A Multi-rater Comparative Study of Automatic Target Localization Methods for Epilepsy Deep Brain Stimulation Procedures

Han Liu[a], Kathryn L. Holloway[b], Dario J. Englot[c], and Benoit M. Dawant[d]

[a]Dept. of Computer Science, Vanderbilt University, Nashville, TN 37235, USA
[b]Dept. of Neurosurgery, Virginia Commonwealth University, Richmond, VA 23284, USA
[c]Dept. of Neurosurgery, Vanderbilt University Medical Center, Nashville, TN 37235, USA
[d]Dept. of Electrical and Computer Engineering, Vanderbilt University, Nashville, TN 37235, USA

ABSTRACT

Epilepsy is the fourth most common neurological disorder and affects people of all ages worldwide. Deep Brain Stimulation (DBS) has emerged as an alternative treatment option when anti-epileptic drugs or resective surgery cannot lead to satisfactory outcomes. To facilitate the planning of the procedure and for its standardization, it is desirable to develop an algorithm to automatically localize the DBS stimulation target, i.e., Anterior Nucleus of Thalamus (ANT), which is a challenging target to plan. In this work, we perform an extensive comparative study by benchmarking various localization methods for ANT-DBS. Specifically, the methods involved in this study include traditional registration method and deep-learning-based methods including heatmap matching and differentiable spatial to numerical transform (DSNT). Our experimental results show that the deep-learning (DL)-based localization methods that are trained with pseudo labels can achieve a performance that is comparable to the inter-rater and intra-rater variability and that they are orders of magnitude faster than traditional methods.

Keywords: Epilepsy, Deep Brain Stimulation, Anterior Nucleus of Thalamus, Localization, Heatmap Matching, DSNT

1. INTRODUCTION

Epilepsy is one of the leading neurological causes of loss of quality-adjusted life years and affects 65 million people worldwide.[1] It has been shown that despite the use of first-choice anti-epileptic drugs and satisfactory seizure outcome rates after resective epilepsy surgery, a considerable percentage of epilepsy patients suffer with persistent seizures.[2] ANT-DBS has emerged as an alternative treatment option for these patients. DBS procedures involve the delivery of a predetermined program of electrical stimulation to deep brain structures via implanted electrodes connected to a pulse generator.[3] One major challenge with this procedure is to determine an optimal trajectory by placing the implant at the proper stimulation target while avoiding sensitive structures such as brain vessels.[5]


Figure 1. An example of ANT-DBS target on the right thalamus. The thalamus mask is displayed in blue, and the target is denoted as the bright spot.

For epilepsy patients, ANT (the Anterior Nucleus of the Thalamus) is the stimulation target as it can influence the brain's predisposition to epileptic seizures.[4] To standardize and facilitate placement and improve outcomes, it is desirable to develop an automatic algorithm for accurate ANT target localization.

The standard approach for automatic target localization is the atlas-based registration. However, this approach suffers from a lack of robustness when the anatomic differences between source and target images are large or when the number of atlases is limited.[7] This can be of importance for our application due to the well documented variability in the ANT anatomy and thalamic atrophy caused by persistent epileptic seizures.[11] Besides, classic registration methods are not very efficient as they require long processing time. Recently, a DL-based approach has been proposed to localize the ANT target via heatmap matching (HM).[12] In this method, heatmaps are generated by rendering a spherical 3D Gaussian centered on the coordinates of pseudo labels and the Mean Squared Error (MSE) loss between the output and synthetic heatmap is minimized during training. During inference, the final ANT targets are obtained from the model's output by computing the argmax of the voxel probabilities. However, as suggested by Nibali *et al.*,[13] the major limitation of HM is that gradient flow begins at the heatmap rather than the numerical coordinates, leading to a mismatch between the optimization objective, i.e., the similarity between heatmaps, and the evaluation metric for target localization, i.e., the distance between the coordinates of the predicted target and the ground truth. To overcome this limitation, a DSNT[13] (differentiable spatial to numerical transform) layer has been proposed to fully preserve the spatial generalization as well as the end-to-end differentiability of the model. In this work, we conduct a comparative study to investigate the ANT target localization performances for DBS with different localization approaches including (1) traditional registration method (2) HM and (3) DSNT. For learning-based methods, we also explore the impact of the amount of training data on the localization performance. Lastly, we compare these results to the manual inter/intra-rater variability observed in a study conducted across two institutions.

2. METHODS

2.1 Data and Materials

Our dataset consists of 309 T1-weighted MRI images from a database of patients who underwent a DBS implantation for movement disorders, i.e., Parkinson Disease or Essential Tremor at Vanderbilt University. The in-plane resolution varies from 0.4356 to 1 mm and slice thickness from 1 to 6 mm. We randomly split the entire dataset into 259 images for training and validation and the rest 50 images for testing. The ANT target on one atlas from a different set is manually annotated by an experienced neurosurgeon. Atlas-based registration[10] is applied between the atlas image and each training sample to generate **pseudo labels**, i.e., 3D coordinates. For all images in our testing set, two experienced neurosurgeons from two different institutions are asked to annotate the ANT targets on both left and right thalamus. One of the neurosurgeons is asked to annotate the testing set twice for intra-rater variability analysis. The 3D coordinates of the annotated targets are collected using an open-source software ITK-SNAP.[6] For image preprocessing, we resample all the images to isotropic voxel size of $1 \times 1 \times 1$ mm^3 by trilinear interpolation and normalize the image intensities to the range of 0 and 1.

2.2 Overview

In this study, we explore several target localization methods for the ANT-DBS application including (1) traditional registration, (2) HM and (3) DSNT. We follow the two-stage strategy proposed in Liu *et al.*,[12] to localize the targets in a coarse-to-fine manner. The entire pipeline of our localization process is illustrated in Figure 2. Specifically, given an unseen testing image, or query image, we detect and crop a region of interest (ROI) around the thalamus in the first stage and localize the ANT target within the cropped ROI in the second stage. For the first stage, to reduce the inference time, we train a 3D U-Net[14] denoted as F_{seg} to segment the left and right thalamus from downsampled MRI scans and then crop an ROI with a size of $64 \times 64 \times 64$ centered at the segmented thalamus. Once we collect the left and right ROIs of the entire dataset, we flip the left ROIs to the right such that all the cropped ROIs have consistent image contents. The consistency of input data reduces difficulty of the downstream task and is thus beneficial for our training. These ROIs are ready to be used as the input volumes for the second-stage localization network. In the following sections, we will discuss each localization method in detail.

Stage II: Localization

Traditional registration

argmax → $(\hat{x}, \hat{y}, \hat{z})$

SyN

Stage I: ROI detection

Atlas

Atlas ROI

Atlas Target

Query image

F_{seg}

Query ROI

F_{hm}

Heatmap matching

Pseudo label supervision

F_{hm}

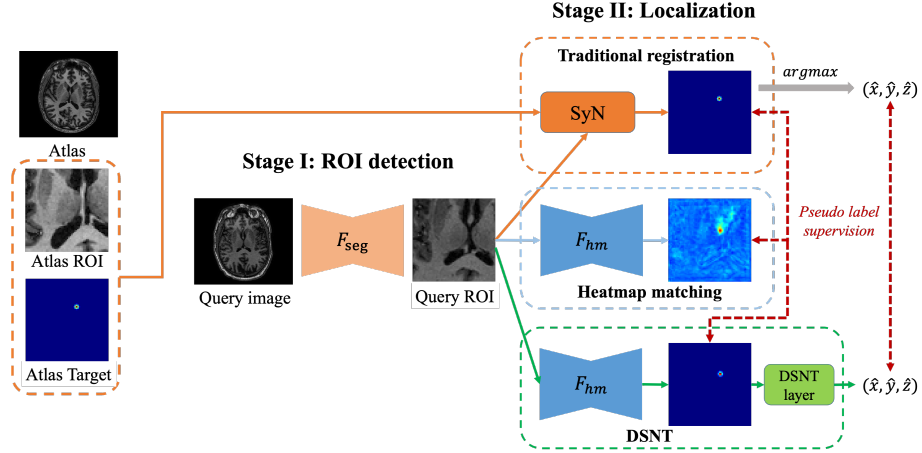DSNT layer → $(\hat{x}, \hat{y}, \hat{z})$

DSNT

Figure 2. An illustration of the two-stage target localization process and the localization methods involved in our comparative study. Note that the F_{hm} in heatmap matching and DSNT share the same network architecture but do not share the weights.

2.3 Registration

Image registration has been widely used to establish a spatial correspondence between biological structures across different images.[8] Registration-based methods are useful for medical image analysis tasks when there are strong anatomical priors in the associated tasks. In our application, the ANT-DBS targets can be localized by registering our atlas ROI to the query ROI. Once the query ROI and the atlas ROI are well-aligned, the obtained deformation field can be used to project the ANT target on the atlas ROI, i.e., the one that is manually annotated by the expert, to the query ROI. Specifically, we use a well-accepted registration algorithm for brain MRI known as SyN[9] to register atlas ROI to each query ROI in the testing set. In implementation, we use the SyN with mutual information as optimization metric in the ANTs[10] package.

2.4 Heatmap matching

For heatmap matching, we first generate synthetic heatmaps y by rendering a spherical 3D Gaussian centered on the pseudo labels; their standard deviation is 1.5 mm. The maximum value of the heatmap is scaled to 1 and any value below 0.05 is set to 0. We train a 3D U-Net F_{hm} to perform a per-voxel regression to perform heatmap matching. A query ROI x is mapped to a probability map $F_{hm}(x)$, where F_{hm} has the same network architecture as the one proposed in Liu *et al.*[12] During training, a weighted MSE (WMSE) loss is used to minimize the heatmap difference between $F_{hm}(x)$ and y. Specifically, in the output probability map, we set the weights associated with positive entries and zero entries from pseudo label maps as W_p and W_z, as described in Equation 1 and 2.

$$W_p = \frac{N_z}{N_p + N_z} \tag{1}$$

$$W_z = \frac{N_p}{N_p + N_z} \tag{2}$$

where N_p and N_z are the number of positive and zero voxels in pseudo label map. During inference, the final predicted target $\widetilde{p}$ is obtained by taking the voxel with the highest activation value from the output probability map, as shown in Equation 3.

$$\widetilde{p} = \arg\max(F_{hm}(x)) \tag{3}$$

2.5 DSNT

To enable end-to-end differentiability for coordinate regression, we build a DSNT-based model F_{dsnt} by adding a DSNT layer on top of the last layer of F_{hm}. The input of the DSNT layer is a normalized localization probability heatmap $F_{hm}(x)$ such that the values of the heatmap are non-negative and sum to 1. In our implementation, we use the softmax activation function to normalize the heatmap. Besides, we extend the original 2D implementation* of DSNT layer to 3D. Specifically, we create three matrices X, Y and Z which contain their own x, y, or z-coordinates and are scaled such that the top-left-front and bottom-right-back corner of the volume are at $(-1, -1, -1)$ at $(1, 1, 1)$, respectively. We can obtain the 3D coordinates $\widetilde{p}$ from the normalized heatmap as described in Equation 4.

$$\widetilde{p} = [\langle F_{hm}(x), X \rangle_F, \langle F_{hm}(x), Y \rangle_F, \langle F_{hm}(x), Z \rangle_F] \tag{4}$$

where $\langle \cdot, \cdot \rangle_F$ denotes the Frobenius inner product. During training, we minimize the MSE loss between the predicted coordinates $\widetilde{p}$ and the ground truth coordinates p as well as a Jensen-Shannon divergence loss to minimize the distribution divergence between the generated heatmap and the ground truth Gaussian distribution, as described in Equation 5.

$$L_{DSNT} = L_{MSE}(\widetilde{p}, p) + \alpha L_{JS}(F_{hm}(x), y) \tag{5}$$

where α is set to be 1 empirically. During inference, we take the argmax of the generated heatmap as our final target because we find that the argmax of the heatmap is slightly more reliable than the coordinate output on our validation set.

2.6 Implementation Details

For pre-processing, we normalize the intensity of each ROI to the range of 0 and 1 before feeding them to the networks. We apply data augmentation to avoid overfitting and improve model generalization ability. Specifically, we augment our training data by random intensity shifts in the range of 0.8 and 1.2, random scaling in the range of 0.9 to 1.1 and random rotation in the range of -10 to 10 degrees along each axis. The best hyperparameters are determined by grid-search within the range of 10^{-2} to 10^{-5}. The best hyperparameters are selected based on the localization performance on the validation set. For training, we use the Adam optimizer[15] with a weight decay of 10^{-4} and a batch size of 16. The CNNs are implemented in PyTorch[16] and MONAI on a Ubuntu desktop with an NVIDIA RTX 2080 Ti GPU. During inference, we first resample the ROI of the testing image to isotropic voxel size of $1 \times 1 \times 1$ mm^3 and obtain the prediction on this resampled image. The actual coordinates of the target are computed based on the resampling factor.

3. EXPERIMENTS AND RESULTS

3.1 Experimental Design

In our experiments, we compare the localization results obtained by each method against the ground truth annotated by the more experienced neurosurgeon. The labels annotated by the same neurosurgeon but at a different time and the other neurosurgeon are used to compute the intra- and inter-rater variability respectively, which can be viewed as the lower bounds of the localization error for the automatic localization algorithms. For quantitative evaluation, we use the mean radial error (MRE) to measure the Euclidean distance between prediction and ground truth as well as the successful detection rate (SDR) for three radii (2 mm, 4 mm and 6 mm). We also compare the inference time per testing image and the impact of the amount of training data for each method.

---

*https://github.com/anibali/dsntnn

3.2 Results

In Table 1, we show the MRE and SDR achieved by different methods with 100%, 50% and 25% of the training data to study the effect of the training set size on the results. The intra-rater and inter-rater variability of MRE are 2.04 ± 0.87 mm and 2.42 ± 1.17 mm, respectively. Our results show that when trained with 100% of the training data, both the DL-based methods, i.e., HM and DSNT, and the registration method SyN have achieved MRE comparable to the inter/intra-rater variability. We perform the paired t-test with a p-value of 0.05 and find no statistically significant difference between each method. In terms of SDR, HM leads to accurate target localization, i.e., targets within 2 mm, more often than SyN and DSNT. We also observe that the SDR achieved by HM and registration are comparable to the inter-rater performance. On the other hand, when trained with less training data, e.g., with only 25% of training data, we observe that DSNT can yield better localization performance than HM as it leads to significantly lower mean MRE with lower standard deviation. Lastly, in terms of inference time, we show that the DL-based methods can reduce the inference time by orders of magnitude compared to the registration method.

Table 1. Comparison of localization performances achieved by different methods

METHODS	MRE($\downarrow$) (mm)			SDR($\uparrow$) (%)			Time (sec)
	100%	50%	25%	100%	50%	25%	
Intra-rater	2.04±0.87	-	-	58, 99, 100	-	-	-
Inter-rater	2.42±1.17	-	-	45, 93, 100	-	-	-
SyN	2.49±1.09	-	-	40, 91, 100	-	-	3.7
HM	2.34±1.12	4.38±3.56	7.13±10.17	42, 92, 100	5, 53, 91	1, 43, 85	0.005
DSNT	2.45±1.03	4.23±1.98	4.17±1.46	29, 95, 100	3, 50, 95	2, 47, 87	0.005

4. BREAKTHROUGH WORK

We have compared several localization methods for ANT-DBS targeting and assessed the performances against the annotations from multi-raters. Our experimental result show that the DL-based localization models can lead to an accuracy that is comparable to the measured intra/inter-rater variability. These models can be trained with pseudo labels that are generated by only one expert-annotated label and can make inference much faster than traditional registration methods. To the best of our knowledge, this work is the first to investigate various localization algorithms for ANT-DBS with multi-rater annotations.

5. DISCUSSION AND CONCLUSION

An interesting observation from our experiment is that the DSNT-layer can improve localization performance when the model is trained with a small amount of data. We conjecture that the outperformance of DSNT might be due to the MSE loss term computed on the 3D coordinates. In our second stage, the ROIs are adjusted to have consistent image contents and thus the 3D coordinates of pseudo label targets do not vary much. With the additional supervision from the coordinate loss, the model equipped with a DSNT layer is more likely to converge faster than the HM model when the amount of training data is limited. With a larger amount of training data, HM is shown to be a better localization algorithm than DSNT as it produces more high-quality targets. In summary, in this work, we perform a multi-rater comparative study for ANT-DBS target localization. Our experimental results demonstrate that DL-based models that are trained with **pseudo labels** can achieve a localization accuracy that is comparable to the intra/inter-rater variability of experienced neurosurgeons with much faster inference time than traditional method. A validation study in which surgeons are asked to assess the quality of the localization, i.e., whether it could be used clinically is ongoing.

6. ACKNOWLEDGMENTS

REFERENCES

[1] Johnson, Emily L. "Seizures and epilepsy." Medical Clinics 103, no. 2 (2019): 309-324.

[2] Van Der Vlis, Tim AM Bouwens, Olaf EMG Schijns, Frédéric LWVJ Schaper, Govert Hoogland, Pieter Kubben, Louis Wagner, Rob Rouhl, Yasin Temel, and Linda Ackermans. "Deep brain stimulation of the anterior nucleus of the thalamus for drug-resistant epilepsy." Neurosurgical review 42, no. 2 (2019): 287-296.

[3] Li, Michael CH, and Mark J. Cook. "Deep brain stimulation for drug-resistant epilepsy." Epilepsia 59, no. 2 (2018): 273-290.

[4] Järvenpää, Soila, Kai Lehtimäki, Sirpa Rainesalo, Timo Möttönen, and Jukka Peltola. "Improving the effectiveness of ANT DBS therapy for epilepsy with optimal current targeting." Epilepsia open 5, no. 3 (2020): 406-417.

[5] Cui, Can, Han Liu, Dario J. Englot, and Benoit M. Dawant. "Brain vessel segmentation in contrast-enhanced T1-weighted MR Images for deep brain stimulation of the anterior thalamus using a deep convolutional neural network." In Medical Imaging 2021: Image-Guided Procedures, Robotic Interventions, and Modeling, vol. 11598, p. 115980K. International Society for Optics and Photonics, 2021.

[6] Yushkevich, Paul A., Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C. Gee, and Guido Gerig. "User-guided 3D active contour segmentation of anatomical structures: significantly improved efficiency and reliability." Neuroimage 31, no. 3 (2006): 1116-1128.

[7] Wang, Hongzhi, Alison Pouch, Manabu Takabe, Benjamin Jackson, Joseph Gorman, Robert Gorman, and Paul A. Yushkevich. "Multi-atlas segmentation with robust label transfer and label fusion." In International Conference on Information Processing in Medical Imaging, pp. 548-559. Springer, Berlin, Heidelberg, 2013.

[8] Wang, Jianing, Yuan Liu, Jack H. Noble, and Benoit M. Dawant. "Automatic selection of landmarks in T1-weighted head MRI with regression forests for image registration initialization." Journal of Medical Imaging 4, no. 4 (2017): 044005.

[9] Avants, Brian B., Charles L. Epstein, Murray Grossman, and James C. Gee. "Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain." Medical image analysis 12, no. 1 (2008): 26-41.

[10] Avants, Brian B., Nicholas J. Tustison, Gang Song, Philip A. Cook, Arno Klein, and James C. Gee. "A reproducible evaluation of ANTs similarity metric performance in brain image registration." Neuroimage 54, no. 3 (2011): 2033-2044.

[11] Keller, Simon S., Mark P. Richardson, Jan-Christoph Schoene-Bake, Jonathan O'muircheartaigh, Samia Elkommos, Barbara Kreilkamp, Yee Yen Goh, Anthony G. Marson, Christian Elger, and Bernd Weber. "Thalamotemporal alteration and postoperative seizures in temporal lobe epilepsy." Annals of neurology 77, no. 5 (2015): 760-774.

[12] Liu, Han, Can Cui, Dario J. Englot, and Benoit M. Dawant. "Uncertainty Estimation in Medical Image Localization: Towards Robust Anterior Thalamus Targeting for Deep Brain Stimulation." In Interpretable and Annotation-Efficient Learning for Medical Image Computing, pp. 130-137. Springer, Cham, 2020.

[13] Nibali, Aiden, Zhen He, Stuart Morgan, and Luke Prendergast. "Numerical coordinate regression with convolutional neural networks." arXiv preprint arXiv:1801.07372 (2018).

[14] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." In International Conference on Medical image computing and computer-assisted intervention, pp. 234-241. Springer, Cham, 2015.

[15] Kingma, Diederik P., and Jimmy Ba. "Adam: A method for stochastic optimization." arXiv preprint arXiv:1412.6980 (2014).

[16] Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen et al. "Pytorch: An imperative style, high-performance deep learning library." Advances in neural information processing systems 32 (2019): 8026-8037.