

UNSURF: Uncertainty Quantification for Cortical Surface Reconstruction of Clinical Brain MRIs

Karthik Gopinath^{1,2,*}, Raghav Mehta^{3,*}, Ben Glocker^{3,**}, and Juan Eugenio Iglesias^{1,2,**}

¹ A. A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital

² Harvard Medical School, Boston, USA.

³ Imperial College London, UK.

raghav.mehta@imperial.ac.uk

Abstract. We propose UNSURF, a novel uncertainty measure for cortical surface reconstruction of clinical brain MRI scans of any orientation, resolution, and contrast. It relies on the discrepancy between predicted voxel-wise signed distance functions (SDFs) and the actual SDFs of the fitted surfaces. Our experiments on real clinical scans show that traditional uncertainty measures, such as voxel-wise Monte Carlo variance, are not suitable for modeling the uncertainty of surface placement. Our results demonstrate that UNSURF estimates correlate well with the ground truth errors and: (i) enable effective automated quality control of surface reconstructions at the subject-, parcel-, mesh node-level; and (ii) improve performance on a downstream Alzheimer’s disease classification task.

Keywords: Uncertainty Estimation · Cortical Analysis · Clinical MRI.

1 Introduction

Cortical surface reconstruction from MRI plays a crucial role in neuroimaging, enabling analyses like parcellation and cortical thickness – a powerful biomarker for studying healthy aging and neurodegenerative diseases such as Alzheimer’s [32]. Traditional methods rely on high-resolution, isotropic MRI scans, typically acquired using T1-weighted (T1) sequences with near 1 mm isotropic resolution. However, clinical MRI scans often exhibit significant variations in orientation, resolution (slice spacing), and contrast, making surface reconstruction difficult.

Recent advancements in machine learning algorithms have enabled analysis of clinical brain MRI acquired at hospitals [4,21]. Uncertainty estimation is a critical aspect of cortical surface reconstruction in these settings, since large slice spacing and potential motion artifacts often make surface placement ambiguous. Quantification of uncertainty is required to identify unreliable measurements and control their impact in downstream analyses like group studies of cortical thickness. Despite its importance, uncertainty quantification remains underexplored in cortical surface reconstruction because standard techniques are not

* Equal contribution

** Equal senior author




Fig. 1. Overview of the proposed uncertainty estimation method (UNSURF).

applicable due to the topological and geometric constraints involved (e.g., lack of self-intersections). Without confidence estimates, it is difficult to determine whether surface errors arise from model limitations or poor-quality input data, limiting the clinical reliability and applicability of automated pipelines.

In this work, we introduce UNSURF, a novel measure of uncertainty in cortical surface reconstructions. UNSURF enables a multi-level Quality Control (QC) strategy at the subject, parcel, and mesh node levels. Experiments on two clinical datasets show that it effectively identifies unreliable regions. We further show that using this strategy for filtering unreliable measurements improves effect sizes in downstream neuroimaging studies of cortical thickness.

Related Work: Classical surface reconstruction methods like FreeSurfer [9] or BrainSuite [33] employ volumetric segmentation followed by topology correction and geometry-constrained surface extraction. These pipelines enable accurate white and pial surface reconstructions from T1 scans by leveraging neuroanatomical priors. However, they require high-resolution isotropic scans and have long processing times, often exceeding a few hours per subject. Several deep learning methods have been proposed to reconstruct cortical surfaces from MRI scans; they either utilize a combination of convolution and graph-based neural networks [17,5,24,18], deformation methods [38], or leverage implicit representations [6,13,24,39] to reconstruct cortical surfaces. Nevertheless, these are supervised methods trained for a specific contrast and resolution (isotropic T1 scans) and do not generalize to clinical MRI. To address these challenges, a domain randomization approach has recently been introduced [14,16], which can extract surfaces from scans of any resolution or contrast without domain adaptation or finetuning. This approach builds on prior work on segmentation, synthesis, and super-resolution [4,21] and relies on training neural networks on synthetic datasets spanning an unrealistically broad range of contrasts and resolutions, enabling models to generalize across heterogeneous clinical data at test time.

Several methods have been proposed for uncertainty quantification in medical imaging [2,20]. Bayesian deep learning [28], Monte Carlo dropout [11], and deep ensembles [23] have been applied to various medical image analysis tasks [26],

such as segmentation [27,3,25], registration [19,7], precision medicine [8], or image reconstruction [35,36]. Test-time augmentation [37] and distribution-aware models [30] further improve reliability by assessing model sensitivity to variations in input data. However, direct application of these methods to cortical surface reconstruction is not straightforward due to geometric constraints; as such uncertainty quantification for cortical surface reconstruction remains unexplored.

2 Methodology

UNSURF is illustrated in Fig. 1 and comprises two modules: one for surface placement and another for uncertainty estimation. The surface placement builds on *recon-all-clinical* [15] and consists of two sub-modules: first, a convolutional neural network (CNN) that predicts four 1 mm isotropic signed distance functions (SDFs) from an input MRI scan – two per hemisphere: pial and white matter (WM) surfaces. The second sub-module uses geometry processing to extract topologically constrained surfaces. The uncertainty estimation module quantifies uncertainty by generating distance maps from the predicted surfaces and using them to compute voxel-wise uncertainty measures. These are mapped back to the surfaces to produce the uncertainty estimates on the reconstructed surfaces.

2.1 Cortical surface placement via SDF prediction

For surface placement, our method uses a hybrid algorithm presented in [15], which combines learning-based estimation of SDFs and classical geometry processing and is agnostic to the resolution and contrast of the input scan. Specifically, given an input brain MRI ($\mathbf{X}^v$), we use a CNN to estimate the four SDFs $\hat{\mathbf{D}}_i^v = f(\mathbf{X}^v; \theta)$, which encode the signed distance of each voxel to the four cortical surfaces. Here, v represents a voxel in a volume, and $i \in [1, 2, 3, 4]$ represents four SDFs. To make the CNN agnostic to resolution and contrast, we employ a domain randomization strategy [4,21] that generates synthetic data on the fly during training. Crucially, the orientation, resolution, and contrast of these synthetic scans are sampled from uniform distributions at every minibatch, thus forcing the CNN to learn contrast- and resolution-agnostic features. The training data are automatically generated from existing datasets with 1 mm T1s: 3D segmentations (used to generate the synthetic images) are extracted with SAMSEG [29], whereas the pial and WM surfaces (used to generate the ground truth SDFs used as targets) are extracted with FreeSurfer [9].

We synthesize training data with the generative model from [4,21], which we summarize here for completeness. Starting from 3D segmentations, random affine and nonlinear spatial transformations are first applied to both the segmentations and SDFs for geometric augmentation. These transformations approximate true SDFs without explicitly deforming surfaces or recomputing distances, offering a computationally efficient solution in practice. Next, it samples a Gaussian mixture model conditioned on the segmentation to generate Gaussian images. The

Gaussian parameters (means, variances) are randomly sampled at every minibatch. Non-isotropic acquisitions are simulated by spatial smoothing in order to mimic a combination of orientation, slice spacing, and thickness – also sampled randomly at every minibatch. Finally, a (randomly sampled) smooth multiplicative bias field is applied to the images. We note that the voxel size of the input image and the target SDFs is 1 mm isotropic; the *intrinsic* resolution of the input depends on the simulated resolution (and is generally lower than 1 mm isotropic), whereas the SDF is always crisp. This strategy enables the CNN to predict high-resolution SDFs independently of the resolution of the input.

The neural network is trained to predict the SDFs of the WM and pial surfaces using an L2 loss: $L_{\text{SDF}} = \frac{1}{|v|} \sum_i (\hat{\mathbf{D}}_i^v - \mathbf{D}_i^v)^2$, where $\mathbf{D}_i^v$ is the ground-truth SDF. We truncate the SDFs at ± 5 mm to constrain learning to cortical regions (e.g., we do not want to waste network capacity by discriminating whether a voxel is 9 vs 10 mm away from the surface). At test time, the input is resampled to 1 mm resolution and the trained model predicts 1 mm isotropic SDFs for real MRI scans irrespective of their resolution and contrast, as explained above.

The predicted distance maps $\hat{\mathbf{D}}_i^v$ are used for surface extraction ($\mathbf{S}_i^n$) using a geometry processing module $g(\cdot, \phi)$ similar to that of [15] (which in turn relies heavily on FreeSurfer tools [9]). In short, the WM surface $\mathbf{S}_i^{\text{wm}}$ is initialized as the level-set of the predicted WM SDF and is iteratively smoothed and deformed to optimize mesh regularity. Topological correction is performed on the mesh to ensure that its Euler characteristic is 2 (i.e., topologically equivalent to a sphere, without holes or handles). The refined mesh is then inflated by expanding the cortical surface into a sphere while maintaining the anatomical features. The pial surface $\mathbf{S}_i^{\text{pial}}$ is then placed by expanding $\mathbf{S}_i^{\text{wm}}$ outward using the pial SDF while preventing self-intersections. The surfaces are spherically registered to a common coordinate system, enabling cortical parcellation using the Desikan-Killiany Atlas. This step allows subject-wise comparisons and statistical analysis of cortical morphology across populations.

2.2 Uncertainty Estimation for Cortical Surface Reconstruction

Baseline - Ensemble Dropout: To estimate uncertainty for cortical surface reconstruction, we utilize ensemble dropout [34] which has been shown to provide better uncertainty estimates for different tasks [26,1,10] compared to Monte Carlo dropout [11] or deep ensembles [23]. Using N independently trained networks, and taking Z dropout samples from each of them, ensemble dropout produces a total $M = N \times Z$ samples of SDFs for the same input. The final SDF estimate $\hat{\mathbf{D}}_i^v$ is computed as: $\hat{\mathbf{D}}_i^v = \frac{1}{M} \sum_{m=1}^M \hat{\mathbf{D}}_{i_m}^v$. The voxel-wise uncertainty is measured as the variance across predictions:

$$(\sigma_i^v)^2 = \frac{1}{M-1} \sum_{m=1}^M (\hat{\mathbf{D}}_{i_m}^v - \hat{\mathbf{D}}_i^v)^2.$$

While the variance is useful for capturing local deviations in $\hat{\mathbf{D}}_i^v$, it does not account for topological corrections introduced by $g(\cdot, \phi)$. An alternative approach

would be to estimate variance uncertainty on the surface mesh node coordinates $\mathbf{S}_i^n$. However, this would be computationally costly, as each geometry processing step $g(\cdot, \phi)$ takes approximately 2 hours; as such, passing M distance maps through it would require $M \times 2$ hours, which is infeasible in a practical scenario. Another approach could be to use surface deformation methods [5, 18, 24] that do not require explicit topological correction, allowing estimation of variance directly on the surface nodes without high computational overhead. However, to our knowledge, these methods are currently limited to high-resolution T1 scans, and Recon-all-clinical [15] is currently the only approach applicable to surface reconstruction of heterogeneous clinical scans.

UNSURF - Uncertainty estimation with difference map: To overcome variance limitations (high computation time and lack of topological correction), we estimate uncertainty using a new measure (UNSURF), which relies on the consistency (or lack thereof) between the SDFs predicted by the neural network ($\hat{\mathbf{D}}_i^v$) and actual SDFs ($\tilde{\mathbf{D}}_i^v$) computed from the extracted surfaces $\mathbf{S}_i^n$ after geometry processing ($g(\cdot, \phi)$). Voxel-level uncertainty as the squared discrepancy:

$$\mathcal{U}_i^v = (\hat{\mathbf{D}}_i^v - \tilde{\mathbf{D}}_i^v)^2.$$

$\mathcal{U}$ serves as a more accurate proxy for uncertainty in surface reconstruction, reflecting geometric corrections and areas with significant discrepancies between the initial voxel-wise prediction ($\hat{\mathbf{D}}_i^v$) and the actual SDF of the final surface ($\tilde{\mathbf{D}}_i^v$). Larger $\mathcal{U}$ indicate greater uncertainty, as they correspond to regions with multiple plausible surfaces.

2.3 Implementation details

Our voxel-wise regression CNN follows the same architecture as *recon-all-clinical* [15], utilizing a 3D U-Net [31] trained with synthetic pairs generated on the fly, as described in Section 2.1. The U-Net consists of five levels with two layers each, employing $3 \times 3 \times 3$ convolutions and exponential linear unit activations. Each level l contains $24 \times 2^{(l-1)}$ features. The final layer includes a dropout layer and uses a linear activation function to model SDFs.

To estimate uncertainty using ensemble dropout variance (σ^2), we train five randomly initialized models on the same dataset using identical hyperparameters and total epochs. We perform 20 forward passes per input for each of the five models at test time, averaging the resulting SDF predictions ($\hat{\mathbf{D}}_i^v$). The geometry processing pipeline remains the same as that of FreeSurfer. Once the surfaces are reconstructed, we estimate the implicit SDF ($\tilde{\mathbf{D}}_i^v$) to compute the $\mathcal{U}_i^v$.

We resampled the uncertainty ($\mathcal{U}^v$ and $(\sigma^2)^v$) values onto the mesh nodes using trilinear interpolation, to obtain node-level uncertainty ($\mathcal{U}^n$ and $(\sigma^2)^n$). Given a parcellation, we compute region-wise uncertainty by averaging it over the nodes in each parcel ($\mathcal{U}^p$ and $(\sigma^2)^p$). Similarly, to obtain a subject-level representation, we compute the global average of uncertainty over all surface nodes ($\mathcal{U}^s$ and $(\sigma^2)^s$). This approach allows us to summarize uncertainty at




Fig. 2. Comparison of uncertainty estimates vs. cortical thickness error for Variance - $(\sigma^2)^s$ (left) and U^s (right) on high-resolution (T1-1mm) and clinical scans. SCC and PCC stand for Spearman and Person correlation coefficient.

different levels of granularity, capturing localized effects at the node level and parcel level while also obtaining an overall subject-wise measure.

3 Experiments and Results

3.1 Datasets

Training Data: We trained the CNN using segmentation maps and SDFs derived from 1,000 1 mm T1 MRI scans. These include 500 scans from the Human Connectome Project (HCP) [12] and 500 scans from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) [22]. FreeSurfer [9] was used to obtain ground-truth cortical surfaces and morphometric measures, such as cortical parcellation and cortical thickness. We emphasize that only the 3D segmentations [29] and estimated SDFs from the white and pial surfaces are used during training; the original T1-weighted MRI scans are disregarded.

Test Datasets: We used two different datasets to evaluate our method. The first dataset (henceforth “synthetically downsampled dataset”) comprises subjects with high-quality scans, which we progressively downsample to analyze the performance of our method as a function of resolution. Specifically, we used 15 subjects chosen at random from the HCP dataset (each with 0.7 mm isotropic T1 and T2 scans) and 15 subjects from the ADNI3 dataset (each with ~ 1 mm isotropic T1 and FLAIR scans). We simulated clinical acquisition on three different orientations (axial, coronal, sagittal) for 5 different slice spacings (2 mm, 3 mm, 4 mm, 5 mm, and 6 mm). The slice thickness was made equal to the slice spacing in all cases. The second dataset comprises 200 randomly selected subjects (age: 74.0 ± 7.4 years; 95 males) from ADNI1. Since ADNI1 includes high-resolution (~ 1 mm isotropic) T1 scans and a 5 mm axial FLAIRs, it enables evaluation on images natively acquired at low-resolution (the axial FLAIRs) using ground truth derived from the high-resolution T1s.

3.2 Comparison Between Variance and $\mathcal{U}$

We first compare UNSURF $\mathcal{U}$ with **ensemble dropout variance** by analyzing their correlation with **thickness errors**. Fig. 2 shows the correlation coefficients (Spearman's - SCC, and Pearson's - PCC) between uncertainty estimates and cortical thickness errors at the subject level ($\mathcal{U}^s$ and $(\sigma^2)^s$). Being unaware of geometry, variance-based uncertainty exhibits weak or even negative correlation with cortical thickness errors across both high-resolution and clinical-quality scans, making it unreliable for identifying surface reconstruction errors. In contrast, $\mathcal{U}$ consistently shows a strong positive correlation with cortical thickness error, indicating that it better captures regions where cortical surface placement is uncertain. The key reason behind $\mathcal{U}$'s improved performance is that variance only reflects prediction dispersion, whereas $\mathcal{U}$ explicitly accounts for the geometric corrections applied during surface reconstruction. This makes $\mathcal{U}$ a more precise measure of uncertainty in cortical thickness estimation.

3.3 Results on synthetically downsampled dataset

To better understand how $\mathcal{U}$ behaves across different cortical regions, we show in Fig. 3(a) the spatial distribution of SCC across the cortex, confirming a strong correlation between $\mathcal{U}^p$ and cortical thickness error in many regions. Furthermore, we analyze four representative parcels that are highly relevant to Alzheimer's disease [32] (parstriangularis, inferiorparietal, supramarginal, and rostralmiddlefrontal). In Fig. 3(b), we observe a clear positive correlation between $\mathcal{U}$ and cortical thickness error for all four selected parcels, further supporting that $\mathcal{U}$ effectively captures reconstruction uncertainty. Interestingly, we see in Fig. 3(c) that SCC values are lower for thinner slices (e.g., 2 mm and 3 mm), i.e., it is harder to accurately predict differences between smaller errors.

We further quantify uncertainty by filtering out the most uncertain predictions and examining the change in thickness errors. In Fig. 3(d), we observe that removing high- $\mathcal{U}^p$ parcels reduces overall parcel-level cortical thickness errors, demonstrating that $\mathcal{U}$ can be used for automated QC. Further, in Fig. 3(e), we see that filtering uncertain ($\mathcal{U}^n$) nodes within already uncertain parcels ($\mathcal{U}^p$) leads to further reduction in cortical thickness errors, reinforcing the utility of multi-level QC.

3.4 Results on the ADNI Dataset (AD vs. CN Classification)

We evaluate the impact of $\mathcal{U}$ -based QC on a AD vs CN classification task, which is representative of group studies in neuroimaging. The goal is to assess whether removing uncertain cortical thickness measurements ($\mathcal{U}^s$) improves statistical power when measuring group differences. We first fit a general linear model for cortical thickness at each parcel, adjusting for age and gender. We then compute Cohen's d as effect size between the AD and CN groups. As shown in Fig. 4(left), the effect size increases as we remove high-uncertainty data, indicating that uncertainty-based filtering enhances group separation. The trend of the




Fig. 3. Evaluation of parcel-level and node-level uncertainty estimates on clinical scans. (a) SCC between uncertainty and cortical thickness error for all parcels. (b) $\mathcal{U}^p$ vs. cortical thickness error for four different parcels and varying slice thickness. (c) SCC across slice thicknesses. (d) Cortical thickness error as a function of percent of kept data based on parcel-level uncertainty ($\mathcal{U}^p$). The x-axis refers to the percent of kept dataset based on uncertainty filtering (i.e. At 100, all subjects are kept when computing thickness error). (e) Cortical thickness error as a function of percent of kept data based on node-level uncertainty ($\mathcal{U}^n$) within the top-10% uncertain subjects ($\mathcal{U}^p$).

parcel-wise effect size in Fig. 4(right) confirms that removing uncertain thickness estimates improves sensitivity across different brain regions.

4 Conclusion

This paper proposed a novel uncertainty estimation measure (UNSURF - $\mathcal{U}$) for cortical parcel reconstruction using implicit surface methods, and demonstrates its importance in clinical applications. Future work will seek to incorporate uncertainty estimation into faster explicit surface reconstruction methods based on deforming templates [5,18,24]. By incorporating uncertainty, our approach improves the robustness and downstream performance of cortical surface analysis on clinical MRI, facilitating large-scale neuroimaging studies on uncured, heterogeneous clinical datasets.




Fig. 4. Effect size (ES) as a function of the kept dataset percentage based on uncertainty ($\mathcal{U}^s$). (Left) Mean ES across all parcels for Recon-All-Clinical (RAC) [15] for both FL and T1 MRIs. The red dashed line represents ES for Recon-All – RA. (Right) ES per parcel for RA and RAC.

Acknowledgment

R.M. is funded through the European Union’s Horizon Europe research and innovation programme under grant agreement 10108030. B.G. acknowledges support from the Royal Academy of Engineering as part of his Kheiron Medical Technologies/RAEng Research Chair in Safe Deployment of Medical Imaging AI. K.G and J.E.I are primarily funded by the National Institute of Aging (1R01AG070988, 1RF1AG080371 and 1R21NS138995). Further support is provided by, BRAIN Initiative (1RF1MH123195, 1UM1MH130981), National Institute of Biomedical Imaging and Bioengineering (1R01EB031114).

Disclosure of interests

B.G. is part-time employee of DeepHealth. No other competing interests.

References

1. Ashukha, A., Lyzhov, A., Molchanov, D., Vetrov, D.: Pitfalls of in-domain uncertainty estimation and ensembling in deep learning. ICLR (2020)
2. Band, N., Rudner, T.G., Feng, Q., Filos, A., Nado, Z., Dusenberry, M.W., Jerfel, G., Tran, D., Gal, Y.: Benchmarking bayesian deep learning on diabetic retinopathy detection tasks. NeurIPS Benchmarking track (2022)
3. Baumgartner, C.F., Tezcan, K.C., Chaitanya, K., Hötter, A.M., Muehlematter, U.J., Schawkat, K., Becker, A.S., Donati, O., Konukoglu, E.: Phiseg: Capturing uncertainty in medical image segmentation. In: MICCAI. pp. 119–127 (2019)
4. 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** (2023)
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: CVPR. pp. 20773–20783 (2022)
6. Cruz, R.S., Lebrat, L., Bourgeat, P., et al.: DeepCSR: A 3D deep learning approach for cortical surface reconstruction. In: WACV. pp. 806–815 (2021)

7. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical image analysis* **57**, 226–236 (2019)
8. Durso-Finley, J., Falet, J.P., Mehta, R., Arnold, D.L., Pawlowski, N., Arbel, T.: Improving image-based precision medicine with uncertainty-aware causal models. In: *MICCAI*. pp. 472–481. Springer (2023)
9. Fischl, B., Sereno, M., Dale, A.M.: Cortical surface-based analysis: II: inflation, flattening, and a surface-based coordinate system. *Neuroimage* **9**(2), 195–207 (1999)
10. Fuchs, M., Gonzalez, C., Mukhopadhyay, A.: Practical uncertainty quantification for brain tumor segmentation. In: *Medical Imaging with Deep Learning* (2021)
11. Gal, Y., Ghahramani, Z.: Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In: *ICML*. pp. 1050–1059 (2016)
12. Glasser, M., Sotiropoulos, S., Wilson, J.A., Coalson, T.S.: The minimal preprocessing pipelines for the Human Connectome Project. *Neuroimage* **80**, 105–24 (2013)
13. Gopinath, K., Desrosiers, C., Lombaert, H.: SegRecon: Learning joint brain surface reconstruction and segmentation from images. In: *MICCAI*. pp. 650–659 (2021)
14. Gopinath, K., Greve, D.N., Das, S., Arnold, S., Magdamo, C., Iglesias, J.E.: Cortical analysis of heterogeneous clinical brain mri scans for large-scale neuroimaging studies. In: *MICCAI*. pp. 35–45. Springer (2023)
15. Gopinath, K., Greve, D.N., Magdamo, C., Arnold, S., Das, S., Puonti, O., Iglesias, J.E.: Recon-all-clinical: Cortical surface reconstruction and analysis of heterogeneous clinical brain mri. *arXiv preprint arXiv:2409.03889* (2024)
16. Gopinath, K., Hoopes, A., et al.: Synthetic data in generalizable, learning-based neuroimaging. *Imaging Neuroscience* **2**, 1–22 (2024)
17. Henschel, L., Kügler, D., Reuter, M.: Fastsurfervinn: Building resolution-independence into deep learning segmentation methods—a solution for highres brain mri. *NeuroImage* **251**, 118933 (2022)
18. 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)
19. Hu, X., Gopinath, K., Liu, P., et al.: Hierarchical uncertainty estimation for learning-based registration in neuroimaging. *ICLR* (2025)
20. Huang, L., Ruan, S., Xing, Y., Feng, M.: A review of uncertainty quantification in medical image analysis: Probabilistic and non-probabilistic methods. *Medical Image Analysis* p. 103223 (2024)
21. Iglesias, J.E., Billot, B., Balbastre, Y., Magdamo, C., Arnold, et al.: SynthSR: A public AI tool to turn heterogeneous clinical brain scans into high-resolution T1-weighted images for 3D morphometry. *Science advances* **9**(5), eadd3607 (2023)
22. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., et al.: The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *Journal of Magnetic Resonance Imaging* **27**(4), 685–691 (2008)
23. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *NeurIPS*. vol. 30 (2017)
24. 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)
25. Mehta, R., Filos, A., et al.: Qu-brats: Miccai brats 2020 challenge on quantifying uncertainty in brain tumor segmentation-analysis of ranking scores and benchmarking results. *The MELBA journal* **2022** (2022)
26. Mehta, R., et al.: Propagating uncertainty across cascaded medical imaging tasks for improved deep learning inference. *IEEE TMI* **41**(2), 360–373 (2021)

27. Nair, T., Precup, D., Arnold, D.L., Arbel, T.: Exploring uncertainty measures in deep networks for multiple sclerosis lesion detection and segmentation. *Medical image analysis* **59**, 101557 (2020)
28. Neal, R.M.: Bayesian learning for neural networks, vol. 118. Springer Science & Business Media (2012)
29. Puonti, O., Iglesias, J.E., Van Leemput, K.: Fast and sequence-adaptive whole-brain segmentation using parametric bayesian modeling. *NeuroImage* **143** (2016)
30. Ren, K., Zou, K., Liu, X., Chen, Y., Yuan, X., Shen, X., Wang, M., Fu, H.: Uncertainty-informed mutual learning for joint medical image classification and segmentation. In: MICCAI. pp. 35–45. Springer (2023)
31. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241 (2015)
32. Salat, D.H., Buckner, R.L., Snyder, A.Z., Greve, D.N., Desikan, R.S., Busa, E., Morris, J.C., Dale, A.M., Fischl, B.: Thinning of the cerebral cortex in aging. *Cerebral cortex* **14**(7), 721–730 (2004)
33. Shattuck, D.W., Leahy, R.M.: Brainsuite: an automated cortical surface identification tool. *Medical image analysis* **6**(2), 129–142 (2002)
34. Smith, L., Gal, Y.: Understanding measures of uncertainty for adversarial example detection. *arXiv preprint arXiv:1803.08533* (2018)
35. Tanno, R., Worrall, D.E., Kaden, E., Alexander, D.C.: Uncertainty modelling in deep learning for safer neuroimage enhancement: Demonstration in diffusion mri. *NeuroImage* **225** (2020)
36. Tezcan, K.C., Karani, N., Baumgartner, C.F., Konukoglu, E.: Sampling possible reconstructions of undersampled acquisitions in mr imaging with a deep learned prior. *IEEE Transactions on Medical Imaging* **41**(7), 1885–1896 (2022)
37. Wang, G., Li, W., Aertsen, M., Deprest, J., Ourselin, S., Vercauteren, T.: Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing* **338**, 34–45 (2019)
38. Wu, W., Qu, R., Shi, D., Xiong, T., Xu, X., Xing, X., Zhang, X.: Cortical surface reconstruction from 2d mri with segmentation-constrained super-resolution and representation learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 99–108. Springer (2024)
39. Zheng, H., Li, H., Fan, Y.: Coupled reconstruction of cortical surfaces by diffeomorphic mesh deformation. *NeurIPS* **36** (2024)