A Versatile Prompt-Driven Framework for Brain Segmentation, Parcellation, and Surface Reconstruction from Diverse Modalities

Zifeng Lian[1], Jiameng Liu[1], Shui Cao[1], Zhichao Liang[1,2], Feng Shi[3], and Dinggang Shen[1,3,4(✉)]

[1] School of Biomedical Engineering & State Key Laboratory of Advanced Medical Materials and Devices, ShanghaiTech University, Shanghai 201210, China
[2] Faculty of Applied Sciences, Macao Polytechnic University, Macao 999078, China
[3] Shanghai United Imaging Intelligence Co., Ltd., Shanghai, China
[4] Shanghai Clinical Research and Trial Center, Shanghai 201210, China
Dinggang.Shen@gmail.com

Abstract. Accurate and comprehensive characterization of brain anatomy across volumetric and surface domains is fundamental to neuroimaging research and clinical practice. Existing methods, however, are often modality-specific and task-restricted, requiring separate models for segmentation, parcellation, or surface reconstruction, thus failing to leverage the intrinsic anatomical complementarity across representations and limiting their utility in heterogeneous clinical settings. In this work, we introduce uBrain, a prompt-driven unified framework that performs simultaneous volumetric segmentation and surface reconstruction within a single model, adaptable to diverse imaging modalities (*i.e.*, T1W, T2W, FLAIR, PD, and CT). Central to our framework is a guided dynamic routing mechanism conditioned on modality- and task-specific prompts, enabling adaptive feature learning and task-oriented decoding across diverse acquisitions. To further enhance generalization under domain shift, we propose a Gaussian Mixture Model-based data synthesis strategy that simulates heterogeneous imaging characteristics. Extensive evaluations on a large-scale cohort of 7,662 subjects demonstrate that uBrain achieves state-of-the-art performance, attaining an average Dice score of 93.21% for volumetric segmentation and an average symmetric surface distance of 0.23 mm for surface reconstruction on unseen external datasets. Moreover, our method maintains robust performance on challenging acquisitions (including thick-slice and ultra-high-field imaging), underscoring its promise for reliable deployment in real-world clinical environments.

Keywords: Prompt-driven · Multimodal framework · Multitask learning · Volumetric segmentation · Surface reconstruction

1 Introduction

Accurate brain tissue segmentation, region-of-interest (ROI) parcellation, and surface reconstruction form the methodological foundation of neuroimaging, en-

abling quantitative characterization of brain morphology and systematic mapping of neurological disorders [1–3]. Traditionally, these tasks are tackled separately using modality-specific and task-restricted models. For instance, researchers typically train separate networks for volumetric segmentation and surface reconstruction, often restricting the input to a single high-resolution modality, such as T1-weighted (T1W) MRI [4–6]. This fragmented paradigm not only overlooks complementary anatomical information across imaging contrasts and structural representations but also leads to a proliferation of models that may hinder efficient clinical deployment.

Considerable efforts have driven the evolution of brain analysis from traditional automated pipelines [7, 8] to advanced learning-based architectures [9–16]. Recent joint methods such as Vox2Cortex [17] and SegRecon [18] couple volumetric and surface representations, but are mainly developed for T1W MRI-based cortical analysis. In parallel, recent Mixture-of-Modality-Experts (MoME) architectures mainly focus on lesion segmentation [19] or multimodal survival prediction [20], while contrast-agnostic methods such as SynthSeg [21] and recon-all-clinical [22] enhance robustness to heterogeneous scans. Despite these advances in architectural refinement and cross-modal generalization, existing methods still address contrast variability and task optimization in isolation. Therefore, there remains a critical need for a unified framework that handles multimodal inputs and multitask outputs while seamlessly bridging volumetric segmentation and surface reconstruction.

In this study, we propose uBrain, a versatile prompt-driven unified framework designed for simultaneous tissue segmentation, ROI parcellation, and surface reconstruction within a single model, adaptable to diverse imaging modalities including T1W, T2-weighted (T2W), fluid-attenuated inversion recovery (FLAIR), proton-density-weighted (PD), and computed tomography (CT). The main contributions of this work are summarized as follows:

- **A Unified Multimodal and Multitask Framework:** We introduce a single-model solution that integrates volumetric brain segmentation, parcellation, and surface reconstruction. This unified design reduces computational overhead while maximizing shared anatomical feature learning across diverse representations.
- **Dynamic Modality and Task Routing:** We propose a novel dual-prompt conditioning mechanism featuring a MoME and an Automatic Pathway (AP) Module. This architecture enables the network to dynamically adapt its computational graph based on the specific input modality and the desired output task.
- **Robustness to Heterogeneous Acquisitions:** By incorporating a Gaussian Mixture Model (GMM)-based data synthesis module during training, uBrain achieves state-of-the-art (SOTA) performance on a large-scale cohort of 7,662 subjects across all tasks. Furthermore, it demonstrates strong generalization on challenging, out-of-distribution acquisitions, including thick-slice clinical scans and ultra-high-field imaging.

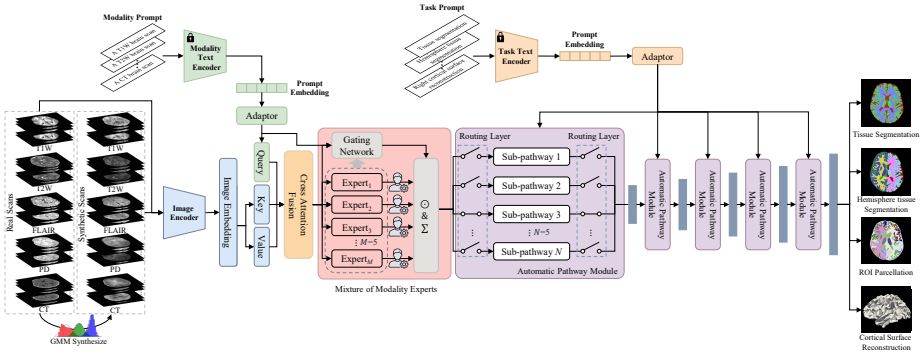

Fig. 1. Overview of the proposed uBrain framework, comprising three core components: (1) a structure-preserving GMM-based data augmentation module, (2) a Modality-Adaptive Encoder, and (3) a Task-Driven Decoder.

2 Methods

2.1 Structure-Preserving GMM Data Augmentation

To improve robustness to heterogeneous data, we design a domain-randomization-based augmentation module that synthesizes realistic training samples with varied contrasts and resolutions [21].

For MRI data synthesis, unlike standard label-to-image translations [21], we use structure-preserving intensity re-mapping to maintain anatomical textures. Specifically, given a source image $\mathcal{I}_{src}$ and its segmentation map $\mathcal{L}$, we model the intensity distribution of each tissue class k as $\mathcal{N}(\mu_{k,src}, \sigma^2_{k,src})$ within a GMM. During training, target statistics $\mathcal{N}(\mu_{k,tgt}, \sigma^2_{k,tgt})$ are randomly sampled from a Look-Up Table, whose entries are class- and modality-wise means and standard deviations computed only from training scans. For all voxels with $\mathcal{L}(v) = k$, the augmented intensity is computed as $\mathcal{I}_{aug}(v) = (\mathcal{I}_{src}(v) - \mu_{k,src})\frac{\sigma_{k,tgt}}{\sigma_{k,src}} + \mu_{k,tgt}$, preserving textures while simulating contrast variations.

For CT data synthesis, where soft-tissue texture is less informative than in MRI, we directly sample voxel intensities from the target distribution to fill anatomical regions.

Following SynthSeg [21], we further apply random bias fields, Gamma transformations, and anisotropic downsampling followed by upsampling to simulate intensity inhomogeneity, nonlinear intensity variations, and thick-slice partial-volume effects.

Together, these structure-preserving synthesis steps enable uBrain to generalize effectively to heterogeneous data without requiring retraining.

2.2 uBrain Architecture

As illustrated in Fig. 1, uBrain is a prompt-driven unified framework for multimodal brain volumetric segmentation (tissue, hemisphere, ROIs) and surface

Table 1. Summary of internal and external multimodal datasets, including the number of scans and the distributions of age and gender.

Source	Dataset	Modality	Number of scans	Age range (min-max, years)	Age distribution (mean ± std, years)	Gender (male/female)
Internal dataset	ADHD	T1W	952	7-26	12.53±3.26	836/116
	HCPY	T2W	1063	22-37	28.89±3.62	485/578
	CBMFM	FLAIR	835	14-85	48.63±18.61	400/435
	IXI	PD	532	20-86	48.67±16.47	234/298
	HuaShan	CT	627	28-89	68.62±6.19	254/373
External dataset	ABIDE	T1W	971	6-64	17.44±8.17	599/372
	ABCD	T2W	485	8-15	9.97±2.89	255/230
	UKBiobank	FLAIR	1355	40-68	46.20±11.93	693/662
	AIBL	PD	772	55-92	69.79±7.47	382/390
	ZhongShan	CT	70	50-88	64.57±7.76	30/40

reconstruction (implicit SDF of bilateral white matter (WM)) across T1W, T2W, FLAIR, PD, and CT. We adopt SwinUNETR [13] as the backbone of uBrain. The architecture comprises three core components: (1) a structure-preserving GMM augmentation module that synthesizes realistic training samples to improve robustness to heterogeneous acquisitions; (2) a Modality-Adaptive Encoder (MAE), which encodes modality prompts with a pretrained text encoder and fuses the resulting embeddings with image features via cross-attention, followed by a MoME module for modality-adaptive expert aggregation; and (3) a Task-Driven Decoder (TDD), composed of AP modules that use task prompts to activate suitable sub-pathways for diverse task-specific outputs.

For each input, uBrain uses controlled prompts from predefined sets of modality and task prompts, e.g., "A T1W brain scan" and "ROI parcellation", rather than unrestricted free-form instructions. The modality prompt is encoded by the pretrained text encoder and fused with bottleneck image features via cross-attention. In MoME, a gating network uses the global image context and modality embedding to produce softmax weights over five expert blocks; their weighted sum forms the modality-adaptive representation.

The proposed TDD consists of five AP modules, each with two routing layers and five parallel sub-pathways. For each AP module l, the task prompt embedding p_{task} is fed into a linear router to produce logits $z_l = W_l p_{\text{task}} + b_l$ over the five sub-pathways, and $\text{softmax}(z_l)$ gives the activation probabilities. During training, hard Gumbel-Softmax sampling with $\tau = 1.0$ generates a one-hot path selection while preserving differentiability through the straight-through estimator. During inference, the sub-pathway with the maximum probability is selected. Thus, AP modules perform hard task-specific routing, using the probabilities for path selection rather than soft mixing across all paths.

2.3 Implementation Details

Before training uBrain, the pretrained modality and task text encoders are independently finetuned from the PubMedBERT-based text encoder in Biomed-

Table 2. Quantitative comparison with SynthSeg for volumetric segmentation and Recon-all-clinical for surface reconstruction on both internal and external datasets. Values represent the task-averaged performance for each modality.

Source	Modality	Volumetric segmentation (Dice score [%] ↑)		Surface reconstruction (ASSD [mm] ↓)	
		SynthSeg	uBrain	Recon-all-clinical	uBrain
Internal dataset	T1W	91.29±4.12	**96.41±1.83**	0.19±0.06	**0.15±0.05**
	T2W	90.14±4.55	**95.11±1.51**	0.20±0.07	**0.16±0.02**
	FLAIR	88.43±5.18	**94.01±1.97**	0.25±0.09	**0.20±0.03**
	PD	86.12±5.66	**91.73±2.49**	0.31±0.11	**0.26±0.05**
	CT	80.42±6.94	**85.86±4.40**	0.88±0.35	**0.47±0.12**
External dataset	T1W	90.90±4.22	**95.62±2.11**	0.18±0.06	**0.16±0.06**
	T2W	90.65±4.32	**94.28±1.93**	0.22±0.08	**0.18±0.02**
	FLAIR	87.86±5.30	**92.89±2.42**	0.25±0.08	**0.23±0.05**
	PD	86.57±5.48	**90.16±2.84**	0.33±0.12	**0.29±0.06**
	CT	80.83±6.15	**83.45±5.20**	0.91±0.38	**0.52±0.14**

CLIP [24] using LoRA [25], and then frozen during uBrain training. The network is trained for 100 epochs using the AdamW optimizer with an initial learning rate of 10^{-4}. To optimize diverse task outputs, we employ a combination of Dice and focal losses [23] to mitigate class imbalance in volumetric segmentation, and use MSE loss to supervise implicit SDF learning for surface reconstruction. For training efficiency, we use random patch cropping with a patch size of 128^3. All experiments are conducted on an NVIDIA A100 GPU (80 GB memory). We extract WM surfaces from the predicted SDFs using Marching Cubes [26], and apply a fast topology correction [16] to ensure the topological correctness of the reconstructed WM surfaces.

3 Experiments and Results

3.1 Dataset

To rigorously validate our framework, we curated a large-scale multimodal dataset comprising 7,662 subjects sourced from 7 publicly available cohorts and 3 in-house cohorts (details in Table 1). We randomly split the internal datasets subject-wise into training (80%), validation (10%), and testing (10%) sets. Ground-truth volumetric labels were initialized using [27] and then underwent visual quality control (QC) and correction by trained raters using ITK-SNAP, with cases showing severe segmentation errors, abnormal intensity distributions, or artifacts excluded. The final QC-approved labels were reviewed by experienced radiologists. We applied a distance transform [28] to these labels to obtain the SDFs and clipped them at an absolute value of 5 mm. The ground-truth WM surfaces were derived from the same QC-approved labels following [29]. All scans were rigidly registered to the UNC-BCP 4D Infant Brain Volumetric Atlas [30] to standardize their orientation and position, followed by z-score intensity normalization.

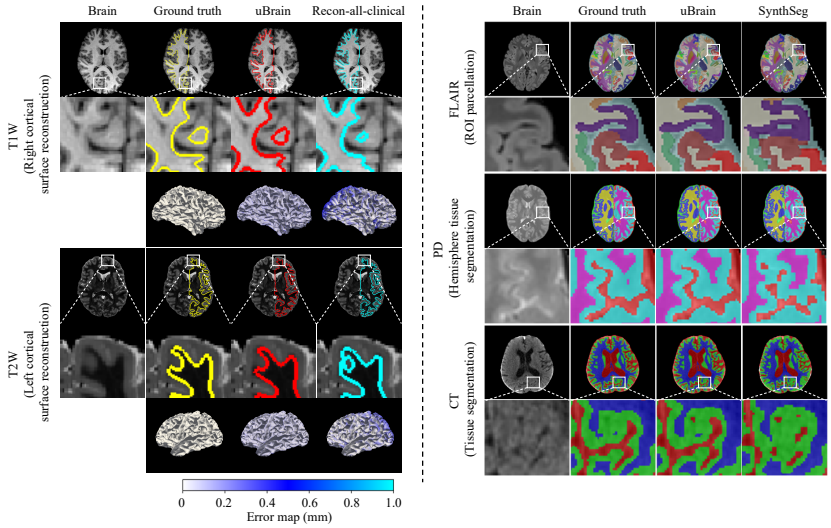

Fig. 2. Qualitative comparisons on the external datasets. White boxes highlight specific regions where uBrain preserves accurate boundaries and anatomical details.

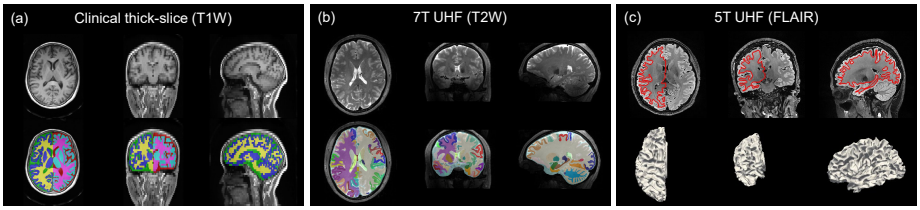

Fig. 3. Qualitative evaluation of uBrain on (a) hemisphere tissue segmentation from a clinical thick-slice T1W scan (5 mm thickness), (b) ROI parcellation from a 7T UHF T2W scan, and (c) left WM surface reconstruction from a 5T UHF FLAIR scan.

3.2 Comprehensive Evaluation of uBrain

Comparison with SOTA Methods We evaluate uBrain against SynthSeg [21] for volumetric segmentation and Recon-all-clinical [22] for surface reconstruction. These methods are strong SOTA baselines for analyzing brain scans with arbitrary contrast and resolution.

Table 2 presents the task-averaged quantitative performance for each modality, while Fig. 2 shows qualitative results on the external datasets. uBrain consistently achieves the highest Dice scores and the lowest average symmetric surface distance (ASSD) values across all modalities on both internal and external datasets. Compared with competing methods, uBrain better preserves fine-grained anatomical details and reduces boundary artifacts, producing segmentation maps and cortical surfaces consistent with the ground truth. Notably, on the challenging CT modality where other approaches suffer significant degradation,

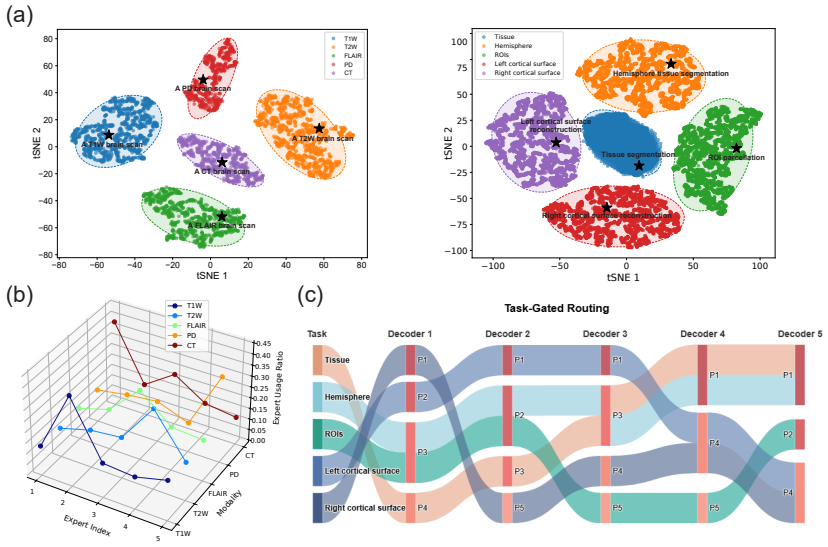

Fig. 4. (a) t-SNE visualization from the pretrained modality (left) and task (right) text encoders. Scatter points represent embeddings of multimodal inputs or task target maps from the internal test set. Black stars denote specific modality or task prompts, while dashed ellipses indicate the 95% CI for each cluster. (b) Expert usage ratios for each modality in the MoME module. (c) Task-gated routing visualization of the five AP modules, with each path representing a sub-pathway activated for specific tasks.

uBrain maintains robust accuracy. This superiority is primarily attributed to our MAE, which utilizes modality prompts to disentangle modality-specific features. Furthermore, the MoME enables the model to leverage rich structural representations learned from high-contrast MRI modalities, effectively compensating for the low soft-tissue contrast inherent in CT data. These consistent quantitative and qualitative results demonstrate the effectiveness of our unified design.

Generalization to challenging acquisitions To further assess the zero-shot generalization capability of uBrain on challenging unseen scans, we evaluate our framework on three additional datasets ($n = 10$ each): clinical thick-slice T1W, 7T Ultra-High-Field (UHF) T2W, and 5T UHF FLAIR scans. uBrain demonstrates robust performance across these diverse domains, yielding average Dice scores of 91.06%, 92.89%, and 90.73% for all volumetric segmentation tasks, alongside ASSD values of 0.31 mm, 0.23 mm, and 0.33 mm for surface reconstruction, respectively. As illustrated in Fig. 3, uBrain handles severe anisotropy in thick-slice data and domain shifts in UHF scans, highlighting its strong potential for real-world clinical deployment.

Analysis of Prompt-Driven Routing To evaluate the domain alignment of our pretrained text encoders, we perform a t-SNE visualization on the inter-

Table 3. Ablation study on the components of uBrain, with quantitative results reported as the average performance across all modalities for each task. Metrics are Dice score [%] for volumetric segmentation and ASSD [mm] for surface reconstruction.

		Baseline	Model II	Model III	Model IV	uBrain
Components	MAE	×	✓	×	✓	✓
	TDD	×	×	✓	✓	✓
	GMM	×	×	×	×	✓
Internal dataset	Tissue	93.78±3.17	93.85±3.49	94.05±3.15	94.43±3.08	**95.53±2.64**
	Hemisphere	93.51±3.21	94.05±3.19	93.79±3.17	94.26±3.14	**95.34±2.68**
	ROIs	89.07±4.12	89.49±3.98	89.40±4.06	90.21±4.02	**91.75±3.71**
	Left cortical surface	15.52±1.01	12.75±0.94	0.26±0.24	0.22±0.14	**0.19±0.10**
	Right cortical surface	18.24±0.98	11.52±0.67	0.26±0.21	0.23±0.13	**0.20±0.10**
External dataset	Tissue	91.24±3.65	91.84±3.58	92.12±3.52	93.45±3.25	**94.74±2.86**
	Hemisphere	90.95±3.72	91.68±3.61	91.55±3.58	93.20±3.31	**94.58±2.87**
	ROIs	85.12±4.85	86.45±4.65	86.20±4.72	88.54±4.25	**90.32±3.92**
	Left cortical surface	16.24±1.45	13.51±1.25	0.35±0.29	0.29±0.19	**0.22±0.12**
	Right cortical surface	18.51±1.52	14.22±1.16	0.36±0.29	0.29±0.19	**0.24±0.11**

nal test set. In Fig. 4(a), the image embeddings form distinct clusters aligned with their respective textual anchors, indicating that the encoders can distinguish among different modalities and task requirements. These results confirm that our pretrained encoders provide robust semantic priors, thereby facilitating modality- and task-specific representation learning in the subsequent MoME and AP modules.

We also visualize the dynamic routing of the MoME and AP modules. As shown in Fig. 4(b), MoME effectively models diverse modality distributions by engaging the most relevant experts for each input. Simultaneously, the AP modules (Fig. 4(c)) adaptively share or separate pathways among tasks based on semantic similarity. For instance, hemisphere tissue segmentation and ROI parcellation share early pathways for low-level structural features but diverge in deeper layers for distinct semantic processing. These results highlight uBrain's flexibility in handling multimodal and multitask learning within a unified framework.

3.3 Ablation Study

We conduct a comprehensive ablation study (Table 3) to disentangle the individual contributions of MAE, TDD, and GMM to the SwinUNETR baseline. While the static baseline struggles with heterogeneous inputs and diverse tasks, especially on surface reconstruction where the lack of TDD causes severe gradient interference and leads to non-convergence of the SDF branch (resulting in ASSD > 10 mm), adding MAE and TDD yields consistent improvements by effectively decoupling modality-specific features and dynamically routing task-specific representations. Crucially, the structure-preserving GMM augmentation forces the network to learn domain-invariant anatomical features, thereby further improving generalization to external datasets. It increases the average Dice

score from 91.73% to 93.21% and reduces the ASSD from 0.29 mm to 0.23 mm. In summary, our full uBrain framework achieves the best overall performance.

4 Conclusion

In this paper, we propose uBrain, a prompt-driven multimodal framework for brain volumetric segmentation and surface reconstruction. Guided by semantic text prompts, the MAE and TDD establish a dynamic routing mechanism to accommodate heterogeneous imaging protocols and diverse task requirements. Furthermore, the proposed structure-preserving GMM augmentation improves robustness to substantial domain shifts. Comprehensive evaluations confirm that uBrain achieves SOTA performance across routine acquisitions and challenging unseen domains, demonstrating its strong potential for real-world clinical deployment.

Acknowledgments. This work was supported in part by National Natural Science Foundation of China (82441023, 62131015, U23A20295, 82394432), the China Ministry of Science and Technology (STI2030-Major Projects-2022ZD0209000, STI2030-Major Projects-2022ZD0213100), The Key R&D Program of Guangdong Province, China (2023B0303040001), and HPC Platform of ShanghaiTech University.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Eickhoff, S.B., Yeo, B.T.T., Genon, S.: Imaging-based parcellations of the human brain. Nature Reviews Neuroscience **19**(11), 672–686 (2018)
2. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical surface-based analysis: I. Segmentation and surface reconstruction. NeuroImage **9**(2), 179–194 (1999)
3. Fischl, B., Sereno, M.I., Dale, A.M.: Cortical surface-based analysis: II: Inflation, flattening, and a surface-based coordinate system. NeuroImage **9**(2), 195–207 (1999)
4. Chen, X., et al.: Brain tissue segmentation across the human lifespan via supervised contrastive learning. arXiv preprint arXiv:2301.01369 (2023)
5. Teng, L., et al.: Knowledge-guided prompt learning for lifespan brain MR image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, pp. 238–248. Springer, Cham (2024)
6. Lian, Z., et al.: UniSurf: Universal lifespan cortical surface reconstruction. Medical Image Analysis **112**, 104090 (2026)
7. Smith, S.M., et al.: Advances in functional and structural MR image analysis and implementation as FSL. NeuroImage **23**(Suppl. 1), S208–S219 (2004)
8. Fischl, B.: FreeSurfer. NeuroImage **62**(2), 774–781 (2012)
9. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2015, pp. 234–241. Springer, Cham (2015)

10. Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature Methods **18**(2), 203–211 (2021)
11. Roy, S., et al.: MedNeXt: Transformer-driven scaling of ConvNets for medical image segmentation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023, pp. 405–415. Springer, Cham (2023)
12. Hatamizadeh, A., et al.: UNETR: Transformers for 3D medical image segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 1748–1758 (2022)
13. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images. In: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries, pp. 272–284. Springer, Cham (2022)
14. Santa Cruz, R., 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 Winter Conference on Applications of Computer Vision (WACV), pp. 806–815 (2021)
15. Lebrat, L., et al.: CorticalFlow: A diffeomorphic mesh transformer network for cortical surface reconstruction. In: Advances in Neural Information Processing Systems (NeurIPS), vol. 34, pp. 29491–29505 (2021)
16. 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 (2023)
17. 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 (CVPR), pp. 20741–20751 (2022)
18. Gopinath, K., Desrosiers, C., Lombaert, H.: SegRecon: Learning joint brain surface reconstruction and segmentation from images. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2021, pp. 650–659. Springer, Cham (2021)
19. Zhang, X., et al.: A foundation model for brain lesion segmentation with mixture of modality experts. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, pp. 379–389. Springer, Cham (2024)
20. Xiong, C., et al.: MoME: Mixture of multimodal experts for cancer survival prediction. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, pp. 318–328. Springer, Cham (2024)
21. Billot, B., et al.: SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. Medical Image Analysis **86**, 102789 (2023)
22. Gopinath, K., et al.: Recon-all-clinical: Cortical surface reconstruction and analysis of heterogeneous clinical brain MRI. Medical Image Analysis **103**, 103608 (2025)
23. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV), pp. 2999–3007 (2017)
24. Zhang, S., et al.: A multimodal biomedical foundation model trained from fifteen million image-text pairs. NEJM AI **2**(1), AIoa2400640 (2025)
25. Hu, E.J., et al.: LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations (ICLR) (2022)
26. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. ACM SIGGRAPH Computer Graphics **21**(4), 163–169 (1987)

27. Liu, J., et al.: BrainParc: unified lifespan brain parcellation from structural magnetic resonance images. Nature Computational Science, 1–15 (2026)
28. Butt, M.A., Maragos, P.: Optimum design of chamfer distance transforms. IEEE Transactions on Image Processing **7**(10), 1477–1484 (1998)
29. Liu, F., et al.: Harmonizing multi-modality biases in infant development analysis with an integrated MRI data processing pipeline. In: Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM), abstract 0238 (2024)
30. Chen, L., et al.: A 4D infant brain volumetric atlas based on the UNC/UMN baby connectome project (BCP) cohort. NeuroImage **253**, 119097 (2022)