

A Curvature-Guided Coarse-to-Fine Framework for Enhanced Whole Brain Segmentation

Fenqiang Zhao^{1,2(✉)}, Yuxing Tang¹, Le Lu¹, and Ling Zhang¹

¹ DAMO Academy, Alibaba Group

² Hupan Lab, 310023, Hangzhou, China

zhaofenqiang.zfq@alibaba-inc.com

Abstract. Whole brain segmentation, which divides the entire brain volume into anatomically labeled regions of interest (ROIs), is a crucial step in brain image analysis. Traditional methods often rely on intricate pipelines that, while accurate, are time-consuming and require expertise due to their complexity. Alternatively, end-to-end deep learning methods offer rapid whole brain segmentation but often sacrifice accuracy due to neglect of geometric features. In this paper, we propose a novel framework that integrates the key curvature feature, previously utilized by complex surface-based pipelines but overlooked by volume-based methods, into deep neural networks, thereby achieving both high accuracy and efficiency. Specifically, we first train a coarse anatomical segmentation model focusing on high-contrast tissue types, i.e., white matter (WM), gray matter (GM), and subcortical regions. Next, we reconstruct the cortical surfaces using the WM/GM interface and compute curvature features for each vertex on the surfaces. These curvature features are then mapped back to the image space, where they are combined with intensity features to train a finer cortical parcellation model. We also simplify the process of cortical surface reconstruction and curvature computation, thereby enhancing the overall efficiency of the framework. Additionally, our framework is flexible and can incorporate any neural network as its backbone. It can serve as a plug-and-play component to enhance the whole brain segmentation results of any segmentation network. Experimental results on the public Mindboggle-101 dataset demonstrate improved segmentation performance with comparable speed compared to various deep learning methods.

1 Introduction

The segmentation of neuroanatomy in brain magnetic resonance imaging (MRI) is a fundamental and critical step in brain MRI analysis [4]. It assigns each voxel a semantic label corresponding to a specific neuroanatomical structure so that quantitative measurements like volume, thickness, and area of various brain regions-of-interest (ROIs) can be obtained, which are key to numerous clinical and research applications, such as brain morphological analysis, surgical planning, and treatment assessment [7].

Existing methods for brain segmentation can generally be categorized into two groups. The first category is the atlas-based approach, which relies on deformable registration techniques to propagate manually segmented atlas labels to individual images, such as FreeSurfer [7] and BrainSuite [19]. It can be improved by using multiple atlases and label fusion strategies [21], or by incorporating additional machine learning classifiers to refine the segmentation results [23]. While these atlas-based methods are effective for segmenting high-contrast tissue types, such as white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF), as well as subcortical structures (SS) like the ventricle, cerebellum, and putamen, they fall short when it comes to the finer parcellation of gray matter, also known as the cerebral cortex, into subdivided regions. This is because the cerebral cortex is a highly convoluted and thin sheet with similar tissue intensities, where traditional intensity-based registration algorithms often fail to align them well [3]. Therefore, to obtain accurate whole brain segmentation (typically >100 labels) with finer cortical ROIs, pipelines like FreeSurfer [6] have developed cortical surface-based registration [23] and parcellation techniques [4] that explicitly use geometric features like cortical curvature to guide the alignment and parcellation of cortical areas, as shown in Fig. 1(a). However, the complexity and time-consuming nature of these pipelines, along with the need for specialized expertise, impede their routine use in clinical applications.

As a result, the second category of whole brain segmentation methods has emerged, leveraging deep learning techniques to develop end-to-end segmentation networks. These methods directly take the whole brain volume as input in a patch-based style [5], slice-based style [1], or sub-volume style [12], and output voxel-wise label classification results, as illustrated in Fig. 1(b). The implementation of these networks on GPUs successfully addresses the speed issue of traditional surface-based pipelines, which now only take minutes to process, thereby satisfying the requirements for real-time clinical applications. Most recently, a nested transformer model based on U-Net was developed [24], enabling local communication among adjacent patches by aggregating them hierarchically, resulting in state-of-the-art performance. Nevertheless, these end-to-end deep learning networks only rely on input images' intensity features, which tend to overlook the intricate geometric nuances of brain anatomy. This problem is particularly evident in cerebral cortex parcellation, which heavily depends on the cortical shape. Note that traditional surface-based pipelines specifically leverage cortical geometric features like cortical curvature to define the cortex parcellation protocols [4]. Therefore, existing deep learning methods, while fast, often sacrifice accuracy due to the neglect of geometric features.

To address these limitations of existing methods, we propose a novel framework that integrates the underexploited cortical geometric features into state-of-the-art deep learning networks. This integration allows the framework to combine the accuracy of traditional complex pipelines with the efficiency of advanced deep neural networks. Specifically, we adopted the cortical curvature feature, a key intrinsic property reflecting the convoluted topology of the brain, which has been widely used to define boundaries of cortical ROIs in various protocols [4,15,7]

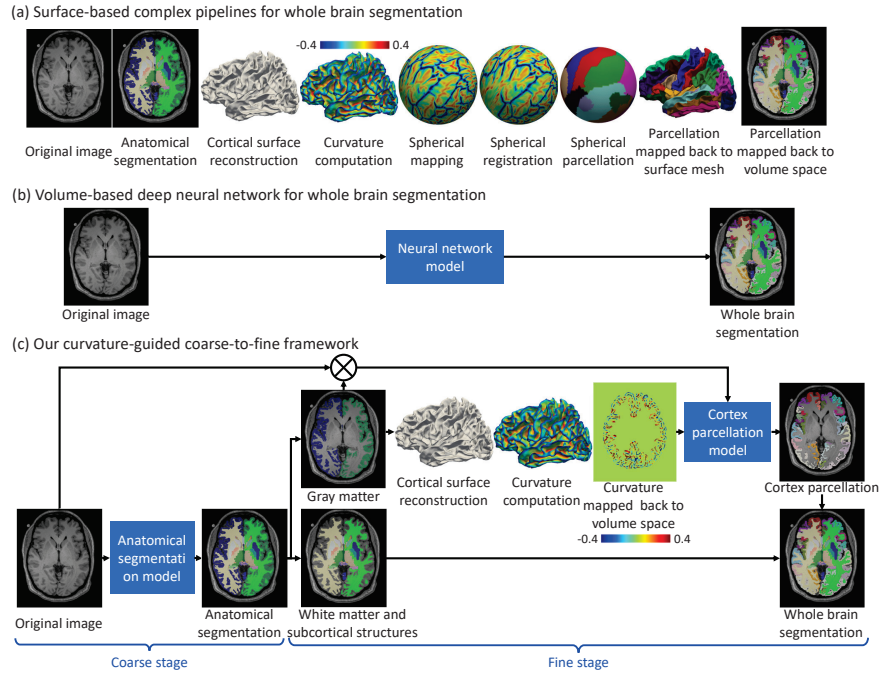


Fig. 1. Illustration of different methods for whole brain segmentation. Note that our framework keeps the key step, i.e., curvature computation of surface-based pipelines, and integrates it with the neural network. This approach allows our framework to benefit from both the accuracy of surface-based methods and the efficiency of deep learning methods. The \otimes symbol in (c) means element-wise multiplication operation.

and drive surface registration [23,25] and parcellation [17,22,27]. Our framework begins with training a coarse anatomical segmentation model targeting the WM, GM, and SS. Then, we reconstruct the cortical surfaces, compute the curvature feature at each vertex, and map the curvature features back to the volumetric image space. Subsequently, these curvature features are combined with intensity features to train a fine cortical parcellation model. Additionally, we optimize the process of cortical surface reconstruction and curvature computation, while eliminating non-essential steps, thereby enhancing the overall efficiency of the framework. Notably, our framework is designed to be flexible, and capable of incorporating any neural network as the backbone. As such, it can serve as a plug-and-play component, enhancing the whole brain segmentation results of any segmentation network.

2 Method

As shown in Fig. 1, motivated by the inherent limitations of previous methods, our framework is designed in a two-step coarse-to-fine manner, aiming to achieve

both high accuracy and efficiency. The first step is coarse anatomical segmentation for GM, WM, and SS like the ventricle, cerebellum, thalamus, caudate, putamen, etc. The second step is fine cortical parcellation to further divide the GM into various ROIs. This approach is justified by the inherent distinctions between anatomical segmentation and cortical parcellation tasks in whole brain segmentation. On one hand, anatomical segmentation aims to differentiate between different anatomical structures (WM, GM, and SS) primarily based on MRI intensities. This is feasible with volume-based methods since these structures often exhibit relatively high intensity contrast in typical structure MRI like T1-weighted (T1w) images. On the other hand, cortical parcellation seeks to partition the GM into different ROIs based on spatial location and shape information. Such information is hard to obtain using only intensity images, necessitating the use of cortical surface-based approaches. Hence, dividing the two tasks into separate steps with specialized models would clarify the training objectives for each model, rendering them more precise and attainable.

2.1 Coarse Anatomical Segmentation

As depicted in Fig. 1(c), we follow the well-established FreeSurfer pipeline [6] to initially conduct coarse anatomical segmentation, but with more efficient deep learning methods instead of the atlas-based method in FreeSurfer. In our implementation, we merge all cortical ROI labels into a single left GM and right GM label, combining them with WM and SS labels to train the model from scratch using the full brain volume as input. Since WM, GM, and SS exhibit relatively high intensity contrast and are distinguishable in T1w images, this model solely emphasizes learning intensity features. In contrast, previous end-to-end networks need to focus on both intensity features for anatomical structures and geometric features for cortical ROIs.

2.2 Curvature Computation

After completing the coarse anatomical segmentation, we obtain the masks of GM, WM, and each SS. To leverage cortical geometric features for more accurate cortical parcellation, we integrate several steps from surface-based analysis to extract essential cortical curvature features [8]. Traditional surface-based pipelines are time-consuming and involve complex steps (see Fig. 1(a)), including surface reconstruction [16], topology correction [9], spherical mapping [8], surface registration [23], parcellation [4], etc. We optimize this process for whole brain segmentation by retaining only the crucial steps necessary for curvature computation while discarding non-essential steps such as topology correction, spherical mapping, and surface registration [23]. This simplification is justified by the observation that most of these steps can now be replaced with their geometric deep learning counterparts, such as [20] for topology correction, [26] for spherical mapping, and [25] for surface registration. Therefore, we assume that these steps can be implicitly learned and addressed to some extent by a single model focusing on anatomical details and geometric features.

Specifically, we begin by removing the cerebellum and brain stem and filling other SS regions. Subsequently, we reconstruct the GM/WM and GM/background surfaces using the fast marching cube method in the scikit library [18], taking less than 1 second. The surfaces are then iteratively smoothed, and a simplified curvature computation method is applied to obtain the mean curvature feature for each vertex. This simplified method converts the complicated maximum and minimum curvature computation of 3D surface to multiple 2D curves’ curvature calculation, which computes the curvature of each curve connecting the 2-ring neighborhood vertex to the center vertex and then averages them. Such methods have been validated to be effective in our experiments, requiring only ~ 1.5 seconds. After curvature computation, we map the features from 3D surface space to 3D volumetric space for fine cortical parcellation.

2.3 Fine Cortical Parcellation

In the fine cortical parcellation stage, we directly concatenate the computed curvature features and original input intensity images in the volumetric space to train the fine cortical parcellation model. To eliminate the influence of WM and SS regions that have already been segmented in the previous stage and to focus only on the cerebral cortex, we crop the intensity images using the predicted GM mask, as illustrated in Fig. 1. Similarly, in our implementation, we keep the cortical ROI labels and merge the WM and SS labels with the background label. After training the model and performing inference to obtain cortical ROI labels on the whole brain volume, we also crop it using the predicted GM mask. Subsequently, we combine it with the previously predicted WM and SS masks to obtain the final whole brain segmentation results.

3 Experiments and Results

3.1 Experimental Settings

We evaluated our method using the publicly available Mindboggle-101 dataset [15]. This dataset contains 101 3D brain MRI T1w volumes, with each image containing 106 manually labeled ROIs. These labels include 31 cortex labels in each hemisphere, 1 WM label in each hemisphere, 41 SS labels, and 1 background label. The manual labeling of cortex ROIs was based on the DKT protocol [15], an improved version of the Desikan-Killiany protocol [4]. The segmentations of SS regions were obtained and corrected using FreeSurfer’s atlas-based method. Detailed information regarding label names and labeling process can be found in [15]. The original size of the data is $182 \times 218 \times 182$ with 1 mm isotropic spacing. The subjects’ ages range from 19 to 60 years old. We randomly selected 30 scans from the dataset as a hold-out test set, while the remaining 71 scans were used for 5-fold cross-validation.

We adopted several popular medical image segmentation methods as our baselines, implemented our method based on their code, and performed experiments within their respective framework, including naive 3D U-Net [2], nnU-Net

Table 1. Comparison of Dice (%) performance for whole brain segmentation on the hold-out test set, with standard deviations calculated across 5 models of the 5-fold cross-validation. SS: Subcortical Structures, WM: White Matter, GM: Gray Matter, and the “All” column is averaged over all regions. #1 means an end-to-end network for all ROIs. #2 represents the proposed two-step coarse-to-fine framework, where the fine stage is without (w/o) or with (w/) curvature (curv) features. “-” indicates that the results are the same as their counterparts w/o curv.

	SS	WM	GM	All
3D U-Net #1	73.64±0.54	85.33±0.93	73.30±1.21	73.51±1.08
3D U-Net #2 w/o curv	75.47±0.77	90.39±0.54	77.56±0.49	77.02±0.64
3D U-Net #2 w/ curv	-	-	79.31±0.69	78.04±0.53
UNETR #1	61.15±0.77	71.77±1.34	58.66±1.56	61.53±1.19
UNETR #2 w/o curv	64.78±0.75	73.24±0.78	63.22±0.63	64.01±0.86
UNETR #2 w/ curv	-	-	66.35±1.09	65.87±1.37
SwinUNETR #1	73.14±0.38	86.63±0.49	73.88±0.52	74.95±0.73
SwinUNETR #2 w/o curv	74.08±0.44	90.09±0.76	77.86±0.55	76.64±0.61
SwinUNETR #2 w/ curv	-	-	79.22±0.36	77.45±0.42
nnU-Net #1	77.33±0.43	93.05±0.23	78.59±0.91	78.36±0.72
nnU-Net #2 w/o curv	77.42±0.62	94.46±0.18	81.04±0.40	79.95±0.24
nnU-Net #2 w/ curv	-	-	83.13±0.49	81.36±0.45

[13], MONAI for UNETR [11] and SwinUNETR [10]. To effectively demonstrate the advantages of the proposed curvature features for two-stage whole brain segmentation, we made minimal modifications to the pre-defined parameters of each method. Specifically, we only adjusted the input and output channels, setting them to 1 and 46 in the coarse stage, and 2 and 63 in the fine stage, respectively. To enhance generalization ability with a relatively small dataset, we opted not to perform any preprocessing on the original dataset. Instead, we conducted extensive data augmentation such as random patch cropping, rotation, scaling, adding Gaussian noise, adjusting brightness and contrast, etc. We trained the models using default parameters of each baseline method, such as a linearly decayed learning rate for nnU-Net and a cosine annealing schedule for UNETR and SwinUNETR, SGD optimizer for nnU-Net and Adam optimizer for UNETR and SwinUNETR, Dice and weighted cross entropy loss for nnU-Net, and only Dice loss for UNETR and SwinUNETR, etc. All the experiments were conducted on NVIDIA Tesla V100 32GB GPUs.

3.2 Results

Quantitative Analysis Table 1 shows the comprehensive quantitative results of whole brain segmentation on the hold-out test set using different baselines and our method. It can be observed the implementation of our coarse-to-fine strategy, even without the inclusion of curvature features (#2 w/o curv) at the fine cortical parcellation stage, already yields enhanced results. Specifically, there is an average 3.8% improvement in Dice scores for GM and a 2.3% improvement for all ROIs compared to the baseline models. This indicates that training anatomical




Fig. 2. Qualitative results of whole brain segmentation using different methods. #1 means the original baseline model trained for all ROIs. #2 w/ curv represents the models enhanced by our proposed curvature-guided coarse-to-fine framework. Boxed areas are enlarged in the lower row. Differences are emphasized with the yellow arrows.

segmentation model and cortical parcellation model separately indeed helps the network to more effectively focus on relevant features, i.e., intensity features for anatomical structures and geometric features for cortical regions, thereby facilitating the achievement of both objectives. Then incorporating explicit curvature features (#2 w/ curv) boosts the results further, with an average 2.1% Dice improvement for GM compared to the model without curvatures (#2 w/o curv). These observations suggest that incorporating curvature information directly into the model is beneficial for delineating cortical ROI boundaries more accurately. Notably, these gains are observed consistently across all baseline models, which strongly suggests that our method can serve as a robust, plug-and-play component to enhance the segmentation results of various whole brain segmentation networks.

Qualitative Analysis Fig. 2 shows some qualitative whole brain segmentation results from the test set. We can see that baseline models generally tend to generate over-smooth segmentations, where CSF among cortical folds are often incorrectly labeled as cortical regions, resulting in over-smooth and unrealistic results. This may be attributed to their limited capability to capture intricate geometric features of the GM, which are crucial for accurately detecting the boundaries of GM ROIs. By integrating our proposed curvature features, there is a noticeable improvement consistently for all baseline models, particularly in challenging cases where two cortical folds are very close. For example, in the last row of Fig. 2, baseline nnU-Net mistakenly labels the left precentral gyrus (in gray color) as its adjacent ROI left pars opercularis (in blue color) due to

Table 2. Comparison with state-of-the-art whole brain segmentation methods.

	SS Dice (%)	WM Dice (%)	GM Dice (%)	All Dice (%)	Time
FreeSurfer [6]	76.28	91.22	86.45	82.84	~5h
UNesT [24]	74.59	88.32	73.49	74.20	~10s
Our method (nnU-Net #2 w/ curv)	77.42	94.46	83.13	81.36	~10s

the difficulty in distinguishing the two close cortical folds. With our proposed curvature feature, nnU-Net #2 w/ curv successfully detects the subtle change of boundary shifts of the two ROIs and correctly labels them.

Comparison with State-of-the-art Methods Table 2 shows the comparison results with current state-of-the-art whole brain segmentation methods on our Mindboggle-101 test set. For UNesT [24], we directly run their public code and pre-trained model on our test set. Since the segmentation protocol with 133 regions they used [14] is an older version of Mindboggle-101’s protocol [15], we merged some labels accordingly and only calculated the Dice for regions presented in both protocols. For FreeSurfer, we ran the “recon-all” command in FreeSurfer software to obtain the whole brain segmentation for each scan. From the table, we can observe that FreeSurfer obtains the best GM parcellation results and overall Dice while taking the longest time. Our method obtains the second-best results on GM since we only used several simplified surface-based processes to obtain geometric features. UNesT’s results are lower possibly due to dataset variance and further fine-tuning on our dataset may improve it. In terms of WM and SS, since they are relatively easy to recognize on T1w images solely based on intensities, our deep learning model in the first step can perform better intensity feature detection and recognition for the segmentation task than the traditional atlas-based method in FreeSurfer. It is also worth noting that our method is as fast as UNesT, taking only 10s to obtain the whole brain segmentation, which is 1000+ times faster than FreeSurfer.

4 Conclusion

In this paper, we propose a novel curvature-guided coarse-to-fine framework to enhance the performance of whole brain segmentation networks. The framework divides anatomical segmentation and cortical parcellation tasks into two steps with specialized models, facilitating the feature learning of each model and improving the segmentation of WM and SS. It also leverages explicit cortical curvature feature to facilitate the detection of ROIs’ boundaries and thus enhance the accuracy of cortical parcellation. Extensive validation with popular segmentation networks demonstrates that our method consistently improves the segmentation performance across various baseline models. More importantly, by employing an optimized surface reconstruction and curvature computation algorithm, our

method can segment a whole brain volume in 10s, which is 1000+ times than FreeSurfer, offering significant potential for real-world clinical applications.

References

1. Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P.A.: Voxresnet: Deep voxelwise residual networks for brain segmentation from 3d mr images. *NeuroImage* **170**, 446–455 (2018)
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2016: 19th International Conference, Athens, Greece, October 17–21, 2016, Proceedings, Part II 19*. pp. 424–432. Springer (2016)
3. Coalson, T.S., Van Essen, D.C., Glasser, M.F.: The impact of traditional neuroimaging methods on the spatial localization of cortical areas. *Proceedings of the National Academy of Sciences* **115**(27), E6356–E6365 (2018)
4. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., et al.: An automated labeling system for subdividing the human cerebral cortex on mri scans into gyral based regions of interest. *Neuroimage* **31**(3), 968–980 (2006)
5. Fang, L., Zhang, L., Nie, D., Cao, X., Rekik, I., Lee, S.W., He, H., Shen, D.: Automatic brain labeling via multi-atlas guided fully convolutional networks. *Medical image analysis* **51**, 157–168 (2019)
6. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
7. Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al.: Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* **33**(3), 341–355 (2002)
8. 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)
9. Han, X., Xu, C., Braga-Neto, U., Prince, J.L.: Topology correction in brain cortex segmentation using a multiscale, graph-based algorithm. *IEEE Transactions on Medical Imaging* **21**(2), 109–121 (2002)
10. 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: *International MICCAI Brainlesion Workshop*. pp. 272–284. Springer (2021)
11. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. pp. 574–584 (2022)
12. Huo, Y., Xu, Z., Xiong, Y., Aboud, K., Parvathaneni, P., Bao, S., Bermudez, C., Resnick, S.M., Cutting, L.E., Landman, B.A.: 3d whole brain segmentation using spatially localized atlas network tiles. *NeuroImage* **194**, 105–119 (2019)
13. Isensee, F., Jaeger, P.F., Kohl, S.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)
14. Klein, A., Canton, T., Ghosh, S., Landman, B., Lee, J., Worth, A.: Open labels: online feedback for a public resource of manually labeled brain images. *16th annu. Meet. Organ. Hum. Brain Mapp* (2010)

15. Klein, A., Tourville, J.: 101 labeled brain images and a consistent human cortical labeling protocol. *Frontiers in neuroscience* **6**, 171 (2012)
16. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: *Seminal graphics: pioneering efforts that shaped the field*, pp. 347–353 (1998)
17. Parvathaneni, P., Bao, S., Nath, V., Woodward, N.D., Claassen, D.O., Cascio, C.J., Zald, D.H., Huo, Y., Landman, B.A., Lyu, I.: Cortical surface parcellation using spherical convolutional neural networks. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III* 22. pp. 501–509. Springer (2019)
18. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. *Journal of machine learning research* **12**(Oct), 2825–2830 (2011)
19. Shattuck, D.W., Leahy, R.M.: Brainsuite: an automated cortical surface identification tool. *Medical image analysis* **6**(2), 129–142 (2002)
20. Sun, L., Zhang, D., Lian, C., Wang, L., Wu, Z., Shao, W., Lin, W., Shen, D., Li, G., Consortium, U.B.C.P., et al.: Topological correction of infant white matter surfaces using anatomically constrained convolutional neural network. *NeuroImage* **198**, 114–124 (2019)
21. Wang, H., Suh, J.W., Das, S.R., Pluta, J.B., Craige, C., Yushkevich, P.A.: Multi-atlas segmentation with joint label fusion. *IEEE transactions on pattern analysis and machine intelligence* **35**(3), 611–623 (2012)
22. Wu, Z., Li, G., Wang, L., Shi, F., Lin, W., Gilmore, J.H., Shen, D.: Registration-free infant cortical surface parcellation using deep convolutional neural networks. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 672–680. Springer (2018)
23. Yeo, B.T., Sabuncu, M.R., Vercauteren, T., Ayache, N., Fischl, B., Golland, P.: Spherical demons: fast diffeomorphic landmark-free surface registration. *IEEE transactions on medical imaging* **29**(3), 650–668 (2009)
24. Yu, X., Yang, Q., Zhou, Y., Cai, L.Y., Gao, R., Lee, H.H., Li, T., Bao, S., Xu, Z., Lasko, T.A., et al.: Unest: Local spatial representation learning with hierarchical transformer for efficient medical segmentation. *Medical Image Analysis* **90**, 102939 (2023)
25. Zhao, F., Wu, Z., Wang, F., Lin, W., Xia, S., Shen, D., Wang, L., Li, G.: S3reg: Superfast spherical surface registration based on deep learning. *IEEE Transactions on Medical Imaging* (2021)
26. Zhao, F., Wu, Z., Wang, L., Lin, W., Li, G.: Fast spherical mapping of cortical surface meshes using deep unsupervised learning. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 163–173. Springer (2022)
27. Zhao, F., Xia, S., Wu, Z., Duan, D., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G.: Spherical u-net on cortical surfaces: methods and applications. In: *International Conference on Information Processing in Medical Imaging*. pp. 855–866. Springer (2019)