On Exploiting Connectomics for Thalamic Nuclei Localization: A Supervised Learning Approach

Dimitris Bolis DBOLIS@EE.ETHZ.CH
András Jakab JAKAB@VISION.EE.ETHZ.CH
Orçun Göksel OGOKSEL@VISION.EE.ETHZ.CH
Gábor Székely SZEKELY@VISION.EE.ETHZ.CH
Computer Vision Laboratory, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland

Abstract

We hypothesize that thalamocortical connections after mapping by diffusion tensor imaging can serve as surrogate markers of individual anatomy, which can then be used for localizing specific neurosurgical targets in the thalamus. A variety of learning schemes, together with pre- and post- processing steps are studied for our goal of thalamic nuclei localization. The training procedure is performed after non-linear registration and probabilistic tractography on diffusion magnetic resonance imaging data. Our results indicate that thalamocortical connectivity data do contain sufficient discriminant internucleus information for thalamic nuclei localization.

1. Introduction

The human thalamus constitutes a subcortical brain structure, consisted principally of gray matter. It serves both for information preprocessing and relaying. In particular, every sensory system, excluding the olfactory one, sends information firstly to the thalamus, which retransmits it to selective cerebral cortical loci, via specialized thalamocortical connections.

The thalamus contains both functionally and anatomically distinct subparts, the thalamic nuclei, which due to their crucial role in the brain function represent an important target for the understanding and treatment of numerous neurological conditions. Specifically, the Parkinson's Disease, is commonly treated either by thalamotomy or thalamus Deep Brain Stimulation (DBS). In the first case, part of the thalamus is destroyed, while in the latter the target is retuned or inactivated, less invasively, by regular electrical stimulation. In both cases, the thalamic target localization is of immense importance for ensuring both the efficacy of the method and the reduction of complications.

During image-guided neurosurgical interventions, localization can be done either indirectly, using atlas registration-based techniques (Otsuki *et al.*, 1994), or directly with patient-specific methods (Wiegell *et al*., 2003), (Deoni *et al*., 2006). The former are often inaccurate due to intersubject variability and potential deformations, while the latter, being usually based on local imaging properties, are limited by the lower thalamus image contrast offered by the current imaging technologies (T1, T2 weighted Magnetic Resonance Imaging (MRI) and Diffusion MRI).

In white matter, diffusion based visualization gains higher contrast due to the higher anisotropic properties of neural fibers. Additionally, the thalamocortical neural network structure is anatomically related to the thalamus subdivision in nuclei. Based on these facts, Behrens *et al.* suggested thalamic nuclei segmentation using thalamocortical connectivity information (2003). Specifically, performing probabilistic tractography on diffusion MRI data, they calculated relative strengths for the thalamocortical projections. A "winner takes all" data reduction scheme was then followed, in which the voxels inside the thalamus were clustered based on their strongest projection on cortical and subcortical targets.

In this paper, a supervised learning approach is proposed for exploiting the thalamocortical connectivity information for thalamic localization. Three standard learning schemes, namely the Support Vector Machines (SVMs), Boosting, and Random Forest (RF), are examined here. Various pre- and post-processing schemes are also considered for increased speed and improved accuracy. Our main goal is to provide individual, patient-specific target maps by effectively incorporating information from thalamocortical connection patterns. We hypothesize that machine learning algorithms are capable in relating connectivity-based information with atlas based segmentation of nuclei.

* Appearing in Proceedings of the 29 the International Conference on Machine Learning, Edinburgh, Scotland, UK, 2012.

2. Methods and Materials

2.1 Construction of the Thalamus Atlas

We exploit recent results of the fusion of multiple histological atlases to generate a three-dimensional statistical model of the human thalamus (Krauth *et al*., 2010), **Fig.1**. The corresponding data from this 3D atlas was non-linearly matched with the MR-based visible thalamic borders of the MNI152 T1-weighted template image. The voxelized 3D volume of ten selected nuclei was then taken as ground truth for training data classifiers and validating the results.

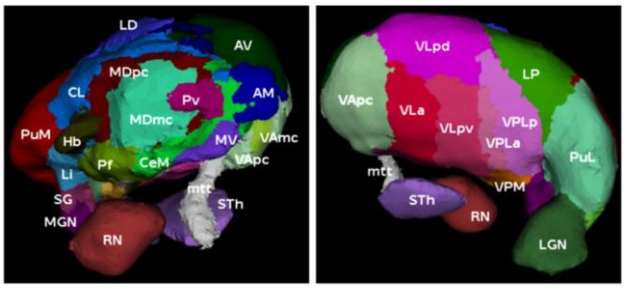

Fig. 1: 3D representation of the mean thalamus atlas used in the study. References: Morel (2007) Stereotactic Atlas of the Human Thalamus and Basal Ganglia. New York: Informa Healthcare USA, Inc.

2.2 Diffusion Tensor Imaging (DTI) and Probabilistic Tractography

For 40 healthy volunteers, diffusion tensor imaging was performed with 25 diffusion weighting directions. Probabilistic fiber tracing was performed similarly to studies describing the connectivity-based parcellation of the thalamus (Behrens *et al*., 2003). Connections between the thalamus and 52 cortical and subcortical areas were mapped; the latter defined using the Harvard-Oxford Cortical Atlas. Such connectivity maps were used as training features for the classifiers.

2.3 Classification

Assuming each thalamic voxel as a training sample and each registered thalamus occupying 9326 voxels in the standard space defines our data as 9326*40 samples of 52 features each. The size of each nucleus is about 1-11% of the total thalamus size.

Three well-known classification algorithms were evaluated for the given task. First, the SVM algorithm (Chang and Lin, 2011), which aims to maximize a functional margin in a projected higher dimensional data space, was examined. Following comparisons of different kernels for SVM, the best performing RBF-kernel is presented here. Second, boosting based methods were considered (Schapire, 1990). This category of methods train multiple simple learners (weak classifiers) by successively assigning higher weights to misclassified instances. The final hypothesis is then given by a weighted majority vote of all the weak hypotheses calculated. Decision trees were used here with 500 ensemble learning cycles. The results are given here for the best performing Gentle Boost implementation. Finally, the Random Forests algorithm was tested (Breiman, 2001), which uses a combination of tree predictors that are inputted with a sub-sampled training set. The sub-sampling is done with replacement. Then, at each node a small fraction of the input variables is used for its splitting. The decision is given by the votes of the individual trees. An implementation with 200 trees is presented here.

In the case of SVM, the normalization of data with a logarithmic transform as a preprocessing step resulted the SVM to converge significantly faster. Additionally, a postprocessing step of 3D median filtering was performed in order to smooth the classifiers' outputs, which were then evaluated in a 5-fold cross-validation scheme.

2.4 Methods Outline

The overall steps of the study are demonstrated in **Fig. 2**.

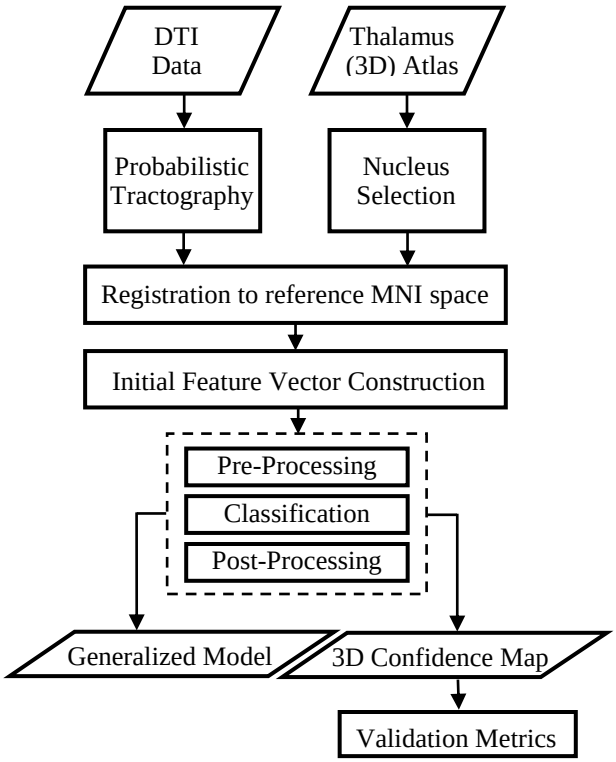

Fig. 2: Predicting the location of thalamic nuclei based on DTI probabilistic mapping of thalamocortical connections: flowchart of the procedure.

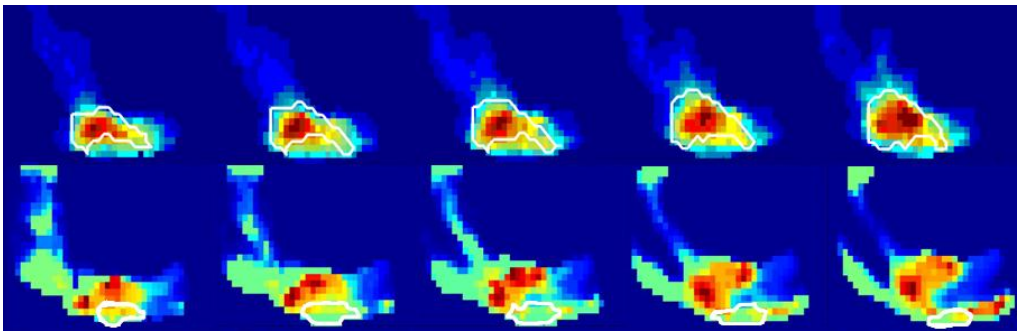

Fig. 3: Confidence Maps. Upper row: Successful case (MDpc nucleus). Lower row: Poorly performing case (MDmc nucleus). White border: ground truth by the Thalamus Atlas.

3. Results

3.1 Confidence Maps

Having applied the classifier combined with the pre- and post-processing steps, confidence maps were constructed, visualizing the classifier outputs. In **Fig. 3**, such confidence maps are presented for a successful (high Dice's Coefficient (DC)) and a poorly-performing case (low DC), respectively. For these examples, a normalized dataset was used for training the Random Forest classifier, the decision values of which were then filtered by a 3D median filter. For a quantitative assessment of such maps, the highest confidence voxels (5% each nucleus size) were selected, and the percentage of those falling inside the target was calculated (**Fig. 4**).

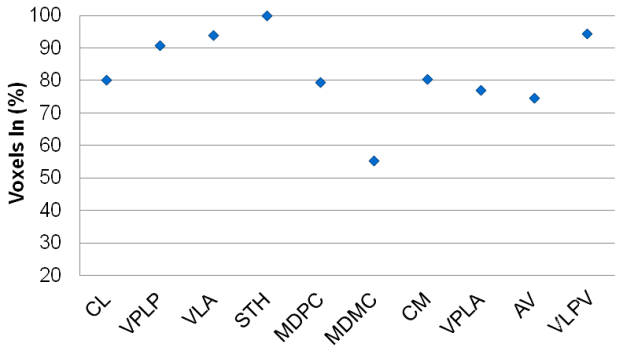

Fig. 4: Success rate for the top 5% confident voxels for ten selected nuclei using the post-processed Random Forest's decision values.

3.2 Comparison on the VLPv nucleus

In **Fig. 5**, the performance of SVM, Boosting and Random Forests is compared for the case of the VLPv nucleus, which is a common neurosurgical target for the treatment of the Parkinson's Disease, both before and after postprocessing (3D median filtering). It is shown that after post-processing the performance generally increases and that in any case the Random Forest performs better than the other two learning schemes.

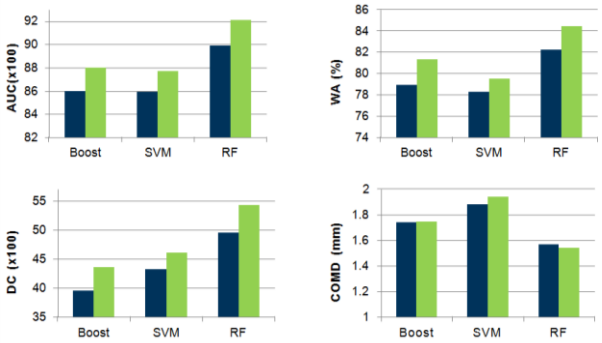

Fig. 5: Comparison of SVM, Boosting (Boost) and Random Forest (RF) for both the original (blue) and post-processed (green) decision values, for the VLPv nucleus. From top left to bottom right: Area Under Curve (AUC) (x100), Weighted Accuracy (WA) (%), Dice's Coefficient (DC) (x100), Center of Mass Deviation (COMD) (mm).

3.3 Comparison on Ten Selected Thalamic Nuclei

In **Fig. 6**, the previous section's two best approaches both in terms of speed and performance, Boosting and Random Forests are compared for ten selected nuclei. The latter performs better in all the cases.

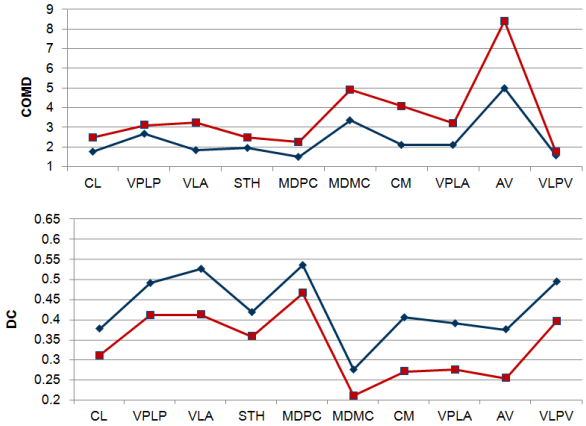

Fig. 6: Comparison of Boosting (red) and Random Forests (blue). Top: Center of Mass Deviation (COMD) (mm). Bottom: Dice's Coefficient (DC).

3.4 Most Effective Approach

The results with Random Forest, the most effective classifier for our problem, are presented in **Fig. 7**.

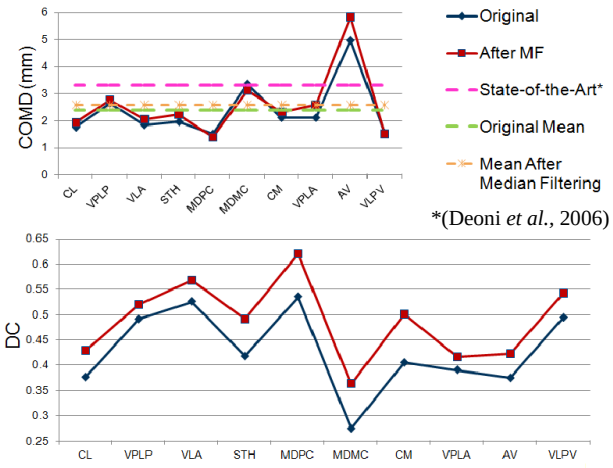

Fig. 7: Random Forests: Before (blue) and after Median Filtering (red). Top: Center of Mass Deviation (COMD) (mm), to Bottom: Dice's Coefficient (DC).

4. Discussion

Confidence maps, such as the ones presented in **Fig. 3**, could be used by clinicians as a complementary visual aid for thalamic nucleus localization. We showed that *confident* voxels fall inside the target with high frequency. For all nuclei excluding MDmc, more than 74.5% of *confident* voxels were classified correctly (**Fig. 4**).

In a first stage, by studying SVM, Boosting and Random Forests for the VLPv nucleus (**Fig. 5**), we concluded that Random Forests generate most accurate results compared to the other two. Additionally, SVM requires significantly more computational time than Boosting and Random Forests, which ran in our tests on the order of hours (considering both training and testing). In a second stage, Random Forests was tested against Boosting for all 10 nuclei (**Fig. 6**). It is seen that the former performs superiorly. Finally, the proposed approach of Random Forests was comparatively presented before and after 3D median filtering (**Fig. 7**). In all the cases, median filtering improved the results. The mean Center of Mass Deviation (COMD) calculated with our approach either before (2.4mm) or after filtering (2.6mm) is seen to be smaller than the previously reported values of 3.3mm-4.4mm, of the closest work (Deoni *et al.*, 2006). Note that our reported Dice coefficients are relatively low, showing that even if the presented method is good for localizing the nuclei, it does not carry sufficient information in order to perfectly delineate them from their surroundings. Different inaccuracies can also be attributed to the lack of *absolute* ground truth, for which we use a registered atlas.

5. Conclusions

We have shown that thalamocortical connectivity does contain internucleus discriminant information. Such information has been exploited for producing confidence maps and localizing individual thalamic nuclei. The nuclei center of mass are localized more accurately using our method compared to the state-of-the-art in the literature. Supervised learning approaches rely greatly on the availability of an absolute ground truth, the lack of which is a common problem. Our research efforts are now channeling into boosting our approach's efficiency by enriching the feature vector with additional information sources (e.g. local imaging and spatial information) and developing an absolute ground truth based on post-mortem histology preparation and diffusion imaging of the same specimen.

Acknowledgments

This work is supported by the NCCR Co-Me of the Swiss National Science Foundation.

References

Schapire RE (1990), The Strength of Weak Learnability. Machine Learning. 5(2): 197-227.

Otsuki T, Jokura H, Takahashi K, *et al.* (1994) Stereotactic Thalamotomy with a Computerized Brain Atlas: Technical Case Report. Neurosurgery 35(4): 764-768.

Breiman L (2001) Random Forests. Machine Learning. 1(45): 5-32.

Behrens TE, Johansen-Berg H, Woolrich MW, *et al.* (2003) Non-invasive Mapping of Connections Between Human Thalamus and Cortex Using Diffusion Imaging, Nature Neuroscience, 6(7): 750–757.

Wiegell MR, Tuch DS, Larson HWB, and Wedeen VJ (2003) Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging. NeuroImage, 19(2): 391–402.

Deoni SCL, Rutt BK, Parrent AG, and Peters TM (2006) Segmentation of thalamic nuclei using a modified k-means clustering algorithm and high-resolution quantitative magnetic resonance imaging at 1.5 T. NeuroImage, 4(1): 117-126.

Krauth A, Blanc R, Poveda A, Jeanmonod D, Morel A, Székely G (2010) A mean three-dimensional atlas of the human thalamus: Generation from multiple histological data. NeuroImage, 49(3): 2053–2062.

Chang CC, Lin CJ (2011) LIBSVM: a library for support vector machines. ACM - Transactions on Intelligent Systems and Technology, 2(3): 1-27.