XSurfer: Reconstructing surface meshes of cerebral and cerebellar cortex from diverse MRI data using untrained neural networks

Haoxiang Li[1,2†], Mingxuan Liu[1†], Divya Varadarajan[3], Zhangxuan Hu[2], Qiyuan Tian[1*], and Jonathan Polimeni[2‡]

[1] School of Biomedical Engineering, Tsinghua University, Beijing, China
[2] Department of Radiology, Stanford University, Stanford, USA
[3] Massachusetts General Hospital, Harvard Medical School, Boston, USA

Abstract. Cortical surface reconstruction (CSR) is widely used in neuroimaging and is crucial for quantitative analyses of cerebral cortical thickness and sulcal morphology. While CSR is a mature technology when applied routinely to adult T_1-weighted brain MRI data of the cerebrum, it remains underexplored for a broad range of MRI contrasts, resolutions, ages, species, and brain structures such as the cerebellum. To address this challenge, we propose **`XSurfer`**, a contrast- and resolution-agnostic CSR framework that performs optimization on single images using an untrained neural network such that training data are not needed. Specifically, **`XSurfer`** starts from segmentations of 3D MRI data, applies sulcal cerebrospinal fluid (CSF) recovery in sulci to address limitations caused by partial-volume effects, identifies CSF voxels in deep sulci, and extracts topologically correct initial white matter (WM) surfaces and pseudo-target pial surfaces. It then employs an untrained neural network to predict velocity fields, which are driven by mesh-based and geometrically-constrained loss function terms to achieve diffeomorphic surface deformation, thereby reconstructing accurate pial surfaces. We evaluated **`XSurfer`** across different contrasts, resolutions and seven datasets spanning adult, fetal, infant, *ex vivo* human, and non-human primate brain; **`XSurfer`** provides accurate CSR for all cases, surpassing state-of-the-art learning-based methods and `recon-all-clinical` recently provided by FreeSurfer as a more general-purpose CSR method. By enabling precise, contrast- and resolution-agnostic CSR without pretraining, **`XSurfer`** facilitates cross-domain multi-modal neuroimaging and standardizes morphometric analysis for neuroscience research and clinical applications. The source code is publicly available at: `https://github.com/birthlab/XSurfer`.

Keywords: Cortical surface reconstruction · Neuroimaging · Brain morphometry · Untrained neural networks

† Equal contribution.
∗ Corresponding author: `qiyuantian@tsinghua.edu.cn`
‡ Senior author.

1 Introduction

Magnetic Resonance Imaging (MRI) is a widely-used non-invasive neuroimaging technique for studying human brain structure and function, providing 3D volumetric images with varying tissue contrasts and spatial resolutions, with most neuroimaging studies investigating the cerebrum. However, the highly folded nature of the cerebral cortex makes it difficult to represent this complex anatomy at the voxel level [33]. Therefore, cortical surface reconstruction (CSR) is commonly used to extract the inner (white-matter, WM) and outer (pial) surfaces of the cortical gray matter (GM) and to model the cortex as a pair of triangular surface meshes embedded in 3D space (Fig. 1) [8]. CSR is a widely used post-processing step in both research and clinical applications. It is essential for (a) visualization of the folded cerebral cortex [11]; (b) quantitative morphometry—e.g., for calculating cortical thickness, a key biomarker of brain development and disease [9]; and (c) surface-based atlasing and parcellation, enabling consistent localization of neuroimaging data across individuals for subject-specific and group-level analyses [41]. CSR is also used for downstream surface-based analyses of diffusion MRI, functional MRI, and other MRI modalities. Beyond adult cerebral imaging, CSR is also beneficial for studying fetal and neonatal brain development [23, 25], and even cortical morphometry in other species for comparative studies [4].

However, obtaining accurate reconstructions across MRI modalities, magnetic field strengths, and resolutions—and across diverse imaging domains (e.g., *ex vivo* 7T MRI of brain specimens, diffusion MRI, fetal and infant MRI, cerebellar MRI, and non-human primate MRI, as shown in Fig. 2a)—remains challenging due to substantial variations in tissue contrast and cortical GM geometry. Existing CSR methods face a fundamental trade-off between accuracy and generalization. Classical methods [11, 16] begin by reconstructing the white surface at the gray-white interface then smoothly deform it to the outer surface at the gray-CSF interface using image intensity gradients combined with geometric constraints including curvature and smoothness priors. Although accurate and robust, these methods are slow (hours per subject) and largely tailored to specific populations (e.g., human adults) and MRI modalities (e.g., T_1-weighted anatomical data). Recent learning-based approaches [19, 22, 25, 28, 30] learn this deformation from extensive training datasets but fail to generalize beyond their training distributions. This limitation

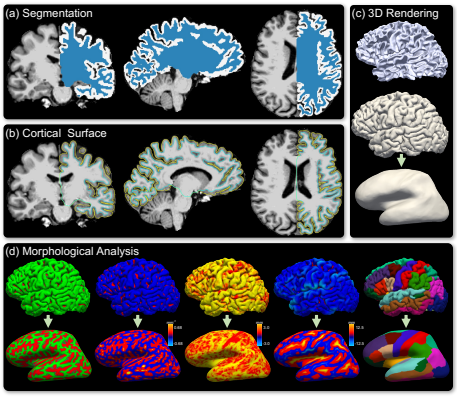

Fig. 1: Overview of CSR and its applications. (a) MRI cross-sections with WM (in blue) and GM (in white) segmentation. (b) CSR representing the WM (cyan) and pial (yellow) surfaces. (c) Rendered and inflated surfaces. (d) Morphological analysis: sulcal/gyral labels, curvature, thickness, sulcal depth, and parcellation.

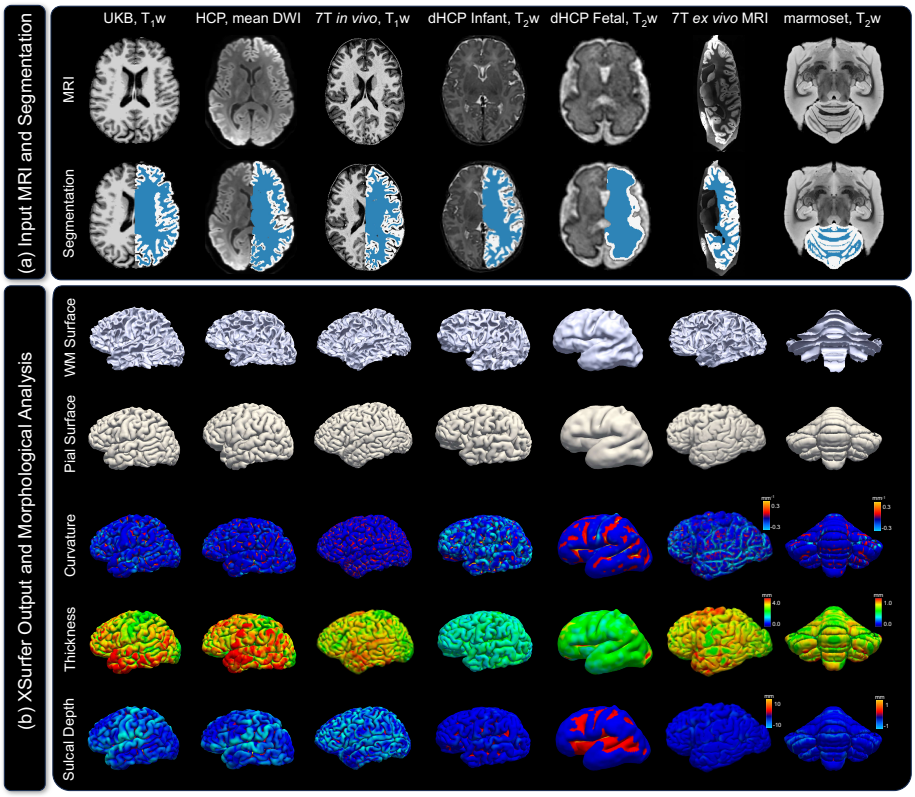

Fig. 2: Example input data and XSurfer CSR results across all seven datasets considered in this study. (a) Input MRI and WM/GM segmentation labels for CSR. (b) XSurfer output (WM and pial surfaces) and morphological analysis (curvature, thickness, and sulcal depth), with two separate colorbars for cerebrum-type and cerebellum-type data.

is particularly problematic when extending CSR to diverse imaging domains where annotated training data is scarce or unavailable.

In essence, a key bottleneck shared by both conventional and learning-based methods is their tight coupling to specific image contrasts or training distributions. Cortical tissue segmentation is a domain-agnostic intermediate representation. Many open-source, pretrained deep learning models [2, 3, 13, 43] already provide robust, domain-agnostic segmentations; segmentation is also easier to annotate manually, and many datasets include high-quality segmentation labels (e.g., UKB [1] and dHCP [10] datasets). The conversion from volumetric segmentation to the corresponding surface reconstruction (that represents the inner and outer surfaces of the cortex with sub-voxel accuracy) is key to achieving general-purpose CSR. However, generating a pair of surface meshes from a voxel-based segmentation remains challenging. Although the classic Marching Cubes [27] directly constructs surface meshes from a 3D segmentation, it struggles to generate an accurate pial surface because of the complex folding of the cortex: sulcal

banks are tightly apposed, causing neighboring folds to become merged, and partial-volume effects are difficult to capture in the segmentation, so Marching Cubes and other isosurface methods fail to produce accurate surface placement in deep sulci.

Recent studies show that untrained neural networks—whose architectures inherently encode geometric regularization—can serve as effective and generalizable tools for solving inverse imaging problems [20, 34], inspiring our training-free CSR approach.

For these reasons, we present **`XSurfer`**, a resolution- and contrast-agnostic cortical surface reconstruction framework that achieves robust "zero-shot" generalization [35] across diverse imaging domains without requiring training data. **`XSurfer`** adopts a generic segmentation-refinement strategy to identify distinct folds of the cortical GM and to use this to generate an improved initial estimate of the pial surface as a target. Then an untrained neural network is used to predict a deformation field that expands the white surface outward to approximately match the target pial surface, producing a more accurate pial surface mesh with an identical mesh topology and vertex correspondence with the white surface. During deformation, our proposed mesh-based and geometrically-constrained loss functions preserve surface topology and geometry while achieving accurate placement at the GM–CSF boundary. Notably, **`XSurfer`** is performed for each MRI volume individually, without any training data, and runs in minutes—substantially faster than classical CSR methods that take hours (e.g., FreeSurfer). Additionally, because the deformation process and loss function are identical across datasets, **`XSurfer`** generalizes CSR across MRI modalities, age groups, species, and brain regions.

To summarize, our contributions are as follows:

- We introduce **`XSurfer`**, a training-free, test-time optimization framework that leverages an untrained neural network to optimize a subject-specific cost function, effectively eliminating the dependency on large, annotated datasets and enabling robust generalization across diverse and unseen data domains (Fig. 2b).
- We design a segmentation refinement and initial surface extraction pipeline that corrects topological errors in the raw segmentation and efficiently generates topologically-correct initial surfaces, providing a robust foundation for the deformation process.
- We formulate a comprehensive loss function that integrates surface-driven constraints and geometric priors inspired by the loss function used in FreeSurfer [11], the benchmark method that has been shown to be robust across an impressive variety of T_1 anatomical data of the adult human brain. Our loss function jointly guides the untrained network to reconstruct pial surfaces with high geometric accuracy and guaranteed topological correctness, all without ground-truth surfaces or supervision.

2 Related work

2.1 Optimization-based CSR

Conventional optimization-based CSR methods [11, 16, 32, 40] are the most widely used and mature approaches. The most commonly used approach is FreeSurfer [11], which is a set of tools including a pipeline for automatic CSR as well as subsequent parcellation and morphometric analyses. These methods typically perform iterative surface deformation from an initial white surface guided by image-intensity features (e.g.,

gradients at the GM–CSF boundary) together with geometric priors (e.g., smoothness and curvature regularization) to enforce anatomical plausibility and topological correctness. Different pipelines are tailored to specific populations and modalities: FreeSurfer [11] and BrainSuite [40] are suitable for adult T_1 anatomical MRI data, while CRUISE [16] and dhcp-structural-pipeline [32] are designed for infant MRI. Recent variants like FastSurfer [18], FetalSurfer [26] and `recon-all-clinical` [15] use deep learning to accelerate segmentation but retain the standard FreeSurfer CSR methodology. While these methods produce highly accurate reconstructions for their target domains, they are computationally expensive (hours per subject) and fail to generalize across imaging modalities or populations.

2.2 Learning-based CSR

Recently, the broad availability of deep learning software has led to new learning-based CSR methods that can be broadly categorized into implicit and explicit frameworks based on their geometric representations, as well as weakly supervised methods that are designed to reduce data dependency.

Implicit CSR Methods, such as DeepCSR [7] and SegRecon [14], typically employ Convolutional Neural Networks (CNNs) to predict Signed Distance Functions (SDFs), from which surface meshes are extracted via Marching Cubes [27] followed by topology correction [24]. Approaches like FastCSR [37] can only predict the SDF for the white surface. Implicit CSR methods struggle with pial surface reconstruction due to the highly convoluted cortical geometry and partial-volume effects in tight sulci. Furthermore, they fail to maintain vertex correspondence between white and pial surfaces, limiting their utility for downstream morphometric analyses that require such correspondence.

Explicit CSR Methods, such as CortexODE [30], CoTAN [28], Vox2Cortex [5] and CorticalFlow++ [39], use neural networks to deform an initial surface toward the target anatomy, often leveraging ODE solvers. They can be classified as *point-based* [5, 30] (predicting deformations from local image features and mesh topology) or *field-based* [6, 28, 39] (predicting 3D deformation fields from volumetric data). These methods achieve fast inference (seconds) once trained on large-scale datasets with ground-truth surfaces. However, their reliance on training data limits generalization: they perform well only on populations and modalities similar to their training distributions, and struggle when applied to novel imaging domains where annotated data is unavailable.

Weakly-supervised learning. CoSeg [29] and SegCSR [46] attempt to bypass the requirement for ground-truth by using instead surfaces generated from synthetic GM segmentations. These segmentations are synthesized by filling the volume between the white and pial surface reconstructions, thus it represents the "perfect" segmentation corresponding to the pair of surfaces. For new data, due to partial-volume effects and segmentation errors, a perfect GM segmentation is practically unattainable unless the full FreeSurfer pipeline is executed (taking 6–7 hours). In the absence of such perfect segmentations, these weakly-supervised methods cannot guarantee a successful reconstruction. Furthermore, they are still learning-based approaches that necessitate large-scale datasets for training.

2.3 CSF enhancement in deep sulci

In tight sulcal folds, partial-volume effects often cause opposite GM banks to appear fused, hindering accurate pial surface placement. To address this, methods like anatomically consistent editing (ACE) [17] modify GM labels by detecting sulcal medial axes, while CLASP [21] uses partial-volume classification to identify buried CSF. These preprocessing techniques have proven effective for specific modalities (e.g., adult T_1 MRI) but require careful parameter tuning and domain-specific adaptations. Our approach incorporates similar principles through a more generalizable segmentation refinement strategy that adapts to different imaging domains without manual tuning.

3 Methods

3.1 Overview of `XSurfer`

Given a 3D brain MRI volume $V \in \mathbb{R}^{D_1 \times D_2 \times D_3}$, **`XSurfer`** reconstructs a pair of topologically correct, point-wise corresponding cortical surfaces (i.e., white surface and pial surface) for each cerebral (or cerebellar) hemisphere. Specifically, each surface is mathematically represented as a triangular mesh $\mathcal{S} = \{\mathcal{V}, \mathcal{F}\}$ with vertex coordinates $\mathcal{V} \in \mathbb{R}^{N \times 3}$ and shared face connectivity $\mathcal{F} \in \mathbb{Z}^{M \times 3}$. **`XSurfer`** achieves domain-agnostic CSR with two cascaded stages (Fig. 3): $V \xrightarrow{\textbf{I}} \{\mathcal{S}_0^{\text{WM}}, \mathcal{S}^{\text{target}}\} \xrightarrow{\textbf{II}} \{\mathcal{S}_0^{\text{WM}}, \hat{\mathcal{S}}^{\text{pial}}\}$, where Stage I (**Topology-Preserving Initialization**) maps the input volume V to a refined anatomical segmentation using morphological simple-point constraints to recover unresolved sulci, and subsequently computes genus-zero signed distance fields Φ' to extract an initial white surface $\mathcal{S}_0^{\text{WM}}$ and a target pial boundary $\mathcal{S}^{\text{target}}$; Stage II (**Neural Surface Deformation**) models the cortical expansion as a diffeomorphic flow $\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t), t)$, relying on an untrained network optimized from scratch to predict the subject-specific time-varying velocity field $\mathbf{u}(\mathbf{x}, t)$. By minimizing a composite geometric loss $\mathcal{L}_{\text{total}}$ over the integration trajectory, the flow progressively warps $\mathcal{S}_0^{\text{WM}}$ towards $\mathcal{S}^{\text{target}}$ to establish the final pial surface $\hat{\mathcal{S}}^{\text{pial}}$ while preserving mesh quality and explicitly preventing self-intersections.

3.2 Stage I: Topology-Preserving Initialization

Volumetric Tissue Segmentation. Given the input MRI volume V, we first obtain a voxel-wise tissue segmentation using a pretrained 3D U-Net [38]. Each voxel is assigned one of five labels: background, left white matter, right white matter, left gray matter, or right gray matter. We denote the resulting label map as $U \in \{0, 1, 2, 3, 4\}^{D_1 \times D_2 \times D_3}$. Here, the WM label implicitly defines the *interior* of the initial cortical surface, encompassing cortical WM and subcortical structures (e.g., basal ganglia, ventricles). The segmentation module in **`XSurfer`** is designed to be domain-agnostic; while the 3D U-Net serves as the default, it can be substituted with modality-specific segmentation methods (e.g., SynthSeg [2] for adult multi-contrast MRI, BOUNTI [43] for fetal MRI, or GOUHFI [12] for 7T imaging) to optimize the downstream initialization boundary.

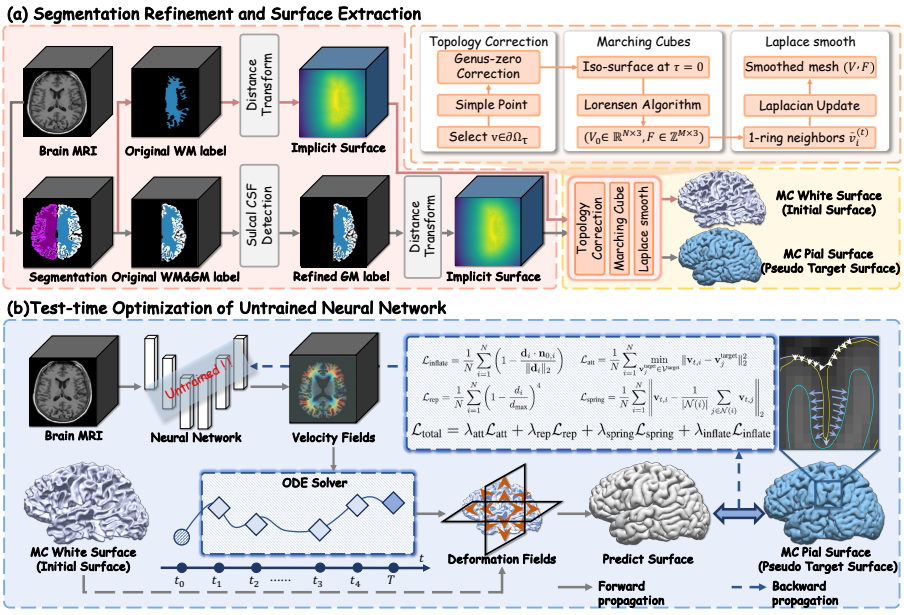

Fig. 3: Overview of `XSurfer` (A) Stage I: Segmentation refinement and initial/target surface generation via the Marching Cubes algorithm. (B) Stage II: An untrained 3D U-Net predicts a multi-scale deformation field, guided by a composite loss function, to evolve the initial white matter surface toward the target pial surface. Dashed lines denote the iterative optimization process.

Morphology-Guided Sulcal CSF Recovery. Partial-volume effects in tight sulci often cause buried CSF voxels to be misclassified as GM, making the two opposing GM banks appear connected in the voxel segmentation [17]. To correct this problem before surface extraction, we refine the GM label by a topology-preserving expansion from the WM label.

Let C denote the current cortical voxel object, initialized by the WM label in U. We then examine neighboring GM-labeled voxels one at a time, expanding outward from the WM boundary. For each candidate voxel $\mathbf{q}$, we first check whether its image intensity supports GM assignment. We then test whether adding $\mathbf{q}$ to C is a *simple-point* update, meaning that the addition does not change the topology of C within the local $3 \times 3 \times 3$ neighborhood. If both conditions are satisfied, $\mathbf{q}$ is accepted as part of the cortical object. If the intensity supports GM but $\mathbf{q}$ is not a simple point, adding it would create a topological bridge between two nearby banks. We therefore reject this voxel from the GM mask and mark it as recovered sulcal CSF.

Formally, a candidate voxel $\mathbf{q}$ is accepted only when

$$\mathcal{I}(\mathbf{q}) > \tau_{\mathrm{GM}} \quad \text{and} \quad T_n(\mathbf{q}, C) = 1 \quad \text{and} \quad T_n(\mathbf{q}, \bar{C}) = 1, \tag{1}$$

where $\mathcal{I}(\mathbf{q})$ is the image intensity at $\mathbf{q}$, τ_{GM} is the GM intensity threshold, $\bar{C}$ is the background complement of C, and T_n denotes the topological number computed in the $3 \times 3 \times 3$ neighborhood of $\mathbf{q}$. This criterion preserves both foreground and background

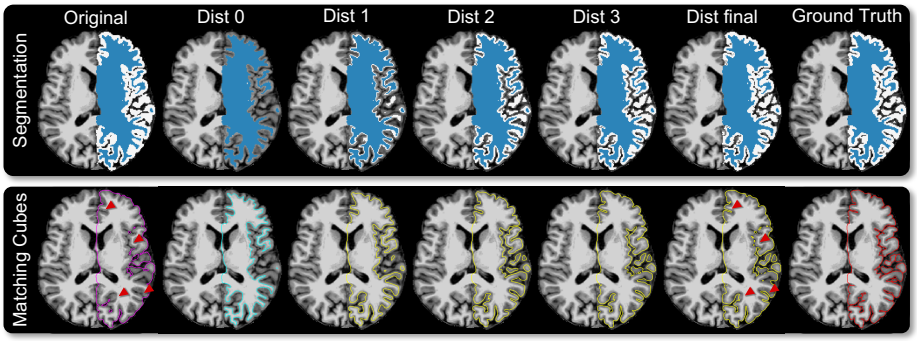

Fig. 4: Sulcal CSF recovery via topology-preserving GM expansion. Top row: GM segmentation (white) progressively expands from the WM (blue) boundary, adding voxels that satisfy both intensity and simple-point criteria. Dist 0 to Dist final illustrate the iterative expansion process. **Bottom row:** Marching Cubes surfaces extracted at each stage. The WM surface (cyan) is largely unaffected by partial-volume effects. Comparing the original segmentation surface (magenta) with the refined surface (yellow) at the red triangle markers reveals that iterative refinement successfully recovers buried CSF in tight sulci, yielding results closer to the ground truth surface (red).

topology during expansion. As a result, this segmentation-agnostic voxel-space refinement separates fused sulcal banks and restores buried CSF regions missed by the initial segmentation (Fig. 4), and can be applied to both cerebral and cerebellar segmentations.

Topological SDF Correction and Surface Extraction. Given the morphologically refined tissue masks, we compute continuous signed distance fields (SDFs) Φ for both the WM mask and the combined WM-GM mask via a Chamfer distance transform, followed by Gaussian regularization ($\sigma = 0.5$). To strictly enforce the genus-zero topological requirement for the subsequent surface deformation, we employ a fast-marching correction framework. Initialized from the boundary, a priority queue $\mathcal{Q}$ propagates through Φ. If a processed voxel $\mathbf{v}$ violates the local $3 \times 3 \times 3$ simple-point criterion (i.e., acts as a critical point creating handles), its distance value is heavily penalized to construct a topologically corrected field Φ':

$$\Phi'(\mathbf{v}) = \min_{j \in \mathcal{N}(\mathbf{v})} \{\Phi(j) + \epsilon\}, \tag{2}$$

where $\epsilon = 10^{-5}$ and $\mathcal{N}(\mathbf{v})$ denotes the processed neighbors. Finally, applying the Marching Cubes algorithm to the zero-level sets of the corrected SDFs yields the **initial white surface** $\mathcal{S}_0^{\text{WM}} = \{\mathcal{V}_0, \mathcal{F}\}$ and the **target pial surface** $\mathcal{S}^{\text{target}} = \{\mathcal{V}^{\text{target}}, \mathcal{F}^{\text{target}}\}$. Staircase artifacts are subsequently suppressed using $n = 2$ iterations of Laplacian smoothing with weight $\lambda = 1.0$:

$$\mathbf{v}_i^{(t+1)} = (1-\lambda)\mathbf{v}_i^{(t)} + \frac{\lambda}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{v}_j^{(t)}. \tag{3}$$

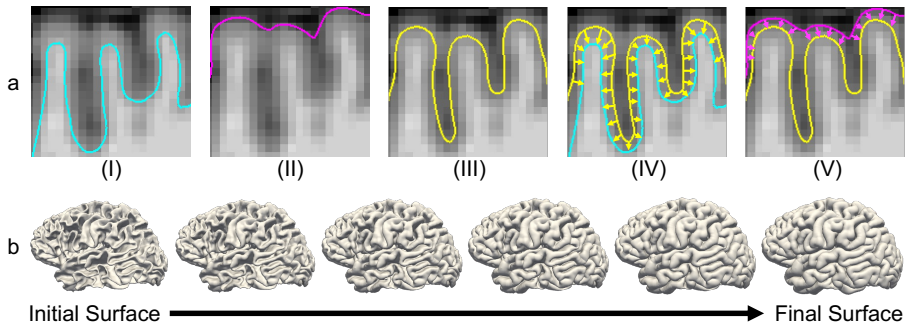

Fig. 5: Visualization of loss terms guiding surface deformation. (a) Cross-sectional views: (I) initial white surface (cyan), (II) target pial surface (magenta), (III) intermediate surface during evolution (yellow), (IV) repulsion loss preventing penetration of the white surface, (V) attraction loss driving toward the target surface. **(b)** Progressive 3D surface deformation from initial white surface to final pial surface.

3.3 Stage II: Neural Surface Deformation and Refinement

Diffeomorphic Surface Flow. Unlike supervised approaches that depend on large-scale datasets and pre-trained shape priors, we adopt a **test-time optimization (TTO)** strategy. This formulation explicitly optimizes a randomly initialized neural network from scratch for each specific subject, thereby implicitly learning a continuous representation of the subject's anatomical manifold and enabling robust generalization across domains. We model the geometric evolution from $\mathcal{S}_0^{\text{WM}}$ to the pial boundary as a continuous diffeomorphic flow governed by an ordinary differential equation (ODE):

$$\dot{\mathbf{x}}(t) = \mathbf{u}(\mathbf{x}(t), t), \quad \mathbf{x}(0) = \mathbf{x}_0, \quad t \in [0, T], \tag{4}$$

where $\mathbf{x}(t) \in \mathbb{R}^3$ tracks the surface vertex trajectory and $\mathbf{u}(\mathbf{x}, t)$ is the time-varying velocity field (TVF). Guided by a temporal-attention architecture, a 3D U-Net predicts $R \times M$ multi-scale stationary velocity fields (SVFs) $\mathbf{u}_i(\mathbf{x})$, which are linearly modulated by a time-conditioned channel-wise attention vector $\mathbf{p}(t)$ (where $\sum p_i(t) = 1$):

$$\mathbf{u}(\mathbf{x}, t) = \sum_{i=1}^{R \times M} p_i(t)\, \mathbf{u}_i(\mathbf{x}). \tag{5}$$

The optimal pial surface $\hat{\mathcal{S}}^{\text{pial}}$ is obtained by integrating this flow using the forward Euler method with step size $h = T/K$: $\mathbf{v}^{k+1} = \mathbf{v}^k + h\, \mathbf{u}(\mathbf{v}^k, hk)$.

Geometric Optimization Objectives. Let $\mathcal{S}_t = \{\mathcal{V}_t, \mathcal{F}\}$ denote the deformed surface at optimization step t. To optimize the network parameters without ground-truth topological supervision, the TTO framework minimizes a composite spatial loss $\mathcal{L}_{\text{total}}$ that strategically combines multiple geometric priors to guide the flow structurally from $\mathcal{S}_0^{\text{WM}}$ toward $\mathcal{S}^{\text{target}}$:

$$\mathcal{L}_{\text{total}} = \lambda_{\text{att}}\mathcal{L}_{\text{att}} + \lambda_{\text{rep}}\mathcal{L}_{\text{rep}} + \lambda_{\text{spring}}\mathcal{L}_{\text{spring}} + \lambda_{\text{inflate}}\mathcal{L}_{\text{inflate}}, \tag{6}$$

where each term uniquely constrains the trajectory and local geometry of the evolving mesh.

Attraction Loss ($\mathcal{L}_{\text{att}}$). This data term drives the deformed vertices toward the target boundary using a single-directional Chamfer distance. We use a single-directional rather than bi-directional Chamfer distance to reduce overfitting to pseudo-target errors. The one-sided term pulls the evolving surface toward the target boundary and into sulcal depths, while avoiding the reverse constraint that would force the pseudo-target surface to be fully matched by the mesh, including noisy partial-volume artifacts. For each vertex $\mathbf{v}_{t,i} \in \mathcal{V}_t$, it minimizes the Euclidean distance to its nearest spatial neighbor on the target pial surface (Fig. 5a-V):

$$\mathcal{L}_{\text{att}} = \frac{1}{N} \sum_{i=1}^{N} \min_{\mathbf{v}_j^{\text{target}} \in \mathcal{V}^{\text{target}}} \|\mathbf{v}_{t,i} - \mathbf{v}_j^{\text{target}}\|_2^2. \tag{7}$$

Repulsion Loss ($\mathcal{L}_{\text{rep}}$). Inspired by FreeSurfer's surface topology constraints, this penalty explicitly prevents the evolving surface from penetrating the initial white surface $\mathcal{S}_0^{\text{WM}}$ or forming localized self-intersections. Let d_i denote the signed distance of $\mathbf{v}_{t,i}$ to the nearest point on $\mathcal{S}_0^{\text{WM}}$ along the inward normal direction. A strong repulsive energy is activated specifically when $d_i < 0$ (indicating geometric penetration):

$$\mathcal{L}_{\text{rep}} = \frac{1}{N} \sum_{i=1}^{N} \left(1 - \frac{d_i}{d_{\text{max}}}\right)^4, \tag{8}$$

where d_{max} establishes the spatial safety margin for the repulsion (Fig. 5a-IV).

Spring Loss ($\mathcal{L}_{\text{spring}}$). To maintain global mesh regularity and prevent vertex clustering during aggressive deformations, we enforce a Laplacian smoothness constraint. It penalizes the spatial deviation of each vertex from the geometric centroid of its 1-ring topological neighbors $\mathcal{N}(i)$:

$$\mathcal{L}_{\text{spring}} = \frac{1}{N} \sum_{i=1}^{N} \left\| \mathbf{v}_{t,i} - \frac{1}{|\mathcal{N}(i)|} \sum_{j \in \mathcal{N}(i)} \mathbf{v}_{t,j} \right\|_2. \tag{9}$$

Inflation Loss ($\mathcal{L}_{\text{inflate}}$). This directional constraint ensures that the cortical expansion honors the underlying developmental anatomy (i.e., growing outward). Let $\mathbf{n}_{0,i}$ be the outward unit normal at vertex i on the initial white surface, and $\mathbf{d}_i = \mathbf{v}_{t,i} - \mathbf{v}_{0,i}$ denote the instantaneous displacement vector. The inflation term is formulated via cosine similarity:

$$\mathcal{L}_{\text{inflate}} = \frac{1}{N} \sum_{i=1}^{N} \left(1 - \frac{\mathbf{d}_i \cdot \mathbf{n}_{0,i}}{\|\mathbf{d}_i\|_2 + \epsilon}\right), \tag{10}$$

which approaches its global minimum when the displacement trajectory perfectly aligns with the initial outward normal.

4 Experiments

4.1 Datasets

A diverse collection of neuroimaging datasets spanning various life stages, species, and scanning protocols was utilized to evaluate the robustness of **XSurfer** across resolutions and imaging contrasts. Specifically: (*i*) For the quantitative evaluation of **XSurfer** against comparison methods, we utilized 50 adult subjects from the UK Biobank (**UKB**; T_1, 1 mm) [1] and 50 infant subjects from the developing Human Connectome Project (**dHCP**; T_2, 0.5 mm) [10]; (*ii*) Due to the lack of established CSR benchmarks for non-standard or ultra-high resolution data, the reconstruction precision of **XSurfer** was further evaluated on several challenging datasets: the Human Connectome Project (**HCP**) diffusion MRI (adult, 1 mm) [44], **dHCP** fetal (fetal T_2, 0.5 mm) [10], *in vivo* 7T (adult, 0.4 mm) [42], *ex vivo* 7T (post-mortem adult brain, 0.12 mm) [45], and the marmoset cerebellum (0.08 mm) [47]. Ground-truth surfaces for UKB and HCP were generated using FreeSurfer [11], for dHCP datasets using the dHCP structural pipeline [10], and for Marmoset using an adapted FreeSurfer [47]; 7T datasets were evaluated qualitatively without ground-truth. All experiments use Adam with learning rate 1×10^{-4} for 600 epochs; the fixed loss weights are $\lambda_{\text{att}} = 1$, $\lambda_{\text{rep}} = 0.5$, $\lambda_{\text{spring}} = 3$, and $\lambda_{\text{inflate}} = 3$, and optimization typically stabilizes after $\sim$400 epochs and takes 15–18 minutes per subject on a single H800 GPU.

4.2 Evaluation Metrics and Comparison Methods

We employ two standard metrics for quantitative evaluation: (1) Average Symmetric Surface Distance (ASSD) measures the mean bidirectional distance between predicted and ground-truth surfaces; and (2) Hausdorff Distance at 90th percentile (HD90) captures the maximum surface deviation while remaining robust to outliers. Both metrics are reported in millimeters (mm), with lower values indicating superior performance. Since **XSurfer** targets training-free, cross-domain CSR, we focus on baselines that do not require domain-specific surface training. Template-dependent or fully supervised methods such as CF++ [39], Vox2Cortex [5]/V2C-Flow+ [6], and V2C-MED [31] can outperform **XSurfer** in-distribution, but require cortical templates, registration training, or domain-specific training data that are unavailable for several domains considered here. **XSurfer** is compared against several traditional optimization-based CSR and learning-based CSR frameworks: (A) DeepCSR [7], a supervised implicit CSR method predicting super-resolution SDFs, retrained on the available 100 training cases because it does not support zero-shot inference; (B) `recon-all-clinical` [15], a domain-randomization method converting predicted SDFs to pseudo-T_1 images for FreeSurfer processing; (C) FreeSurfer [11] and InfantFS [48], the standard adult and infant surface generation pipelines. To avoid confusion with the FreeSurfer-derived reference surfaces used as ground truth on UKB, the FreeSurfer baseline in Table 1 was not run on the original T_1 image used to generate the reference surface. Instead, following common cross-modal practice, we synthesized a pseudo-T_1 volume from the segmentation labels by assigning fixed intensities to each tissue class: WM voxels were assigned an intensity of 110, GM voxels an intensity of 70, and CSF/background voxels an intensity of 0. We

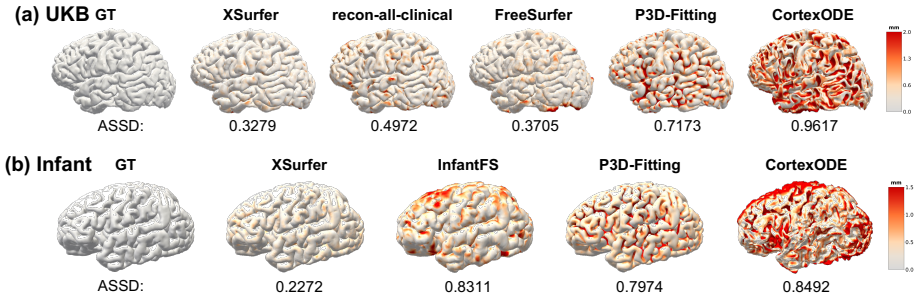

Fig. 6: CSR error across methods. Point-to-surface distance errors relative to ground-truth CSR (in units of mm) are rendered on the corresponding pial surfaces. The ASSD for this subject is shown below. Errors appear both on gyral crowns and within sulci.

Table 1: Quantitative comparison of pial surface reconstruction errors across methods. **Bold** indicates the best result per dataset.

		ASSD (mm)		HD90 (mm)	
Dataset	**Method**	mean	stdev.	mean	stdev.
UKB T_1	`recon-all-clinical` [15]	0.498	0.054	0.962	0.102
	FreeSurfer [11]	0.379	0.124	0.864	0.237
	DeepCSR [7]	0.519	0.112	2.320	0.493
	CortexODE [30]	1.030	0.099	2.246	0.217
	P3D-Fitting [36]	0.667	0.143	1.850	0.333
	`XSurfer` (Ours)	**0.303**	0.043	**0.763**	0.115
dHCP Infant	InfantFS [48]	0.930	0.185	2.650	0.683
	CortexODE [30]	1.132	0.466	2.452	0.771
	P3D-Fitting [36]	0.733	0.142	1.688	0.304
	`XSurfer` (Ours)	**0.241**	0.021	**0.502**	0.048

then ran FreeSurfer on this synthesized image using the standard surface reconstruction pipeline; (D) CortexODE [30], a point-based explicit CSR method evaluated in the same zero-shot TTO setting as **`XSurfer`**, i.e., with randomly initialized weights and optimization on each test subject only; and (E) P3D-Fitting [36], a direct stochastic gradient descent optimization using our loss function but without the untrained network prior.

4.3 Quantitative Evaluation Results

Adult Brain CSR (UKB T_1): `XSurfer` outperforms all comparison methods on the **UKB T_1** dataset, achieving the lowest ASSD of 0.303 mm and HD90 of 0.763 mm for the pial surface (Table 1, Fig. 6a). **`XSurfer`** improves upon the FreeSurfer pipeline (ASSD 0.379 mm, HD90 0.864 mm) by 20.0% and 11.7%, respectively. Furthermore,

it demonstrates a substantial advantage over learning-based baselines like DeepCSR (ASSD 0.519 mm) and standard pipelines augmented with deep learning, such as `recon-all-clinical` (ASSD 0.498 mm). Notably, the comparison to P3D-Fitting (ASSD 0.667 mm) demonstrates that, while the refined segmentation and loss function aid in CSR performance, the untrained network used in **XSurfer** aids in optimization by adding an implicit structural prior that contributes regularization.

Infant Brain CSR (dHCP): The infant brain CSR task presents greater challenges due to ongoing neurodevelopment, reduced tissue contrast, and dynamic anatomical proportions. Despite these hurdles, **XSurfer** maintains robustness, yielding an ASSD of 0.241 mm and HD90 of 0.502 mm (Table 1, Fig. 6b). This surpasses the performance of the InfantFS pipeline (ASSD 0.930 mm, HD90 2.650 mm) with a 74.1% reduction in ASSD error. **XSurfer** also outperforms explicit deformation methods like Cortex-ODE (ASSD 1.132 mm) and direct optimization via P3D-Fitting (ASSD 0.733 mm), confirming the generalizability of **XSurfer** across age and imaging contrasts.

4.4 Generalization capability of XSurfer

Table 2: Generalization of **XSurfer** across diverse datasets. ASSD and HD90 (mm) are reported for each stage of the reconstruction pipeline. **Bold** indicates the best result per dataset.

	UKB [1]		**HCP** [44]		**dHCP Infant** [10]		**dHCP Fetal** [10]		**Marmoset** [47]	
Method	ASSD	HD90	ASSD	HD90	ASSD	HD90	ASSD	HD90	ASSD	HD90
Marching Cubes	0.592	2.565	1.066	3.683	0.608	2.868	0.915	2.347	0.508	1.689
Refined MC	0.411	1.335	0.527	2.359	0.345	1.146	0.915	2.344	0.123	0.278
XSurfer (Ours)	**0.303**	**0.763**	**0.430**	**0.938**	**0.241**	**0.502**	**0.717**	**1.424**	**0.091**	**0.196**

To further evaluate the generalization of **XSurfer** across life stages, imaging modalities, and species, we evaluated CSR performance through the stages of our pipeline on five datasets. Specifically, we compare the MC surfaces reconstructed from the raw segmentation, the MC surfaces from our intermediate segmentation refinement ("Refined MC"), and the CSR of **XSurfer**.

As shown in Table 2, MC yields surface errors across all datasets due to initial segmentation inaccuracies. The Refined MC approach reduces both ASSD and HD90. The **XSurfer** prediction achieves the highest performance across all five datasets. Notably, on extreme out-of-distribution data such as the dHCP Fetal and Marmoset datasets, **XSurfer** achieves a lower ASSD (0.717 mm and 0.091 mm, respectively) than MC (0.915 mm and 0.508 mm). Furthermore, morphological measurements (curvature, thickness, and sulcal depth) derived from the reconstructed surfaces are shown in Fig. 2b.

Table 3: Ablation study on the **UKB** T_1 dataset. Each row disables or modifies one component of **`XSurfer`**. **Bold** indicates the best result per metric.

	ASSD (mm)		HD90 (mm)	
Variant	mean	stdev.	mean	stdev.
bi-Chamfer	0.416	0.058	1.146	0.291
w/o spring	0.332	0.022	**0.711**	0.041
w/o inflation	0.381	0.029	1.011	0.179
w/o CSF Recovery	0.407	0.036	1.048	0.143
`XSurfer` (w/pre-training)	0.346	0.021	0.733	0.047
`XSurfer` (pre-trained+TTO)	0.347	0.020	0.757	0.063
`XSurfer` (Full)	**0.303**	0.043	0.763	0.115

4.5 Ablation Study

We conducted an ablation study using the UKB dataset. This was performed by disabling specific components and computing the overall CSR accuracy. The results are summarized in Table 3, and qualitative comparisons (e.g., pre-trained vs. randomly initialized weights) are shown in Fig. 7. Specifically:

(1) Directional Chamfer Distance. Replacing our unidirectional Chamfer loss with a bi-directional variant degrades ASSD from 0.303 mm to 0.416 mm (a 37.3% error increase) while increasing HD90 (1.146 mm). This confirms that the unidirectional constraint prevents the optimization from being distracted by noise, ensuring more precise surface placement. (2) Inflation Loss. Omitting the inflation loss (*w/o inflation*) yields a notable performance drop, increasing ASSD to 0.381 mm (a 25.7% degradation). This verifies its necessity in penalizing self-intersections and maintaining appropriate vertex spacing, which is especially critical for capturing complex cortical folding patterns without topological errors. (3) CSF Recovery. Removing the initial segmentation refinement module (*w/o CSF Recovery*) causes ASSD to increase to 0.407 mm. This demonstrates that refining the initial label aids in correcting ambiguous tissue boundaries and ensuring accurate topological initialization before mesh deformation. (4) Spring Energy Term. Disabling the spring energy term (*w/o spring*) leads to an ASSD increase to 0.332 mm. While the HD90 metric is marginally lower without it, the overall global accuracy (ASSD) suffers. This regularization is vital for penalizing excessive edge-length variations, preserving local smoothness, and preventing spurious high-frequency artifacts that do not align with actual cortical anatomy. (5) Test-Time Optimization Strategy. We additionally compare our subject-specific optimization against two variants involving pre-training on a 100-subject training dataset: direct inference with the pre-trained model (*XSurfer w/pre-training*) and test-time optimization initialized from the pre-trained weights (*XSurfer pre-trained+TTO*). Direct use of the pre-trained model yields an inferior ASSD of 0.346 mm and HD90 of 0.733 mm, highlighting the limitation of relying only on population-level learned weights. Starting TTO from the pre-trained weights produces similar performance (ASSD = 0.347 mm, HD90 = 0.757 mm), but still does not improve over the fully subject-specific randomly initialized TTO (ASSD = 0.303 mm,

HD90 = 0.763 mm). These results suggest that individualized TTO is the key factor for adapting to subject-specific anatomy, rather than pre-training.

5 Discussion and Conclusions

We present **XSurfer**, a training-free CSR framework that employs untrained neural networks as subject-specific optimizers. By shifting from population-based learning to per-subject optimization, **XSurfer** achieves high accuracy without requiring ground-truth surfaces, demonstrating robust generalization across life stages (fetus to adult), imaging modalities (T_1, T_2, diffusion MRI), spatial resolutions (0.08–1.0 mm), and species (human and non-human primates) without domain-specific tuning. The generalization capability stems from three key properties of untrained networks: (i) architectural bias toward smooth surfaces provides implicit regularization; (ii) unified differentiable framework integrates image features and geometric constraints in a single gradient descent optimization; (iii) progressive optimization through time-variable conditioning enables coarse-to-fine refinement.

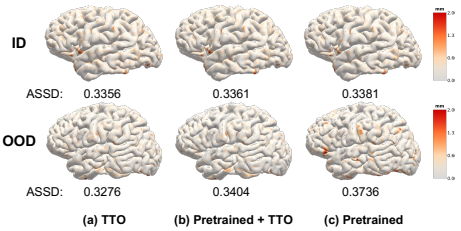

Fig. 7: Demonstration that XSurfer achieves highest performance with subject-specific optimization, i.e., without training. (a) Test-time optimization (TTO) from randomly initialized network weights. (b) TTO starting from network weights pre-trained on 100 subjects. (c) Direct inference using pre-trained weights without TTO. Rows show in-distribution (ID; subject in pre-training set) and out-of-distribution (OOD; subject not in pre-training set) cases.

Despite high performance across datasets, **XSurfer** has limitations. As our method relies on an initial tissue segmentation, the final CSR quality depends on segmentation accuracy. Although we apply CSF enhancement in deep sulci to mitigate partial-volume effects, **XSurfer** cannot fully compensate for severely inaccurate segmentations. Additionally, the current framework makes limited use of MRI intensity information beyond segmentation. Future work should explore strategies to better leverage intensity information for improved accuracy.

Acknowledgements

This work was supported by the Tsinghua University Startup Fund and the Tsinghua University Dushi Program (Grant Nos. 20241080026 and 20251080056); in part by the NIH NIBIB (grants P41-EB030006 and R01-EB019437), NINDS (grant R01-NS136385), by the *BRAIN Initiative* (NIH NINDS grants U19-NS123717 and U19-NS128613), and by the Department of Radiology at the Stanford University School of Medicine.

References

1. Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., et al.: Image processing and quality control for the first 10,000 brain imaging datasets from UK Biobank. NeuroImage **166**, 400–424 (2018)
2. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E., et al.: SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. Medical Image Analysis **86**, 102789 (2023)
3. Billot, B., Magdamo, C., Cheng, Y., Arnold, S.E., Das, S., Iglesias, J.E.: Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets. Proceedings of the National Academy of Sciences **120**(9), e2216399120 (2023). https://doi.org/10.1073/pnas.2216399120
4. Boch, M., Karadachka, K., Loh, K.K., Benn, R.A., Roumazeilles, L., Bertelsen, M.F., Manger, P.R., Wriggelsworth, E., Spiro, S., Spocter, M.A., Johnson, P.J., Avelino-de Souza, K., Patzke, N., Lamm, C., Miller, K.L., Sallet, J., Khrapitchev, A.A., Tendler, B.C., Mars, R.B.: Between-species variation in neocortical sulcal anatomy of the carnivoran brain. eLife **13**, RP100851 (jan 2026). https://doi.org/10.7554/eLife.100851
5. Bongratz, F., Rickmann, A.M., Pölsterl, S., Wachinger, C.: Vox2Cortex: Fast explicit reconstruction of cortical surfaces from 3D MRI scans with geometric deep neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20773–20783 (2022)
6. Bongratz, F., Rickmann, A.M., Wachinger, C.: Neural deformation fields for template-based reconstruction of cortical surfaces from MRI. Medical Image Analysis **93**, 103093 (2024)
7. Cruz, R.S., Lebrat, L., Bourgeat, P., Fookes, C., Fripp, J., Salvado, O.: DeepCSR: A 3D deep learning approach for cortical surface reconstruction. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 806–815 (2021)
8. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. segmentation and surface reconstruction. Neuroimage **9**(2), 179–194 (1999)
9. Dong, X., Tang, K., Hu, D., Wu, Z., Wang, L., Lin, W., Li, G.: Lifespan cortical surface reconstruction from thick-slice clinical MRI. In: Gee, J.C., Alexander, D.C., Hong, J., Iglesias, J.E., Sudre, C.H., Venkataraman, A., Golland, P., Kim, J.H., Park, J. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2025. pp. 303–313. Springer Nature Switzerland, Cham (2026)
10. Edwards, A.D., Rueckert, D., Smith, S.M., Abo Seada, S., Alansary, A., Almalbis, J., Allsop, J., Andersson, J., Arichi, T., Arulkumaran, S., et al.: The developing human connectome project neonatal data release. Frontiers in Neuroscience **16**, 886772 (2022)
11. Fischl, B.: FreeSurfer. NeuroImage **62**(2), 774–781 (2012)
12. Fortin, M.A., Kristoffersen, A.L., Larsen, M.S., Lamalle, L., Stirnberg, R., Goa, P.E.: GOUHFI: A novel contrast-and resolution-agnostic segmentation tool for ultra-high field MRI. arXiv preprint arXiv:2505.11445 (2025)
13. Fortin, M.A., Kristoffersen, A.L., Larsen, M.S., Lamalle, L., Stirnberg, R., Goa, P.E.: GOUHFI: A novel contrast- and resolution-agnostic segmentation tool for ultra-high-field MRI. Imaging Neuroscience **3**, IMAG.a.960 (10 2025). https://doi.org/10.1162/IMAG.a.960
14. Gopinath, K., Desrosiers, C., Lombaert, H.: SegRecon: Learning joint brain surface reconstruction and segmentation from images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 650–659. Springer (2021)
15. Gopinath, K., Greve, D.N., Magdamo, C., Arnold, S., Das, S., Puonti, O., Iglesias, J.E., Initiative, A.D.N., et al.: Recon-all-clinical: Cortical surface reconstruction and analysis of heterogeneous clinical brain MRI. Medical Image Analysis **103**, 103608 (2025)

16. Han, X., Pham, D.L., Tosun, D., Rettmann, M.E., Xu, C., Prince, J.L.: CRUISE: Cortical reconstruction using implicit surface evolution. NeuroImage **23**(3), 997–1012 (2004)
17. Han, X., Xu, C., Tosun, D., Prince, J.L.: Cortical surface reconstruction using a topology preserving geometric deformable model. In: Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA 2001). pp. 213–220. IEEE (2001)
18. Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M.: FastSurfer–a fast and accurate deep learning based neuroimaging pipeline. NeuroImage **219**, 117012 (2020)
19. Hoopes, A., Iglesias, J.E., Fischl, B., Greve, D., Dalca, A.V.: TopoFit: Rapid reconstruction of topologically-correct cortical surfaces. Proceedings of Machine Learning Research **172**, 508 (2022)
20. Kang, S.H., Kim, J., Cho, J., Cooley, C.Z., Heydari, A., Bilgic, B., Kim, T.H.: Enhanced partial Fourier MRI with zero-shot deep untrained priors. IEEE Access **12**, 182157–182170 (2024)
21. Kim, J.S., Singh, V., Lee, J.K., Lerch, J., Ad-Dab'bagh, Y., MacDonald, D., Lee, J.M., Kim, S.I., Evans, A.C.: Automated 3-D extraction and evaluation of the inner and outer cortical surfaces using a Laplacian map and partial volume effect classification. NeuroImage **27**(1), 210–221 (2005)
22. Lebrat, L., Santa Cruz, R., de Gournay, F., Fu, D., Bourgeat, P., Fripp, J., Fookes, C., Salvado, O.: CorticalFlow: A diffeomorphic mesh transformer network for cortical surface reconstruction. In: Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems. vol. 34, pp. 29491–29505. Curran Associates, Inc. (2021), https://proceedings.neurips.cc/paper_files/paper/2021/file/f6b5f8c32c65fee991049a55dc97d1ce-Paper.pdf
23. Lee, J., Bae, W., Ji, U.G., Yang, H., Lee, J.M.: Fetal and neonatal cortical surface reconstruction with anatomical normal-guidance and perceptual enhancements. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 7759–7768 (March 2026)
24. Lewiner, T., Lopes, H., Vieira, A.W., Tavares, G.: Efficient implementation of Marching Cubes' cases with topological guarantees. Journal of Graphics Tools **8**(2), 1–15 (2003)
25. Li, H., Liu, M., Bai, X., Liao, Y., Zheng, J., Yang, H., Li, Z., Qu, H., Tian, Q.: FetalCSR: Multi-input attention fusion network for neural ODE-based fetal cortical surface reconstruction. In: ICLR 2025 Workshop AI4CHL. Singapore, Singapore (Apr 2025), https://hal.science/hal-05039904
26. Li, H., Liu, M., Zheng, J., Yang, H., Li, Z., Tian, Q.: FetalSurfer: Automated fetal cortical surface reconstruction. In: 2024 ISMRM & ISMRT Annual Meeting & Exhibition (2024)
27. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. In: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques. p. 163–169. SIGGRAPH '87, Association for Computing Machinery, New York, NY, USA (1987). https://doi.org/10.1145/37401.37422, https://doi.org/10.1145/37401.37422
28. Ma, Q., Li, L., Kyriakopoulou, V., Hajnal, J.V., Robinson, E.C., Kainz, B., Rueckert, D.: Conditional temporal attention networks for neonatal cortical surface reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 312–322. Springer (2023)
29. Ma, Q., Li, L., Robinson, E.C., Kainz, B., Rueckert, D.: Weakly supervised learning of cortical surface reconstruction from segmentations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 766–777. Springer (2024)
30. Ma, Q., Li, L., Robinson, E.C., Kainz, B., Rueckert, D., Alansary, A.: CortexODE: Learning cortical surface reconstruction by neural ODEs. IEEE Transactions on Medical Imaging **42**(2), 430–443 (2022)

31. Madlindl, P., Bongratz, F., Wachinger, C.: Template-based cortical surface reconstruction with minimal energy deformation. In: International Workshop on Shape in Medical Imaging. pp. 140–149. Springer (2025)
32. Makropoulos, A., Robinson, E.C., Schuh, A., Wright, R., Fitzgibbon, S., Bozek, J., Counsell, S.J., Steinweg, J., Vecchiato, K., Passerat-Palmbach, J., et al.: The developing human connectome project: A minimal processing pipeline for neonatal cortical surface reconstruction. NeuroImage **173**, 88–112 (2018)
33. Pham, D.L., Bazin, P.L., Prince, J.L.: Digital topology in brain imaging. IEEE Signal Processing Magazine **27**(4), 51–59 (2010)
34. Qayyum, A., Ilahi, I., Shamshad, F., Boussaid, F., Bennamoun, M., Qadir, J.: Untrained neural network priors for inverse imaging problems: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence **45**(5), 6511–6536 (2022)
35. Rahman, S., Khan, S., Porikli, F.: A unified approach for conventional zero-shot, generalized zero-shot, and few-shot learning. IEEE Transactions on Image Processing **27**(11), 5652–5667 (2018). https://doi.org/10.1109/TIP.2018.2861573
36. Ravi, N., Reizenstein, J., Novotny, D., Gordon, T., Lo, W.Y., Johnson, J., Gkioxari, G.: Accelerating 3D deep learning with PyTorch3D. arXiv preprint arXiv:2007.08501 (2020)
37. Ren, J., Hu, Q., Wang, W., Zhang, W., Hubbard, C.S., Zhang, P., An, N., Zhou, Y., Dahmani, L., Wang, D., et al.: Fast cortical surface reconstruction from MRI using deep learning. Brain Informatics **9**(1), 6 (2022)
38. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 234–241. Springer (2015)
39. Santa Cruz, R., Lebrat, L., Fu, D., Bourgeat, P., Fripp, J., Fookes, C., Salvado, O.: CorticalFlow++: Boosting cortical surface reconstruction accuracy, regularity, and interoperability. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 496–505. Springer (2022)
40. Shattuck, D.W., Leahy, R.M.: BrainSuite: An automated cortical surface identification tool. Medical Image Analysis **6**(2), 129–142 (2002)
41. Tang, K., Dong, X., Wu, Z., Ma, L., Hung, S.C., Zhang, H., Lin, W., Li, G.: Surface-guided construction of 4D volumetric atlases of fetal brains. In: Cui, Z., Rekik, I., Suk, H.I., Ouyang, X., Sun, K., Wang, S. (eds.) Machine Learning in Medical Imaging. pp. 266–276. Springer Nature Switzerland, Cham (2026)
42. Tardif, C.L., Schäfer, A., Trampel, R., Villringer, A., Turner, R., Bazin, P.L.: Open Science CBS Neuroimaging Repository: Sharing ultra-high-field MR images of the brain. NeuroImage **124**, 1143–1148 (2016)
43. Uus, A.U., Kyriakopoulou, V., Makropoulos, A., Fukami-Gartner, A., Cromb, D., Davidson, A., Cordero-Grande, L., Price, A.N., Grigorescu, I., Williams, L.Z., et al.: BOUNTI: Brain vOlumetry and aUtomated parcellatioN for 3D feTal MRI. bioRxiv (2023)
44. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., WU-UMinn HCP Consortium: The WU-Minn human connectome project: An overview. NeuroImage **80**, 62–79 (2013)
45. Zeng, X., Puonti, O., Sayeed, A., Herisse, R., Mora, J., Evancic, K., Varadarajan, D., Balbastre, Y., Costantini, I., Scardigli, M., et al.: Segmentation of supragranular and infragranular layers in ultra-high-resolution 7T ex vivo MRI of the human cerebral cortex. Cerebral Cortex **34**(9), bhae362 (2024)
46. Zheng, H., Chen, X., Li, H., Chen, T., Liang, P., Fan, Y.: SegCSR: Weakly-supervised cortical surfaces reconstruction from brain ribbon segmentations. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–4. IEEE (2025)
47. Zhu, X., Yan, H., Zhan, Y., Feng, F., Wei, C., Yao, Y.G., Liu, C.: An anatomical and connectivity atlas of the marmoset cerebellum. Cell Reports **42**(5) (2023)

48. Zöllei, L., Iglesias, J.E., Ou, Y., Grant, P.E., Fischl, B.: Infant FreeSurfer: An automated segmentation and surface extraction pipeline for T1-weighted neuroimaging data of infants 0–2 years. NeuroImage **218**, 116946 (2020)