Sulcus-Aware Hippocampal Surface Modeling with Training-time Sulcus-guided Learning

Yuyeon Kim[1], Banseok Kim[1], Jiwon Son[1], Mingyu Joo[2], Kangmin Kim[2], Minji Ha[3,4], Minah Kim[3,4,5], and Ilwoo Lyu[1,2] (✉)

[1] Graduate School of Artificial Intelligence, Pohang University of Science and Technology, Pohang, Republic of Korea
ilwoolyu@postech.ac.kr
[2] Department of Computer Science and Engineering, Pohang University of Science and Technology, Pohang, Republic of Korea
[3] Institute of Human Behavioral Medicine, SNU-MRC, Seoul, Republic of Korea
[4] Department of Psychiatry, Seoul National University College of Medicine, Seoul, Republic of Korea
[5] Department of Neuropsychiatry, Seoul National University Hospital, Seoul, Republic of Korea

Abstract. Analyzing hippocampal structure is important for characterizing localized geometry that is poorly captured by volumetric summaries. Yet existing surface reconstruction pipelines often suppress the hippocampal sulcus—a thin internal indentation—because it is under-resolved by partial-volume effects or introduce incorrect topology. Deformation based methods have a risk of self-intersections. We present a sulcus-aware, topology-consistent surface reconstruction framework that anchors subjects to a sulcus-preserving template and learns a smooth template-to-subject deformation to transfer the surface with vertex-wise correspondence. Sulcus masks are used only in training to encourage sulcal alignment in a training-time auxiliary manner. Inference hence uses only standard structural volume and hippocampus masks without requiring sulcus masks. Experimental results show improved sulcus fidelity while maintaining competitive overall boundary agreement and stable meshes with minimal geometric failures. The resulting topology-consistent and sulcus-preserved surfaces enable surface-based sulcus morphometry; in a pilot downstream task, we show the proposed method captures sulcus-related surface statistics in cognitively normal (CN) and Alzheimer's disease (AD) cohorts. The software is available at https://github.com/Shape-Lab/Sulcus-Aware-Hippo-Surface.

Keywords: Surface Reconstruction · Hippocampal sulcus.

1 Introduction

The hippocampus supports episodic memory and spatial cognition, and its morphology is altered across many neurological and psychiatric disorders. Beyond global atrophy, hippocampal shape varies along the longitudinal axis and exhibits complex folding patterns [25], motivating morphometric methods that localize subtle, region-specific changes [15, 23, 24]. A particularly important yet technically challenging feature is the hippocampal sulcus, a thin internal cleft whose visibility and variability have been reported in clinical MRI studies of aging and Alzheimer's disease [2,18]. However, reliable sulcus delineation typically requires labor-intensive expert annotation, limiting scalability in large cohorts. Because the sulcus is narrow and highly susceptible to partial-volume effects, it is easily attenuated by downstream processing, making sulcus-preserving modeling critical for sensitive and anatomically meaningful morphometry.

Most hippocampal morphometric studies rely on volumetric measures, which detect global size differences but provide limited insight into localized geometric variation. Surface-based representations complement volumetric analysis by enabling direct assessment of local shape properties and have proven useful for capturing subtle hippocampal morphology [5, 6, 29]. However, many surface reconstruction pipelines do not retain sulcal geometry: surfaces are typically extracted from binary masks and then smoothed, which can wash out thin internal indentations including the sulcus [5, 28, 30]. As a result, sulcus-related analyses are often performed in the volumetric domain [7, 9, 22], the sulcus is frequently under-resolved at typical clinical resolution (e.g., $1\,\mathrm{mm}^3$), which can cause resolution-limited settings [26], and even when it is visible in the mask, conventional surface extraction and post-processing can still wash it out [12].

This gap reflects a trade-off between preserving fine-scale geometry including thin sulcal indentations and enforcing valid geometry. Deformation-based surface methods can represent complex folding and fine-scale variation [3, 16, 17, 27, 33], but they must be carefully designed to avoid self-intersections and preserve valid surface geometry. Harmonic-based approaches tend to provide valid geometry [28, 30], yet their reliance on mask-derived surfaces and smoothing can oversimplify anatomically meaningful details, including sulcal indentations [5, 28]. Recent learning-based mesh reconstruction methods can predict meshes end-to-end from volumes [32], but sulcus-aware extensions have received limited attention. This is likely due to the difficulty of retaining such a narrow indentation under the regularization needed for stable mesh reconstruction, as well as the limited availability of high-resolution sulcus supervision.

To address this trade-off, we present a sulcus-aware hippocampal surface modeling framework that i) preserves thin sulcal indentations while enforcing a fixed topology and ii) provides reliable inter-subject correspondence by learning a single template-based diffeomorphic deformation field that jointly warps the template volume and mesh to produce subject-specific surfaces with vertex-wise correspondence; expert-validated sulcus annotations (with a pseudo-sulcus cue) guide sulcus preservation during training, whereas inference requires only volumetric structure and hippocampus masks and obtains the subject-specific sulcus by propagating the template sulcus through the predicted deformation field, which we validate using surface-distance-based measures and demonstrate in surface-based sulcus morphometry and population-level statistical analysis with improved anatomical fidelity and statistical sensitivity.

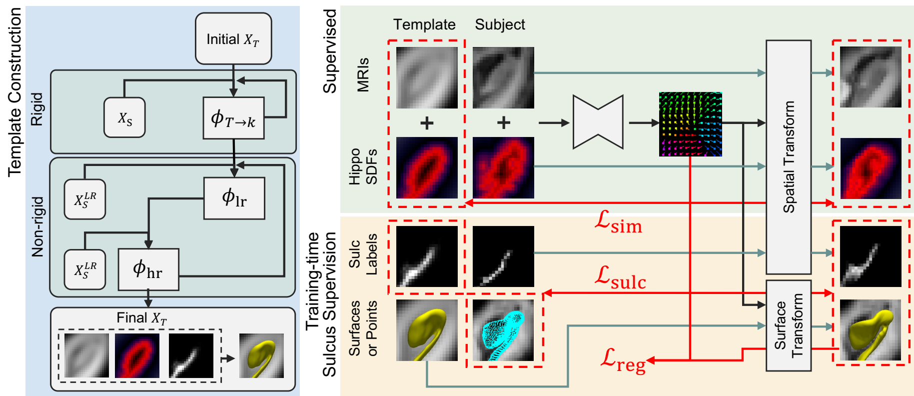

Fig. 1. Sulcus-aware reconstruction framework. Left: Build a sulcus-aware template by groupwise alignment and averaging, and extract a template surface. Right: Learn a dense template-to-subject deformation, applied to volumes and surfaces for supervised and training-time auxiliary sulcus supervision signals with regularization.

2 Method

We propose a sulcus-aware hippocampal surface reconstruction pipeline (Fig. 1). We first construct a sulcus-aware template mesh that explicitly preserves sulcal geometry. Given this, we reconstruct a subject mesh by learning a template-to-subject deformation field from inputs available at inference (volumetric structure and hippocampus mask), while using sulcus masks only during training.

2.1 Problem setting: surface reconstruction from a template

We are given a template and subject volume $X_T : \Omega_T \to \mathbb{R}^C, X_S : \Omega_S \to \mathbb{R}^C$ with C channels, and a template mesh $\mathcal{M}_T = (\mathcal{V}_T, \mathcal{F})$ with vertices $\mathcal{V}_T = \{\mathbf{v}_i^T\}_{i=1}^{N_v}$ and connectivity $\mathcal{F}$. By deforming $\mathcal{M}_T$, our goal is to reconstruct a subject mesh $\mathcal{M}_S = (\mathcal{V}_S, \mathcal{F})$, where $\mathcal{V}_S = \{\mathbf{v}_i^S\}_{i=1}^{N_v}$. We model the reconstruction via deformation field $\phi_{T\to S} : \Omega_T \to \Omega_S$ that maps template coordinates to subject's. This field induces a warping operator for any subject-space volume:

$$(X_S \circ \phi_{T\to S})(\mathbf{x}) = X_S(\phi_{T\to S}(\mathbf{x})), \quad \mathbf{x} \in \Omega_T, \tag{1}$$

and the subject mesh vertices are obtained by sampling $\phi_{T\to S}$ at $\mathcal{V}_T$:

$$\mathbf{v}_i^S = \phi_{T\to S}(\mathbf{v}_i^T), \quad i = 1, \ldots, N_v. \tag{2}$$

We model $\phi_{T\to S}$ using a neural network f_θ with learnable parameters θ:

$$\phi_{T\to S} = f_\theta(X_T, X_S). \tag{3}$$

2.2 Sulcus-aware template construction

As discussed, constructing a sulcus-aware template surface $\mathcal{M}_T$ is central to our framework. The template should capture the overall hippocampal shape while containing an explicit sulcal indentation, so that subsequent template-to-subject deformation only requires moderate local inflation/deflation. Given N subjects with structural intensity volumes $\{I_k\}$, hippocampus masks $\{Y_k\}$, and sulcus masks $\{L_k\}$, we estimate a sulcus-aware template in the volumetric domain and then extract a template mesh. We build three template channels: (i) structural intensity I_T, (ii) hippocampus soft mask Y_T, and (iii) sulcus soft mask L_T. We adopt unbiased iterative groupwise registration [20]. At iteration t, we estimate a transform $\phi^{(t)}_{T\to k}$ for each subject and update the template:

$$(I_T^{(t+1)}, Y_T^{(t+1)}, L_T^{(t+1)}) = \frac{1}{N}\sum_{k=1}^{N}(I_k \circ \phi^{(t)}_{T\to k},\, Y_k \circ \phi^{(t)}_{T\to k},\, L_k \circ \phi^{(t)}_{T\to k}). \tag{4}$$

To ensure the thin sulcus drives the alignment, we add an upweighted sulcus-channel similarity term (MSE) between $L_T^{(t)}$ and $L_k \circ \phi^{(t)}_{T\to k}$. Mathematically, ϕ defines a continuous mapping, so all channels can be compared on any common grid by resampling. However, in practice the sulcus masks are available on a much finer grid than the structural intensity volume; naively upsampling all structural channels to the sulcus resolution is memory-prohibitive. We therefore compute intensity/mask similarity on the structural grid, while evaluating the sulcus similarity on the sulcus grid by sampling ϕ there, which preserves thin sulcal signals without the cost of high-resolution volumetric warping. Following the two-stage procedure in [20], we run rigid alignment followed by non-rigid alignment. Rigid transforms are estimated from (I_k, Y_k, L_k). For non-rigid registration, we use a two-resolution cascade of [1]: a low-resolution network predicts a coarse field $\phi^{(t)}_{\mathrm{lr}}$ from the paired template/subject channels $(I_T^{(t)}, Y_T^{(t)})$ and (I_k, Y_k); this field is upsampled and refined by a high-resolution residual field $\phi^{(t)}_{\mathrm{hr}}$ driven by $L_T^{(t)}$ and $L_k \circ \phi^{(t)}_{lr}$. The final non-rigid deformation is $\phi^{(t)}_{T\to k} = \phi^{(t)}_{\mathrm{lr}} \circ \phi^{(t)}_{\mathrm{hr}}$.

After volumetric template estimation, we explicitly encode the sulcal indentation with $\widetilde{Y}_T = \max(Y_T - L_T, 0)$. We then extract an iso-surface of $\widetilde{Y}_T$ with marching cubes and apply topology fix and icosphere-based remeshing [9] to obtain the smooth, topology-ensured sulcus-aware $\mathcal{M}_T$.

2.3 Supervised deformation learning

Once $\mathcal{M}_T$ is obtained, $\mathcal{M}_S$ can be reconstructed from $\mathcal{M}_T$. To avoid self-intersections, we model $\phi_{T\to S}$ as a diffeomorphic trajectory via integrating a stationary velocity field on the volume grid. To this end, we adapt a deformation network [1] to predict the subject-specific deformation fields from inputs available at inference (MRI and hippocampus mask). We use signed distance fields (SDFs) $D : \Omega \to \mathbb{R}$ rather than binary masks because they provide continuous shape embedding and define the alignment objective directly in the continuous domain. The network receives paired template/subject voxel features

$X_T = [I_T, D(Y_T)]$ and $X_S = [I_S, D(Y_S)]$, and predicts a dense deformation $\phi_{T\to S}$. A spatial transformer [14] warps subject channels into the template domain for voxel-wise supervision.

2.4 Training-time sulcus-guided deformation learning

To attain the sulcus correctly, direct supervision of sulcus alignment is critical, but the acquisition of expert-verified sulcus masks is expensive at inference. We hence use sulcus masks only during training to define auxiliary losses (not as network inputs). For high sulcus-label resolution, computing $\phi_{T\to S}$ is expensive. To address this issue, we use two complementary training-time supervision signals: (i) a low-resolution sulcus probability map for voxel-wise soft supervision to avoid the cubic memory/runtime cost of high-resolution 3D supervision, and (ii) a boundary point cloud extracted from the sulcus masks, used only to compute surface-distance penalty against the deformed template surface.

As the first signal, we downsample the binary sulcus masks $L_S, L_T \in \{0,1\}$ by average pooling with a $3\times3\times3$ kernel and stride 3. We obtain the soft sulcus occupancy maps $P_S, P_T \in [0,1]$ on a common low-resolution grid, enabling direct voxel-wise comparison without upsampling. As the second signal, we extract a boundary point cloud from the high-resolution masks for geometric supervision against the deformed template surface. We form a sulcus-aware mask $\widetilde{Y}_k = \max(Y_k - L_k, 0)$, sample points in a narrow band from its SDF $\{\mathbf{z}_j : |D(\mathbf{z}_j)| < \varepsilon\}$, and project them onto the zero level set as $\mathbf{b}_j = \mathbf{z}_j - \frac{D(\mathbf{z}_j)}{\|\nabla D(\mathbf{z}_j)\|_2^2 + \eta}\nabla D(\mathbf{z}_j)$, yielding stable boundary points $\mathcal{B}_k = \{\mathbf{b}_j\}_{j=1}^{N_B}$, where ε and η are small constants.

2.5 Learning objectives

We optimize a weighted sum of similarity, sulcus, and regularization:

$$\mathcal{L} = \omega_{\text{sim}}\mathcal{L}_{\text{sim}} + \omega_{\text{sulc}}\mathcal{L}_{\text{sulc}} + \mathcal{L}_{\text{reg}}, \tag{5}$$

where ω_{sim} and ω_{sim} are scalar weights. The supervised term $\mathcal{L}_{\text{sulc}}$ uses image- and shape-based alignment signals available at inference. We use an intensity similarity term (NCC) [1] and an SDF-based shape term (weighted MAE):

$$\mathcal{L}_{\text{sim}} = \text{NCC}(I_T,\, I_S \circ \phi_{T\to S}) + \lambda_{\text{sim}}\ \text{MAE}_{\bar{w}}(D(Y_T),\, D(Y_S) \circ \phi_{T\to S})\,, \tag{6}$$

where $\bar{w}(\mathbf{x}) = \frac{w(\mathbf{x})}{\sum_{\mathbf{x}} w(\mathbf{x})}$ and $w(\mathbf{x}) = \exp(-|D(Y_T)(\mathbf{x})|/\sigma)$, where σ is constant, emphasizes voxels near the template surface. The sulcus term $\mathcal{L}_{\text{sulc}}$ encourages sulcus consistency by NCC and symmetric Chamfer distance [8]:

$$\mathcal{L}_{\text{sulc}} = \text{NCC}(P_T,\, P_S \circ \phi_{T\to S}) + \lambda_{\text{sulc}}\,\text{CF}\Big(\mathcal{B}_S, \{\phi_{T\to S}(\mathbf{v}_i^T)\}_{i=1}^{N_v}\Big)\,. \tag{7}$$

Finally, $\mathcal{L}_{\text{reg}}$ includes a deformation-field smoothness term [1] and mesh regularizers that discourage edge-length/area distortion and promote normal consistency, with reduced normal penalties in high-curvature regions such as the sulcus. To keep optimization balanced, we rescale each constituent loss with fixed normalization constants using $\lambda_{sim}, \lambda_{sulc}$ so that their magnitudes remain comparable throughout training.

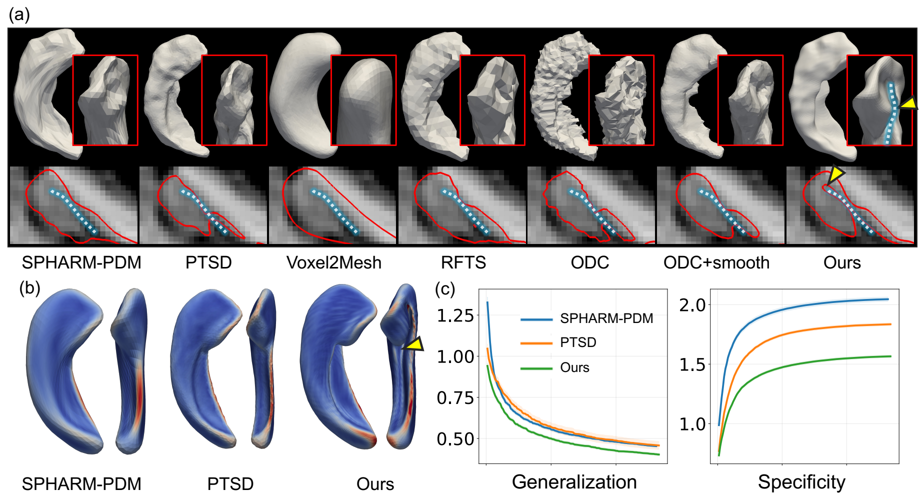

Fig. 2. Quantitative and correspondence results. (a) Representative reconstructions; our method better preserves subject-specific sulcal folding (yellow arrows) and local curvature than baselines. (b) Mean surfaces colored by Koenderink's curvedness. (c) PCA-based correspondence evaluation using generalization and specificity following Styner et al. [31]; lower values indicate better shape-modeling performance.

3 Experiments and Results

3.1 Dataset and implementation details

We used T1-weighted structural MRI scans from [21] acquired at 1 mm isotropic resolution (68 at 1.5T and 67 at 3T), each paired with an expert manual hippocampus mask following [4]. Hippocampal sulcus masks were initialized using the adaptive segmentation tool of [13] to obtain 0.33 mm isotropic sulcus masks, which were subsequently refined and verified by our hippocampal experts. For comparison with baseline methods, we used hippocampus-specific toolboxes, SPHARM-PDM [30] and PTSD [16]. We further included generic deformable-surface [16, 32] and level-set–based [12, 26] reconstruction methods.

For each subject, we derived a soft ROI mask by dilating and smoothing the hippocampus mask, and applied it to the T1 scan. The ROI-masked scan and mask were then rigidly registered to the sulcus-aware template space (Sec. 2.2), resampled on the $64 \times 64 \times 64$ voxels template grid. All $\mathcal{M}_T$ was made only by training set. We sampled $\mathcal{B}_S$ in $40,000$ points. We trained the proposed model on 120 out of 135 subjects (each contains left and right hemispheres). The deformation network followed the default U-Net architecture of VoxelMorph [1], with encoder filters $[16, 32, 32, 32]$ and decoder filters $[32, 32, 32, 32, 32, 16, 16]$. We trained the network on 120 out of 135 subjects (each contains left and right hemispheres) with Adam (learning rate 10^{-4}, batch size 8) for 100 epochs. We set loss weights to $\omega_{\text{sup}} = 1$ and $\omega_{\text{sulc}} = 5$. All quantitative results were computed

Table 1. Quantitative comparison of surface reconstruction methods. We reported mean±std boundary-to-surface distances from hippocampal and sulcus masks (mm), self-intersections (count/mean intersect area ratio for each mesh), topology errors (non–genus-0 count), and reconstruction failures.

Method	Hippo dist. ↓	Sulc dist. ↓	Self-inter.	Topo err	Recon err
SPHARM-PDM [30]	0.89±0.07	1.62±0.29	–	–	1
PTSD [16]	0.64±0.03	1.06±0.23	9/10.9	–	–
Voxel2Mesh [32]	1.42±0.27	2.35±0.34	2/2.39	–	–
RFTS [26]	0.90±0.03	1.39±0.23	1/0.32	–	–
ODC [12]	**0.43**±0.01	0.97±0.18	–	14	–
ODC+smooth [12]	0.49±0.02	1.09±0.22	–	9	–
Ours	0.68±0.05	**0.66**±0.17	1/0.01	–	–

on the held-out 15 subjects × 2 hemispheric meshes. The same dataset was used to validate the all baseline methods.

3.2 Experiments and results

Qualitative results. Visual inspection showed that our reconstructed surface follows subject-specific sulcal patterns more faithfully, producing sulcus-consistent folding instead of overly smoothed geometry. Compared to baseline methods, the resulting meshes better captured fine-scale hippocampal curvature and local indentations (Fig. 2(a)). In addition, we compared our method with those that explicitly provide correspondences during reconstruction. We computed a mean surface across subjects after the rigid Procrustes alignment (Fig. 2(b)). We observed that our mean shape preserved sharper anatomical details with sulcal indentations, whereas competing correspondence-aware baselines yielded noticeably smoother averages.

Quantitative results. We measured geometric accuracy using closest-point distances from anatomical boundary points of a volumetric mask $\mathbf{b}$ to $\mathcal{M}_S$. Hippocampus and sulcus boundary points were extracted from the hippocampus mask [4] and sulcus mask boundary points were extracted from our sulcus masks. In Table 1, boundary differences in most methods were below 1 mm. However, from a sulcal reconstruction view, even such small gaps could cause opposing sulcal walls to adhere to each other (Fig. 2). Meanwhile, all the methods exhibited at least one failure in reconstruction, self-intersection, or topology. SPHARM-PDM [30] may suppress high-frequency geometry and are sensitive to mask-topology violations, deformation-based approaches [16,32] tend to produce self-intersections, and level set-based approaches [12, 26] are often resolution-limited and topology-ambiguous, yielding artifact-prone surfaces. Our method also exhibited a single self-intersection case, which is highly localized and has negligible area (Fig. 3(a)). It was noteworthy that this behavior was consistent with discretization and interpolation effects arising from voxel-grid SVF integration [19].

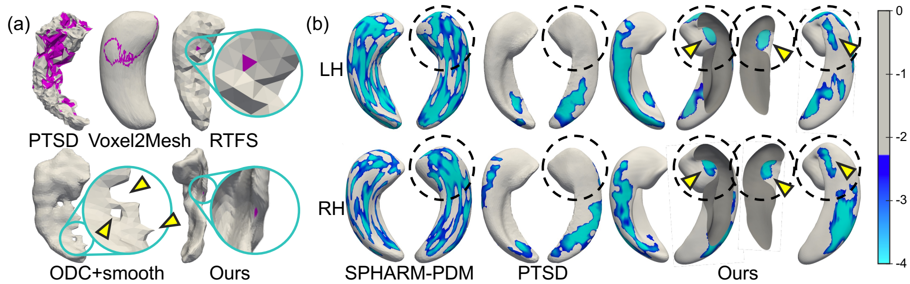

Fig. 3. Failure cases and statistical analysis. (a) Worst case examples of self-intersection and topology error. (b) AD–CN surface-based group analysis on per-vertex area, shown as cluster-wise corrected t-value maps on LH/RH surfaces ($p < 0.005$).

We further evaluated cross-subject correspondence against SPHARM-PDM [30] and PTSD [16]. Voxel2Mesh [32] and level-set–based methods [12, 26] were excluded because they do not provide inter-subject vertex-wise correspondence during surface reconstruction. Due to the absence of ground-truth correspondences, we evaluated correspondence quality through the shape modeling criteria of Styner *et al.* [31], focusing on *generalization* and *specificity* (Fig. 2(c)). Briefly, generalization measures how well a shape model built from the established correspondences can reconstruct unseen subjects. Specificity measures whether shapes sampled from the model remain anatomically plausible and close to the observed population. Both metrics capture two complementary types of reconstruction errors in the PCA space. Our method achieved the best performance, indicating tighter and more consistent correspondences across subjects with stronger suitability for statistical shape modeling.

Clinical study. To maximize data usage in the clinical study, we used five-fold cross-validation and a held-out test set (unseen in all folds), and aggregated out-of-fold predictions for 120 subjects. We compared AD–CN vertex-wise surface area using a FreeSurfer [10] and correspondence-aware baselines [16, 30], which is used in 3.2, because correspondence is required for vertex-wise area statistics. To analyze the same cohort across methods, we excluded subjects that failed the quantitative evaluation (Fig. 3), removing 10 AD cases (2 reconstruction failures in [30]; 9 self-intersections in [16] with self-intersection area ratio $> 5\%$). This yielded 30 AD (13F/17M; age 72.8 ± 6.6 years; eTIV $1.53 \times 10^6 \pm 1.94 \times 10^5\,\mathrm{mm}^3$) and 39 CN (19F/20M; age 76.8 ± 7.5 years; eTIV $1.50 \times 10^6 \pm 1.83 \times 10^5\,\mathrm{mm}^3$) participants. Group differences were tested with cluster-wise inference (covariates: sex, age, eTIV; cluster-forming threshold $p < 0.005$). Fig. 3 shows similar clusters in the hippocampal tail, but stronger method dependence in the head and sulcal indentation; our method produced more sulcus-localized clusters than both baselines. Similar sulcus-related alterations in AD, including hippocampal sulcus widening, have been reported in [2, 11, 18]. It is noted that these differences require further validation due to the absence of ground truth. However, they indicate that reconstruction and correspondence choices can affect downstream surface-based statistics particularly within the hippocampal sulcus.

4 Conclusion

We presented a sulcus-aware hippocampal surface modeling framework that preserves thin sulcal indentations while maintaining fixed topology and consistent correspondence. We build a sulcus-aware template and learn a single template-to-subject deformation that jointly warps volumetric features and the mesh, using sulcus cues only for template construction and training. Experiments show competitive hippocampal boundary agreement with improved sulcus fidelity and fewer self-intersections than baselines, enabling more reliable surface-based sulcus morphometry. Future work will scale to broader datasets and imaging protocols and strengthen geometric guarantees under partial-volume effects.

Acknowledgments. This work was supported in part by the National Research Foundation of Korea (NRF) under RS-2024-00333931, RS-2025-02216257, RS-2026-25470841, and in part by the Institute for Information & Communications Technology Planning & Evaluation (IITP) under RS-2019-II191906.

Disclosure of Interests. The authors have no competing interests in the paper.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. IEEE transactions on medical imaging **38**(8), 1788–1800 (2019)
2. Bastos-Leite, A., Van Waesberghe, J., Oen, A., Van Der Flier, W., Scheltens, P., Barkhof, F.: Hippocampal sulcus width and cavities: comparison between patients with alzheimer disease and nondemented elderly subjects. American journal of neuroradiology **27**(10), 2141–2145 (2006)
3. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International journal of computer vision **61**, 139–157 (2005)
4. Boccardi, M., Bocchetta, M., Morency, F.C., Collins, D.L., Nishikawa, M., Ganzola, R., Grothe, M.J., Wolf, D., Redolfi, A., Pievani, M., et al.: Training labels for hippocampal segmentation based on the eadc-adni harmonized hippocampal protocol. Alzheimer's & Dementia **11**(2), 175–183 (2015)
5. Chang, C., Huang, C., Zhou, N., Li, S.X., Ver Hoef, L., Gao, Y.: The bumps under the hippocampus. Human brain mapping **39**(1), 472–490 (2018)
6. DeKraker, J., Köhler, S., Khan, A.R.: Surface-based hippocampal subfield segmentation. Trends in neurosciences **44**(11), 856–863 (2021)
7. DeKraker, J., Palomero-Gallagher, N., Kedo, O., Ladbon-Bernasconi, N., Muenzing, S.E., Axer, M., Amunts, K., Khan, A.R., Bernhardt, B.C., Evans, A.C.: Evaluation of surface-based hippocampal registration using ground-truth subfield definitions. ELife **12**, RP88404 (2023)

8. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 605–613 (2017)
9. Fischl, B.: Freesurfer. NeuroImage **62**(2), 774–781 (2012)
10. FreeSurfer Wiki: Surface-Based Group Analysis (FsTutorial/GroupAnalysisV6.0). https://surfer.nmr.mgh.harvard.edu/fswiki/FsTutorial/GroupAnalysisV6.0 (2026), last modified: 2026-02-06; Accessed: 2026-02-25
11. Hekmatnia, A., Basiratnia, R., Koohi, R., Barekatein, M., Ahrar, H., Hekmatnia, F., Ghazavi, A.: Prevalence of hippocampal morphologic variants between healthy elderly subjects and patients with alzheimer's disease. Advanced Biomedical Research **3**(1), 59 (2014)
12. Hwang, J., Sung, M.: Occupancy-based dual contouring. In: SIGGRAPH Asia 2024 Conference Papers (2024)
13. Iglesias, J.E., Augustinack, J.C., Nguyen, K., Player, C.M., Player, A., Wright, M., Roy, N., Frosch, M.P., McKee, A.C., Wald, L.L., et al.: A computational atlas of the hippocampal formation using ex vivo, ultra-high resolution mri: Application to adaptive segmentation of in vivo mri. Neuroimage **115**, 117–137 (2015)
14. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks. In: Advances in Neural Information Processing Systems (NeurIPS) (2015)
15. Joseph, J., Warton, C., Jacobson, S.W., Jacobson, J.L., Molteno, C.D., Eicher, A., Marais, P., Phillips, O.R., Narr, K.L., Meintjes, E.M.: Three-dimensional surface deformation-based shape analysis of hippocampus and caudate nucleus in children with fetal alcohol spectrum disorders. Human brain mapping **35**(2), 659–672 (2014)
16. Kim, J., Valdes-Hernandez, M.d.C., Royle, N.A., Park, J.: Hippocampal shape modeling based on a progressive template surface deformation and its verification. IEEE transactions on medical imaging **34**(6), 1242–1261 (2014)
17. Lee, J.M., Kim, S.H., Jang, D.P., Ha, T.H., Kim, J.J., Kim, I.Y., Kwon, J.S., Kim, S.I.: Deformable model with surface registration for hippocampal shape deformity analysis in schizophrenia. NeuroImage **22**(2), 831–840 (2004)
18. Li, Y., Li, J., Segal, S., Wegiel, J., De Santi, S., Zhan, J., De Leon, M.: Hippocampal cerebrospinal fluid spaces on mr imaging: Relationship to aging and alzheimer disease. American journal of neuroradiology **27**(4), 912–918 (2006)
19. Liu, Y., Chen, J., Wei, S., Carass, A., Prince, J.: On finite difference jacobian computation in deformable image registration. International journal of computer vision **132**(9), 3678–3688 (2024)
20. Lyttelton, O., Boucher, M., Robbins, S., Evans, A.: An unbiased iterative group registration template for cortical surface analysis. Neuroimage **34**(4), 1535–1544 (2007)
21. Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L.: The alzheimer's disease neuroimaging initiative. Neuroimaging Clinics **15**(4), 869–877 (2005)
22. Patenaude, B., Smith, S.M., Kennedy, D.N., Jenkinson, M.: A bayesian model of shape and appearance for subcortical brain segmentation. NeuroImage **56**(3), 907–922 (2011)
23. Qiu, A., Miller, M.I.: Multi-structure network shape analysis via normal surface momentum maps. NeuroImage **42**(4), 1430–1438 (2008)
24. Raamana, P.R., Rosen, H., Miller, B., Weiner, M.W., Wang, L., Beg, M.F.: Three-class differential diagnosis among alzheimer disease, frontotemporal dementia, and controls. Frontiers in neurology **5**, 71 (2014)

25. Ropireddy, D., Bachus, S.E., Ascoli, G.A.: Non-homogeneous stereological properties of the rat hippocampus from high-resolution 3d serial reconstruction of thin histological sections. Neuroscience **205**, 91–111 (2012)
26. Sellán, S., Batty, C., Stein, O.: Reach for the spheres: Tangency-aware surface reconstruction of sdfs. In: SIGGRAPH Asia 2023 conference papers. pp. 1–11 (2023)
27. Shen, D., Moffat, S., Resnick, S.M., Davatzikos, C.: Measuring size and shape of the hippocampus in mr images using a deformable shape model. Neuroimage **15**(2), 422–434 (2002)
28. Shi, Y., Lai, R., Morra, J.H., Dinov, I., Thompson, P.M., Toga, A.W.: Robust surface reconstruction via laplace-beltrami eigen-projection and boundary deformation. IEEE transactions on medical imaging **29**(12), 2009–2022 (2010)
29. Shi, Y., Thompson, P.M., de Zubicaray, G.I., Rose, S.E., Tu, Z., Dinov, I., Toga, A.W.: Direct mapping of hippocampal surfaces with intrinsic shape context. NeuroImage **37**(3), 792–807 (2007)
30. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J.J., Shenton, M.E., Gerig, G.: Framework for the statistical shape analysis of brain structures using spharm-pdm. The insight journal (1071), 242 (2006)
31. Styner, M.A., Rajamani, K.T., Nolte, L.P., Zsemlye, G., Székely, G., Taylor, C.J., Davies, R.H.: Evaluation of 3d correspondence methods for model building. In: Biennial international conference on information processing in medical imaging. pp. 63–75. Springer (2003)
32. Wickramasinghe, U., Remelli, E., Knott, G., Fua, P.: Voxel2mesh: 3d mesh model generation from volumetric data. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 299–308. Springer (2020)
33. Zhao, J., Liu, S., Ahmad, S., Yap, P.T.: Meshdeform: Surface reconstruction of subcortical structures via human brain mri. International Conference on Information Processing in Medical Imaging pp. 536–547 (2023)