

Patient-Specific Real-Time Segmentation in Trackerless Brain Ultrasound

Reuben Dorent¹, Erickson Torio¹, Nazim Haouchine¹, Colin Galvin¹, Sarah Frisken¹, Alexandra Golby¹, Tina Kapur¹, and William Wells^{1,2}

¹ Harvard Medical School, Brigham and Women’s Hospital, Boston, MA, USA

² Massachusetts Institute of Technology, Cambridge, MA, USA

rdorent@bwh.harvard.edu

Abstract. Intraoperative ultrasound (iUS) imaging has the potential to improve surgical outcomes in brain surgery. However, its interpretation is challenging, even for expert neurosurgeons. In this work, we designed the first patient-specific framework that performs brain tumor segmentation in trackerless iUS. To disambiguate ultrasound imaging and adapt to the neurosurgeon’s surgical objective, a patient-specific real-time network is trained using synthetic ultrasound data generated by simulating virtual iUS sweep acquisitions in pre-operative MR data. Extensive experiments performed in real ultrasound data demonstrate the effectiveness of the proposed approach, allowing for adapting to the surgeon’s definition of surgical targets and outperforming non-patient-specific models, neurosurgeon experts, and high-end tracking systems. Our code is available at: <https://github.com/ReubenDo/MHVAE-Seg>.

Keywords: Intraoperative Ultrasound · Image Segmentation · Cross-Modal Synthesis · Neurosurgery

1 Introduction

Intraoperative ultrasound (iUS) has recently raised considerable interest as it is an affordable, real-time, intraoperative imaging technology that can be easily integrated into existing surgical workflows. Despite these advantages, the interpretation of iUS imaging is challenging due to the presence of artifacts, oblique orientations, reduced field of view, and low and variable contrast between different tissues, which notably complicates the identification of brain tumor margins and potential areas of residual tumor during neurosurgery [17].

To improve the interpretability of iUS, efforts have been made to overlay pre-operative data with iUS images using tracked iUS probes. However, relying on complex and costly navigation systems to track iUS detracts from the inherent simplicity and low cost of ultrasound imaging. Moreover, neuronavigation based on pre-operative data becomes invalid as the surgery progresses due to brain shift and non-linear deformations caused by tissue resection, leading to misalignments between iUS and overlaid pre-operative data [14], as shown in Figure 1. For these reasons, the use of tracked ultrasound systems remains uncommon.




Fig. 1. An illustrative example showing a) the variability between neurosurgeon and BraTS annotation protocols; b) the challenge of identifying tumor boundaries in untracked 2D iUS; c) misalignment overlay issues of tracked 2D iUS with navigation system (red arrows); d) predictions of our patient-specific models in untracked 2D iUS.

Alternatively, image segmentation methods have been proposed to analyze iUS images directly [15,10,1,18,6]. Traditionally, this involves supervised learning with expert-annotated ultrasound images. However, segmenting ultrasound images demands rare clinical expertise and even in expert hands is subject to inter-rater variability [19]. Consequently, the size of publicly available datasets is limited [5], affecting the generalization capability and performance of these methods. Moreover, to handle data ambiguity in ultrasound imaging, segmentation often relies on 3D reconstructed images from tracking systems [10,1], not native 2D iUS images, as it leads to significant performance improvements [15,6]. However, this introduces two major limitations: the reliance on navigation systems and the inability to operate in real-time due to the reconstruction process.

Another significant challenge in the surgical context is the selection of an appropriate annotation protocol. For instance, the BraTS protocol [3] for brain tumors in MRI combines tumor-induced edema and infiltrated tissues into a single category, which often does not correspond to the neurosurgeon’s delineation of resectable areas. Furthermore, determining areas deemed safe for resection depends on the location of critical surrounding structures. Consequently, surgical target definitions vary across subjects and surgeons. This variability highlights the need for patient-specific segmentation strategies tailored to the patient’s anatomy, and consistent with the surgeon’s pre-operative planning.

In this work, we propose a novel segmentation approach for iUS imaging that addresses the challenges of data disambiguation and the need for patient-specificity. To our knowledge, this is the first work that segments trackerless 2D brain iUS by leveraging pre-operative data. This approach involves training a segmentation model specific to each patient by generating 2D synthetic iUS images using surgical planning data. Specifically, we assume that pre-operative MR data and the segmentation of the surgical target are available prior to surgery, which is typically the case in clinical routine. Our contribution is three-fold. First, we designed a novel approach to generate virtual ultrasound sweeps in pre-operative




Fig. 2. Overview of our approach to create synthetic iUS images from pre-operative MRI. First, a generic sweep is virtually positioned on MR brain surface to simulate iUS acquisition. Then, MHVAE [7] synthesizes iUS from generated 2D pre-op MRI slices.

MR data space, effectively mimicking the iUS acquisition process. Second, we exploit state-of-the-art Multi-Modal Hierarchical Variational Auto-Encoder [7] to synthesize iUS data from 2D MR slices in the virtual sweeps’ field of view. To generate variability in our synthetic data, we take advantage of the stochasticity of MHVAE and its ability to handle incomplete sets of MR data. Third, extensive experiments were conducted on real ultrasound data, highlighting the need for patient-specificity and demonstrating the effectiveness of our approach by outperforming experts and tracking-system-based segmentation methods.

2 Patient-specific segmentation in brain ultrasound

Problem overview. The overall objective is to train a segmentation network f_{θ^s} parameterized by weights θ^s that automatically segments 2D brain ultrasounds for a given subject s . As the acquisition of adult brain ultrasound images requires skull opening (craniotomy), ultrasound images are not available prior to the surgery. Instead, we assume that we have access to a set of M multi-parametric co-registered pre-operative MRI data $I^s = \{I_j^s\}_{j=1}^M$ with the segmentation of the surgical target T^s . In this work, we propose to train the patient-specific segmentation network f_{θ^s} using a set of synthetic 2D ultrasound images generated from pre-operative imaging data (I^s, T^s) . Our approach involves three steps: 1) simulating the positioning and trajectory of an ultrasound probe on the surface of the brain (sweep); 2) synthesizing iUS images using pre-operative MR data in the field of view of a virtual sweep; and 3) training the segmentation network f_{θ^s} using the produced synthetic iUS images and MR-based annotations.

2.1 Step 1: Simulation of ultrasound acquisitions

In this section, we introduce a novel approach for simulating the acquisition of ultrasound sweeps derived from preoperative MRI data with segmented brain

tumors. To accurately position the virtual iUS probe on the brain surface of our patient data I^s and simulate the probe’s trajectory during acquisition, we assume that: 1) the probe should be positioned on the brain surface near the tumor; 2) the sweep’s trajectory should cover as much of the tumor as possible.

First, a 3D ultrasound sweep acquired during the surgery of a previous patient is randomly selected to serve as the reference for the acquisition trajectory and probe geometry. The principal trajectory direction of this reference sweep is represented by the vector $\mathbf{n}_1 \in \mathbb{R}^3$. The extremities of the probe on the median 2D slice, orthogonal to the direction of the sweep, are denoted as $\mathbf{L}_1 \in \mathbb{R}^3$ and $\mathbf{R}_1 \in \mathbb{R}^3$ for the left and right sides, respectively, as illustrated in Figure 2. The next objective is to find a rigid transformation R to place this 3D sweep in the space of the patient data I^s .

To ensure that the virtual sweep’s trajectory covers most of the tumor, we postulate that the primary direction of the virtual ultrasound sweep in the MR space should align with the first principal component $\mathbf{n}_2 \in \mathbb{R}^3$ of the tumor’s spatial distribution. Therefore, the virtual median 2D slice is within a plane \mathcal{P} orthogonal to the vector \mathbf{n}_2 .

To fully characterize the plane \mathcal{P} of the virtual median slice, an additional point \mathbf{C}_2 within this plane must be identified. Since the objective of the iUS acquisition is to visualize the surgical target, we propose that the plane \mathcal{P} passes through a point \mathbf{C}_2 on the cortical surface near the surgical target. To select a candidate \mathbf{C}_2 , we first automatically segment the brain in I^s [12] and extract the set of points on the cortical surface $\mathcal{S} = \{\mathbf{S}_i\}$. Then, we randomly select a point $\mathbf{C}_2 \in \mathcal{S}$, with the probability inversely proportional to the exponential of the Euclidean distance to the tumor’s centroid \mathbf{M}_2 in the segmentation T^s :

$$\mathbf{S}_i \in \mathcal{S}, p(\mathbf{C}_2 = \mathbf{S}_i) \propto \exp(-\|\mathbf{S}_i - \mathbf{M}_2\|^2). \quad (1)$$

Thus, the virtual median 2D slice is in the plane \mathcal{P} passing through \mathbf{C}_2 and orthogonal to \mathbf{n}_2 .

To accurately position the virtual median 2D slice in the plane \mathcal{P} , we aim to determine the optimal locations of the virtual median slice extremities \mathbf{L}_2 and \mathbf{R}_2 with three criteria: \mathbf{L}_2 and \mathbf{R}_2 should be 1/ on the cortical surface; 2/ at a similar distance to \mathbf{C}_2 , 3/ their Euclidean distance should be equal to the Euclidean distance between the extremities of the median 2D slice, i.e. $\|\mathbf{L}_2 - \mathbf{R}_2\| = \|\mathbf{L}_1 - \mathbf{R}_1\|$. This optimization problem can be formulated as follows:

$$\begin{cases} \mathbf{L}_2 = \arg \min_{\mathbf{S}_i \in \mathcal{S}} (\|\mathbf{S}_i - \mathbf{C}_2\| - 0.5 \times \|\mathbf{L}_1 - \mathbf{R}_1\|), \\ \mathbf{R}_2 = \arg \min_{\mathbf{S}_i \in \mathcal{S}} (\|\mathbf{S}_i - \mathbf{L}_2\| - \|\mathbf{L}_1 - \mathbf{R}_1\|). \end{cases} \quad (2)$$

Finally, the rigid transformation R that places the reference sweep in the MR data can be estimated by aligning the sets $(\mathbf{L}_1, \mathbf{R}_1, \mathbf{n}_1)$ from the reference sweep and $(\mathbf{L}_2, \mathbf{R}_2, \mathbf{n}_2)$ from the MR data using the Least-Squares Fitting algorithm[2].

2.2 Step 2: Synthesizing ultrasound data from MR data

To create synthetic MR 2D slices, we apply the rigid transformation R to the reference sweep, enabling the generation of a series of 2D MR slices within the

field of view of an intraoperative ultrasound probe. An illustrative example of a 2D MR slice is presented in Figure 2. To synthesize iUS images from these 2D MR slices, we utilize the Multi-Modal Hierarchical Variational Auto-Encoder (MHVAE), the current state-of-the-art MRI to iUS synthesis framework [7]. The MHVAE architecture leverages a hierarchical latent space to capture complex and multi-level multi-modal data representations. Moreover, MHVAE has the flexibility to perform image synthesis in the presence of incomplete MR data. We use the pre-trained MHVAE model that performs iUS synthesis from any combination of the primary MR sequences for brain tumor surgery: contrast-enhanced T1 (ceT1), T2, and T2 Fluid Attenuation Inversion Recovery (FLAIR). Figure 2 shows an example of synthetic iUS generated using a 2D T2 slice.

To introduce variability in our synthetic data, we employ two strategies. First, we generate iUS by randomly omitting sequences from the MR data inputs. Variations in synthetic iUS are observed for different sets of input MR sequences, as shown in supplementary materials. Second, we exploit the stochastic and probabilistic nature of MHVAE to generate variability in the images. MHVAE involves sampling from Gaussian distributions at each level of the hierarchy. To introduce local variability in the generated iUS, we draw samples with various sampling temperatures. This modulates the synthetic images’ local aspect, with more speckles for higher temperatures, as shown in supplementary materials.

2.3 Step 3: Training a patient-specific segmentation network

Training set. To train the patient-specific segmentation network, we use the two previously described steps to create a paired dataset of synthetic ultrasound with manual annotations. Given the pre-operative MR data I^s and manual segmentation T^s , for each possible combination of input MR images I_π^s (6 combinations in total), we randomly select K reference sweeps from other cases and simulate an ultrasound acquisition for each of them. Note that each acquisition simulation leads to a different positioning due to the stochastic selection of the cortical point C_2 . This process results in the creation of a paired set of MR 2D slices and 2D surgical target segmentations for each virtual iUS sweep acquisition. Then, we generate synthetic iUS 2D slices using MHVAE with 4 different sampling temperatures $\tau \in \{0.3, 0.5, 0.7, 1.0\}$. In total, series of synthetic 2D iUS and surgical target annotations from $24K$ virtual sweeps are created.

Network training. Thanks to this patient-specific synthetic dataset, we can train the patient-specific segmentation network in a supervised manner. To perform real-time inference, we use a standard 2D Unet network presented in supplementary materials. Given the challenge of distinguishing tumor boundaries in iUS images, the network is deliberately trained to “overfit” to the surgical target segmentation. However, to ensure generalization to real iUS images, extensive data augmentation techniques are implemented, including rotations, elastic deformations, scaling, mirroring, and the application of additive Gaussian noise, as well as adjustments to brightness, contrast, and gamma. At each

training iteration, a complete series of synthetic iUS from a randomly selected virtual sweep is fed to the network, setting the batch size to approximately 100 2D images. Optimization is conducted using deep supervision with the Dice loss function, with parameters θ^s optimized through stochastic gradient descent employing Nesterov momentum (momentum set to 0.9). The model is trained for 100 epochs with an initial learning rate of 0.01, which gradually decreases to zero following a polynomial learning rate schedule. A 24GB NVIDIA A5000 GPU was used for model training. Our code is publicly available at: <https://github.com/ReubenDo/MHVAE-Seg>.

3 Experiments

To evaluate the performance of our proposed patient-specific approach, extensive experiments were conducted on real 2D iUS images from 7 cases acquired during brain surgery. Three sets of brain tumor annotations (two neurosurgeons, automated) were used for evaluation. Our patient-specific framework was compared with non-patient-specific automated methods, pre-operative segmentation overlaid by navigation systems, and manual expert iUS segmentation.

Dataset. The publicly available ReMIND dataset [13] was used in our experiments. 7 cases not included during MHVAE’s training were randomly selected. Each case comprised all pre-operative MR sequences (ceT1, T2, and FLAIR) and 3D pre-dura iUS reconstructed from sweeps of a tracked handheld 2D probe, with one case lacking the FLAIR sequence. Note that MR images were acquired at different institutions with different protocols [13]. The testing set corresponds to 2D slices obtained by slicing the 3D ultrasound sweeps in the direction of sweep acquisition, resulting in an average of 70 iUS 2D slices per case. All images were resampled to an isotropic 0.5mm resolution, padded for a matrix size of (192, 192), and normalized in $[-1, 1]$.

Pre-operative segmentations. Three sets of pre-operative MR-based segmentations were used in our experiments. Two neurosurgeons (N_1 and N_2) manually segmented the surgical target in either pre-operative ceT1 or T2 images, depending on the tumor grade and location. The segmentations by N_1 correspond to publicly available ReMIND annotations acquired during surgical planning. Additionally, to simulate the *BraTS* protocol, automated segmentation was conducted on ceT1 and T2 images using nnUnet [11], trained on BraTS data [3]. These MR-based annotations served as the training segmentation data for our framework, which is specific to the case and the annotation protocol.

Evaluation. Since the segmentation of iUS is error-prone and ambiguous, groundtruth 2D iUS segmentations were instead obtained by propagating pre-operative MR-based segmentations via image registration between MR and 3D




Fig. 3. Examples of automated surgical target segmentations using non-patient-specific models (*RESECT Unet* and *BraTS Unet*), high-end tracking system and *Ours*. MR scan used for manual surgical target annotation is shown.

iUS. This process involved affinely registering pre-operative scans with 3D pre-dura iUS using NiftyReg [16], following the pipeline described in [9]. The registration outputs were checked by three neurological experts. To assess the accuracy of each segmentation method, Dice Similarity Coefficient (Dice) and Average Symmetric Surface Distance (ASSD) were used, with statistical significance assessed via Wilcoxon signed-rank tests ($p < 0.01$).

Determining the number of virtual sweeps. To evaluate the impact of the total number of virtual iUS sweeps ($24K$), we evaluated the framework’s performance with various K values using N_1 ’s annotations. Since the duration of the framework optimization, which includes synthetic iUS dataset generation ($\approx 5\%$) and network training ($\approx 95\%$), linearly increases with K , the development time was tracked. As reported in Table 1, **the level of performance increases with the number of virtual sweeps K , highlighting the need for variability during training.** To balance development time with performance, we set an upper limit of 4 hours for model development, thus fixing $K = 10$ in the next experiments.

Baseline approaches. We compared our patient-specific approach against three baselines. First, we trained two non patient-specific 2D Unet approaches: *RESECT Unet*, using the only publicly available dataset of brain iUS with manual tumor annotations (RESECT-SEG [5], $N = 23$ 3D iUS) and *BraTS Unet*, using a synthetic dataset of iUS generated by applying our proposed method to a large, annotated subset of the BraTS MRI dataset (UPenn-GBM [4], $N = 611$

Table 1. Impact of the number K of generated synthetic iUS in terms of model performance and model development duration using N_2 's segmentations. Median and inter-quartile range are reported. Arrows indicate favorable direction of each metric.

| | $K = 1$ | $K = 3$ | $K = 5$ | $K = 7$ | $K = 10$ | $K = 15$ |
|-----------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| DSC (%)↑ | 79.7 (19.4) | 78.2 (22.9) | 80.5 (19.1) | 83.6 (17.2) | 84.2 (18.1) | 84.9 (17.9) |
| ASSD (mm)↓ | 2.4 (2.2) | 2.4 (2.0) | 2.2 (1.8) | 1.9 (1.7) | 1.9 (1.7) | 1.8 (1.7) |
| Development Duration (min)↓ | 27 (4) | 60 (4) | 103 (8) | 148 (10) | 210 (8) | 309 (23) |

Table 2. Quantitative evaluation of different approaches for surgical target segmentation for three annotation protocols. Median and inter-quartile range are reported. * denotes significant improvement provided by a Wilcoxon test ($p < 0.01$).

| | BraTS (487 2D iUS slices) | | Neurosurgeon 1 (487 2D iUS slices) | | Neurosurgeon 2 (70 2D iUS slices) | |
|--------------------|------------------------------|-------------------|---------------------------------------|-------------------|--------------------------------------|-------------------|
| | DSC (%)↑ | ASSD (mm)↓ | DSC (%)↑ | ASSD (mm)↓ | DSC (%)↑ | ASSD (mm)↓ |
| RESECT Unet | 71.7 (17.4) | 4.2 (3.5) | 59.8 (33.2) | 5.4 (3.5) | 58.5 (30.4) | 5.3 (2.8) |
| Brats Unet | 76.8 (12.8) | 3.9 (2.4) | 56.1 (28.8) | 6.0 (3.1) | 53.8 (30.2) | 5.8 (3.8) |
| Tracking System | 85.4 (9.2) | 1.9* (0.7) | 80.0 (16.9) | 2.2 (1.4) | 80.1* (16.4) | 2.2* (1.0) |
| Manual | × | × | × | × | 67.4 (29.4) | 3.0 (2.2) |
| Ours (K=10) | 87.2* (8.5) | 2.1 (1.7) | 84.2* (18.1) | 1.9* (1.7) | 84.2* (36.6) | 1.8* (2.4) |

cases). Second, we compared our trackless method to the patient-specific segmentation derived from the high-end optical tracking of the ‘‘Curve’’ navigation system (Brainlab AG). Third, an experienced neurosurgeon (N_2) with expertise in iUS imaging manually segmented the surgical target in preoperative MRI and subsequently in 10 random iUS slices per case.

Results. Quantitative and qualitative results are shown in Table 2 and Figure 3. Firstly, we can observe that *RESECT Unet* and *BraTS Unet* obtained comparable performance on the ReMIND dataset, showing that our framework circumvents the need for iUS manual annotations. Secondly, the comparison between *Ours* and *BraTS Unet* highlights the need for patient-specificity. On the one hand, leveraging pre-operative MR data allows disambiguating iUS images, as shown in Figure 3, significantly increasing the Dice by +11.4pp and reducing the ASSD by a factor of 2 for BraTS annotations. On the other hand, the performance of *BraTS Unet* drops when compared with N_1 and N_2 's annotations, illustrating the variability in annotation protocols. In contrast, our patient-specific approach obtained stable results across protocols, demonstrating the advantage of leveraging neurosurgeon pre-operative planning during training. Third, our trackerless approach surprisingly reached comparable performance to the tracking-system-based method, and even outperformed it in some cases. This suggests our approach may provide a viable alternative to expensive and complex tracking systems without compromising performance. Finally, the mediocre results achieved by the iUS expert emphasize the inherent challenges of segmenting brain surgical targets in iUS. Notably, our approach significantly outperforms

the expert neurosurgeon, suggesting its potential utility in assisting surgeons. Overall these experiments demonstrate the effectiveness of our approach, which leverages a real-time 2D Unet (200 FPS) and pre-operative data, to create a patient-specific segmentation model adapted to the neurosurgeon.

4 Conclusion

In this work, we introduced a novel approach to design patient-specific models on the challenging problem of surgical target segmentation in brain intraoperative ultrasound. The proposed framework can be seen as a patient-specific cross-modality domain adaptation approach [8]. Our approach, which only requires access to pre-operative planning, could be a viable alternative to complex tracking systems to assist with iUS interpretation. In line with recent medical imaging research, this work advocates for patient-specific models leveraging pre-operative data to improve performance during surgery. Future work will expand automated segmentation to additional brain structures, incorporate temporal consistency during iUS sweeps acquisition and focus on reducing the development duration. From an application perspective, we will explore whether our framework can automatically detect significant misalignments between pre-operative and intra-operative data. This could alert surgeons to inaccuracies in the navigation system.

Acknowledgement This work was supported by the National Institutes of Health (R01EB032387, R01EB027134, P41EB015902, P41EB028741 and K25EB035166). We also thank the philanthropic support from the Jennifer Oppenheimer Cancer Research Initiative and Karen and Brian McMahon.

References

1. Angel-Raya, E., Chalopin, C., Avina-Cervantes, J.G., Cruz-Aceves, I., Wein, W., Lindner, D.: Segmentation of brain tumour in 3d intraoperative ultrasound imaging. *The International Journal of Medical Robotics and Computer Assisted Surgery* **17**(6), e2320 (2021)
2. Arun, K.S., Huang, T.S., Blostein, S.D.: Least-squares fitting of two 3-D point sets. *IEEE Transactions on pattern analysis and machine intelligence* (5), 698–700 (1987)
3. Bakas, S., et al: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge (2019)
4. Bakas, S., Sako, C., Akbari, H., Bilello, M., Sotiras, A., Shukla, G., Rudie, J.D., Santamaría, N.F., Kazerooni, A.F., Pati, S., et al.: The University of Pennsylvania glioblastoma (UPenn-GBM) cohort: Advanced MRI, clinical, genomics, & radiomics. *Scientific data* **9**(1), 453 (2022)
5. Behboodi, B., Carton, F.X., Chabanas, M., De Ribaupierre, S., Solheim, O., Munkvold, B.K., Rivaz, H., Xiao, Y., Reinertsen, I.: Resect-seg: Open access annotations of intra-operative brain tumor ultrasound images. *arXiv preprint arXiv:2207.07494* (2022)

6. Carton, F.X., Chabanas, M., Le Lann, F., Noble, J.H.: Automatic segmentation of brain tumor resections in intraoperative ultrasound images using u-net. *Journal of Medical Imaging* **7**(3), 031503–031503 (2020)
7. Dorent, R., Haouchine, N., Kogl, F., Joutard, S., Juvekar, P., Torio, E., Golby, A.J., Ourselin, S., Frisken, S., Vercauteren, T., et al.: Unified brain mr-ultrasound synthesis using multi-modal hierarchical representations. In: *International conference on medical image computing and computer-assisted intervention*. pp. 448–458. Springer (2023)
8. Dorent, R., Kujawa, A., Ivory, M., Bakas, S., Rieke, N., et al.: CrossMoDA 2021 challenge: Benchmark of cross-modality domain adaptation techniques for vestibular schwannoma and cochlea segmentation. *Medical Image Analysis* **83** (2023)
9. Drobný, D., Vercauteren, T., Ourselin, S., Modat, M.: Registration of MRI and iUS data to compensate brain shift using a symmetric block-matching based approach. In: *CuRIOUS* (2018)
10. Ilunga-Mbuyamba, E., Avina-Cervantes, J.G., Lindner, D., Arlt, F., Ituna-Yudonago, J.F., Chalopin, C.: Patient-specific model-based segmentation of brain tumors in 3d intraoperative ultrasound images. *International journal of computer assisted radiology and surgery* **13**, 331–342 (2018)
11. 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)
12. Isensee, F., Schell, M., Pflueger, I., Brugnara, G., Bonekamp, D., Neuberger, U., Wick, A., Schlemmer, H.P., Heiland, S., Wick, W., et al.: Automated brain extraction of multisequence MRI using artificial neural networks. *Human brain mapping* **40**(17), 4952–4964 (2019)
13. Juvekar, P., Dorent, R., Kögl, F., Torio, E., Barr, C., Rigolo, L., Galvin, C., Jowkar, N., Kazi, A., Haouchine, N., Cheema, H., Navab, N., Pieper, S., Wells, W.M., Bi, W.L., Golby, A., Frisken, S., Kapur, T.: ReMIND: The Brain Resection Multimodal Imaging Database. medRxiv (2023)
14. Juvekar, P., Torio, E., Bi, W.L., Bastos, D.C.D.A., Golby, A.J., Frisken, S.F.: Mapping resection progress by tool-tip tracking during brain tumor surgery for real-time estimation of residual tumor. *Cancers* **15**(3) (2023). <https://doi.org/10.3390/cancers15030825>
15. Milletari, F., Ahmadi, S.A., Kroll, C., Plate, A., Rozanski, V., Maiostre, J., Levin, J., Dietrich, O., Ertl-Wagner, B., Bötzel, K., et al.: Hough-cnn: Deep learning for segmentation of deep brain regions in mri and ultrasound. *Computer Vision and Image Understanding* **164**, 92–102 (2017)
16. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. *Computer Methods and Programs in Biomedicine* **98** (2010)
17. Šteňo, A., Buvala, J., Babková, V., Kiss, A., Toma, D., Lysak, A.: Current limitations of intraoperative ultrasound in brain tumor surgery. *Frontiers in Oncology* **11**, 659048 (2021)
18. Valanarasu, J.M.J., Yasarla, R., Wang, P., Hacihaliloglu, I., Patel, V.M.: Learning to segment brain anatomy from 2d ultrasound with less data. *IEEE Journal of Selected Topics in Signal Processing* **14**(6), 1221–1234 (2020)
19. Weld, A., Dixon, L., Anichini, G., Patel, N., Nimer, A., Dyck, M., O’Neill, K., Lim, A., Giannarou, S., Camp, S.: Challenges with segmenting intraoperative ultrasound for brain tumours. medRxiv (2023). <https://doi.org/10.1101/2023.12.13.23299820>

5 Supplementary materials




Fig. 4. Impact of the temperature τ on the local variability (speckles).




Fig. 5. Our 2D Unet segmentation network.