *Brain*, 107(1):328–361, 1984. 1
- [40] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 3
- [41] Anne-Marie Rickmann, Abhijit Guha Roy, Ignacio Sarasua, and Christian Wachinger. Recalibrating 3d convnets with project & excite. *IEEE transactions on medical imaging*, 39(7):2461–2471, 2020. 1
- [42] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 3
- [43] Abhijit Guha Roy, Sailesh Conjeti, Nassir Navab, C. Wachinger, and Alzheimer's Disease Neuroimaging Initiative. Quicknat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186:713–727, 2019. 1
- [44] Stewart Shipp. Structure and function of the cerebral cortex. *Current Biology*, 17(12):R443–R449, 2007. 1
- [45] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 7462–7473. Curran Associates, Inc., 2020. 3
- [46] Edward Smith, Scott Fujimoto, Adriana Romero, and David Meger. GEOMETRICS: Exploiting geometric structure for graph-encoded objects. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5866–5876, Long Beach, California, USA, 09–15 Jun 2019. PMLR. 5
- [47] Stephen M Smith, Mark Jenkinson, Mark W Woolrich, Christian F Beckmann, Timothy EJ Behrens, Heidi Johansen-Berg, Peter R Bannister, Marilena De Luca, Ivana Drobnjak, David E Flitney, et al. Advances in functional and structural mr image analysis and implementation as fsl. *NeuroImage*, 23:S208–S219, 2004. 2
- [48] J. Vollmer, R. Mencl, and H. Muller. Improved Laplacian Smoothing of Noisy Surface Meshes. *Computer Graphics Forum*, 18(3):131–138, Sept. 1999. 4
- [49] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XI*, volume 11215 of *Lecture Notes in Computer Science*, pages 55–71. Springer, 2018. 3, 5, 6

- [50] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Hang Yu, Wei Liu, Xiangyang Xue, and Yu-Gang Jiang. Pixel2mesh: 3d mesh model generation via image guided deformation. *IEEE transactions on pattern analysis and machine intelligence*, 2020. [2](#), [3](#)
- [51] Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. 3dn: 3d deformation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. [3](#)
- [52] Chao Wen, Y. Zhang, Zhuwen Li, and Yanwei Fu. Pixel2mesh++: Multi-view 3d mesh generation via deformation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1042–1051, 2019. [2](#), [3](#)
- [53] Udaranga Wickramasinghe, Edoardo Remelli, Graham Knott, and Pascal Fua. Voxel2mesh: 3d mesh model generation from volumetric data. In Anne L. Martel, Purang Abolmaesumi, Danail Stoyanov, Diana Mateus, Maria A. Zuluaga, S. Kevin Zhou, Daniel Racoceanu, and Leo Joskowicz, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*, pages 299–308, Cham, 2020. Springer International Publishing. [3](#), [5](#), [6](#), [7](#)
- [54] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomír Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *NeurIPS*, 2019. [3](#)
- [55] Guodong Zeng, Xin Yang, Jing Li, Lequan Yu, Pheng-Ann Heng, and Guoyan Zheng. 3d u-net with multi-level deep supervision: Fully automatic segmentation of proximal femur in 3d mr images. In Qian Wang, Yinghuan Shi, Heung-II Suk, and Kenji Suzuki, editors, *Machine Learning in Medical Imaging*, pages 274–282, Cham, 2017. Springer International Publishing. [3](#)
- [56] Zhengxin Zhang, Qingjie Liu, and Yunhong Wang. Road extraction by deep residual u-net. *IEEE Geoscience and Remote Sensing Letters*, 15(5):749–753, 2018. [3](#)