Tractogram filtering of anatomically non-plausible fibers with geometric deep learning

Pietro Astolfi^{1,2,3*}, Ruben Verhagen², Laurent Petit⁴, Emanuele Olivetti^{1,2},
Jonathan Masci⁶, Davide Boscaini⁵ and Paolo Avesani^{1,2}

¹ NeuroInformatics Laboratory (NILab), Bruno Kessler Foundation, Trento, Italy

² Center for Mind and Brain Sciences (CIMeC), University of Trento, Italy

³ PAVIS, Italian Institute of Technology, Genova, Italy

⁴ GIN, IMN, CNRS, CEA, Universit de Bordeaux, Bordeaux, France

⁵ Technologies of Vision (TeV), Bruno Kessler Foundation, Trento, Italy

⁶ NNAISENSE, Lugano, Switzerland

pastolfi@fbk.eu avesani@fbk.eu

Abstract. Tractograms are virtual representations of the white matter fibers of the brain. They are of primary interest for tasks like presurgical planning, and investigation of neuroplasticity or brain disorders. Each tractogram is composed of millions of fibers encoded as 3D polylines. Unfortunately, a large portion of those fibers are not anatomically plausible and can be considered artifacts of the tracking algorithms. Common methods for tractogram filtering are based on signal reconstruction, a principled approach, but unable to consider the knowledge of brain anatomy. In this work, we address the problem of tractogram filtering as a supervised learning problem by exploiting the ground truth annotations obtained with a recent heuristic method, which labels fibers as either anatomically plausible or non-plausible according to well-established anatomical properties. The intuitive idea is to model a fiber as a point cloud and the goal is to investigate whether and how a geometric deep learning model might capture its anatomical properties. Our contribution is an extension of the Dynamic Edge Convolution model that exploits the sequential relations of points in a fiber and discriminates with high accuracy plausible/non-plausible fibers.

1 Introduction

Tractography represents a powerful method to reconstruct the white matter fibers from diffusion magnetic resonance (MR) recordings [1]. While this method provides a good approximation of the brain connectivity structure, there is an open issue of artifactual fibers [12,11], i.e. anatomically non-plausible pathways. We address the task of filtering out such artifactual fibers using a deep learning model.

* We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

The structural connectivity of the brain can be reconstructed from diffusion MR signal by a step of local estimation of the diffusivity model [21] and a subsequent step of fiber tracking [6]. The outcome is a tractogram, a virtual representation of the axonal pathways in the white matter, where each fiber is encoded as 3D polylines, commonly referred to as *streamlines*. Typically, a whole brain tractogram is composed of millions of streamlines.

Tractograms are providing valuable contributions to critical tasks like presurgical intervention planning, the detection of biomarkers for brain disorders and the investigation of neuroplasticity. For these purposes, the accuracy of fiber tracking is of paramount importance. While in the preprocessing of diffusion data it is common practice to denoise the signal before estimating the diffusivity model [7,25], there is no similar step after fiber tracking to filter out noisy streamlines.

The assessment of the accuracy of fiber tracking has been approached a few years ago with an open contest⁷ involving many (20) research groups. The contest was designed as a task of bundle detection on a dataset composed of simulated diffusion MRI data. The joint effort allowed the evaluation of the quality and limits of the most common tracking methods [12]. The positive outcome has been the lack of false negative streamlines, the critical issue has concerned the many false positive errors, e.g. artifactual streamlines.

The occurrence of false positive streamlines is not surprising. The general strategy of tracking methods is to oversample the possible pathways to preserve the property of coverage of all true positive fibers [19]. The tacit assumption is to postpone the task of filtering false positive streamlines to a subsequent post-processing step. The reason for this strategy is the difficulty of encoding anatomical priors into tracking algorithms.

The most common approach to reduce artifactual streamlines are methods based on the inverse problem of signal reconstruction, e.g. Life [16], SIFT [20], Commit [3]. The intuitive idea is to estimate how much the orientation of a streamline explains the diffusion signal. Despite the principled criterion, these methods do not take into account the knowledge of brain anatomy, like another method based on the topological properties of streamlines [24].

A recent rule-based method, namely ExTractor [10], has been proposed to filter out artifactual streamlines from tractograms by following anatomical principles. The rules that it proposes encode the geometrical and spatial properties of the streamlines with respect to the basics of white matter neuroanatomy in terms of association, projection, and commissural fibers. As output, ExTractor labels streamlines as *anatomically plausible* or *anatomically non-plausible*.

In this work, we address for the first time the problem of tractogram filtering as a supervised learning problem. We need to train a binary classifier to discriminate between two classes: anatomically plausible (P) and non-plausible (nP) fibers. The intuitive idea is to exploit the labeling of fibers provided by the rule-based method [10] and to adopt deep learning models to learn the features of streamlines underlying the rules. The ultimate goal is to have a fast run-time solution to filter large tractograms and a flexible method to transpose new expert annotations.

⁷ tractometer.org/ismrm_2015_challenge

Given the sequential structure of a streamline, we have chosen as reference learning model a bidirectional LSTM (bLSTM) neural network [9]. Although this model can exploit the sequential information, it requires forcing the streamline representation to a fixed number of points. For this reason, we consider Geometric Deep Learning (GDL) models [13] that support more flexible and appropriate data representations. We investigate a PointNet (PN) model [18], where a streamline can be represented as a point cloud i.e. set of 3D points, and a Dynamic Edge Convolution (DEC) model [23], which in addition to PN considers the relations between points belonging to the same spherical local context. Our experiments shows that both GDL models provide an improvement with respect to bLSTM.

Despite the better results provided by PN and DEC, these models are invariant to permutations of points in the input point cloud. It means that if we permute the points in a streamline classified as plausible, the model will continue to classify it as plausible albeit the altered sequence of points. To overcome this issue we propose a Sequence Dynamic Edge Convolution (sDEC) model, an extension of the DEC model that introduces the property of being sequence sensitive.

While the sDEC model provides only a modest increase in accuracy with respect to PN in classifying P and nP fibers, the analysis of error distribution shows different behaviours. sDEC is more robust when fibers are long and curved: those are the type of fibers where PN performs worst and produces significantly more false positive errors. In addition, visual inspection of the false negatives errors made by sDEC shows that these fibers look truly anatomically non-plausible. We may argue that this apparent mismatch might be related to the noise of the labelling process.

2 Materials

Dataset. Our reference dataset is composed of a collection of tractograms of 20 subjects randomly selected from the HCP dataset [22], which is a publicly available curated MRI dataset. The tractograms were obtained computing (i) the constrained spherical deconvolution (CSD) model [21] on the 3T DWI (1.25mm, 270 multishell gradients), and the (ii) Particle Filtering Anatomically Constrained Tractography (PF-ACT) algorithm [6]. Specifically, the tracking generated around ~ 1 M fibers for each tractogram by seeding 16 points for each voxel with step size 0.5mm. The tractogram of each subject has been non linearly coregistered to the MNI standard space, and, for computational purposes, all the fibers have been compressed to the most significant points [17].

ExTractor labelling. According to the premise of a supervised learning approach, we created a dataset where for each tractogram we labeled the fiber as either plausible (P) or non-plausible (nP). The procedure of labelling followed the heuristic rules defined by ExTractor [10], a tool that encodes the current knowledge on the anatomical pathways of white matter structures. In particular, ExTractor carries out a 2-step procedure. First, fibers are marked as nP when they are either (i) shorter than 20 mm, (ii) contain a loop, or (iii) are truncated, i.e.,

Fig. 1: The DEC architecture adopted. Green, gray, and red blocks represent input, intermediate, and output tensors, respectively. Parametric layers are colored in blue, while fixed layers in white.

they terminate in the deep white matter. The second step is concerned with the labeling of fibers marked as P, by splitting the main pathways into the three macro classes: association, projection, and commissural. This selection of fibers is further refined to filter out non-plausible pathways using a clustering method [2]. The outcome of the ExTractor procedure is a balanced partition between P and nP fibers, irrespective of the different tracking algorithms and data sources [10]. On our dataset, Extractor resulted in $49.7 \pm 1.5\%$ of P streamlines on average on the 20 subjects, but more significantly it showed the presence of a massive percentage, $31.8 \pm 1.2\%$, of nP streamlines shorter than 20mm.

3 Methods

The intuition underlying our work is to treat a streamline as an undirected sequence of 3D points, aiming to learn geometric and spatial features relevant for the tractography filtering task. Currently, the best way to achieve such a goal is by employing GDL models [13], which are designed to learn geometric features of graphs and point clouds. Among the existing GDL approaches, PointNet [18] is the most adopted method both for its simplicity and effectiveness. Nevertheless, such an approach does not consider relations between points. It learns global properties of a point cloud just by encoding all the points separately and then aggregating them in a single descriptor through global pooling. Conversely, Dynamic Edge Convolution (DEC) [23] considers the points relations, by encoding them as edges of a graph dynamically induced by the point cloud.

Dynamic Edge Convolution. Considering a point cloud $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathbb{R}^3$, the DEC model first guesses an initial graph structure for the point cloud by retrieving for each point $\mathbf{x}_i$ the set of k nearest neighbors, $knn(\mathbf{x}_i) = \{\mathbf{x}_{j_{i_1}}, \dots, \mathbf{x}_{j_{i_k}}\}$, in terms of Euclidean distance (see Figure 2a). Hence, DEC builds a k -nn graph, $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{X}$ is the set of nodes, and an edge $(i, j) \in \mathcal{E}$ exists iff $\mathbf{x}_j \in knn(\mathbf{x}_i)$. Then, each point representation, $\mathbf{x}_i$, is enriched with the representation of each of its neighbors $\mathbf{x}_{j_i}$, creating edges features $\mathbf{e}_{ij}$, which are learnt through a neural network h_{Θ} , i.e. $\mathbf{e}_{ij} = h_{\Theta}(\mathbf{x}_i \oplus (\mathbf{x}_j - \mathbf{x}_i))$, where $\oplus$ denotes the concatenation operator. Finally, a new representation of a point, $\mathbf{x}'_i$,

Fig. 2: Comparison between Euclidean k -nn (a), and k -nn based on the sequence graph (b).

is obtained by aggregating all the learned edge features with a pooling operator, i.e. $\mathbf{x}'_i = \text{pool}(\mathbf{e}_{ij})$, $j: (i, j) \in \mathcal{E}$, where pool is either max or mean. The sequence of operations that from $\mathcal{X}$ produce $\mathcal{X}'$ define an Edge Convolution (EC) layer, which in the DEC model is repeated multiple times. Referring to our architecture, see Figure 1, two EC layers are stacked in depth to produce new representations $\mathcal{X}'$ and $\mathcal{X}''$ with 64 and 128 features respectively. Each of the two EC layers computes its own k -nn graph in order to adjust the local neighborhood of points to its input representation, i.e. $\mathcal{X}$ and $\mathcal{X}'$. The re-computation of knn is what defines the DEC model as *dynamic*. Finally, the different learned representations are concatenated, encoded with a learning layer, g_{Φ} , to 1024 features, and pooled to obtain a single descriptor of the point cloud, $\mathcal{Z} = \text{pool}(g_{\Phi}(\mathcal{X}' \oplus \mathcal{X}''))$, which is classified using a fully connected (FC) network.

Dynamic Edge Convolution of a Sequence. A remarkable property of the DEC model (also shared by PN) is the invariance to the permutation of the points in the input point cloud. Indeed, such models make use only of operators invariant to the order e.g., FC layers, max / mean pooling layers. Although this property is fundamental in the point cloud domain, it becomes undesired if the input is a sequence as in our case. To solve this issue, we propose a simple but well-motivated modification of the DEC model, in which we impose the input point cloud to have a graph structure, without needing a Euclidean k -nn to induce it. According to the streamline structure, we impose the input to be a bidirectional sequence graph where each non-terminal point, $\mathbf{x}_{i \neq 0, n}$, has two neighbors: the previous and the next point in the sequence, while the terminal points, $\mathbf{x}_0, \mathbf{x}_n$, have just one neighbor (see Figure 2b):

$$\mathcal{G}(\mathcal{V}, \mathcal{E}'), e'_{ij} \in \mathcal{E}': \mathbf{x}_i \rightarrow \mathbf{x}_j, j = i + 1 \vee j = i - 1.$$

It is important to remark that we impose this structure only in the first sDEC layer, which is enough to lose the invariance with respect to the input permutations while maintaining the invariance with respect to the flipping of the sequence, thanks to the bidirectionality of the sequence graph. Also, imposing this structure only in the first layer preserves the dynamicity of the model, which remains able to re-organize the structure of the point cloud in the latent space according to the task at hand.

Method	Accuracy	Precision	Recall	DSC
bLSTM	92.8 (± 0.3)	93.7 (± 0.5)	96.1 (± 0.2)	94.9 (± 0.3)
PN	94.5 (± 0.1)	95.4 (± 0.2)	96.8 (± 0.2)	96.1 (± 0.2)
DEC	94.3 (± 0.1)	95.4 (± 0.3)	96.5 (± 0.2)	95.9 (± 0.2)
sDEC	95.2 (± 0.1)	96.2 (± 0.3)	96.9 (± 0.1)	96.6 (± 0.2)

Table 1: Average scores for the 4 HCP subjects of test set. Standard deviation among the 4 subjects is reported between brackets.

4 Experiments and Results

The main purpose of the experiments is to perform an empirical assessment of the behaviour of different models for the task of tractogram filtering. In our comparison, we consider a reference model, bLSTM, and three competing models based on geometric deep learning: PN, DEC, and sDEC.

Before setting the training of learning models, we carried out a few preprocessing steps. The first step was the pruning of fibers. An analysis of the distribution of fibers between P and nP according to the ExTractors rules, highlighted a massive percentage of very short streamlines, i.e. length below 20 mm. All of them were labeled as nP. To reduce this potential bias for the learning models we removed such streamlines from the tractograms. After the pruning, the average distribution of P and nP classes was 68% and 32% respectively. The second preprocessing step concerned the resampling of the points of streamlines. Traditional learning models, e.g. bLSTM, require input to be represented as fixed vectorial representations. Therefore, despite GDL models can deal with a varying number of points, we have to resample all streamlines to have the same number of points [5,14,8]. We computed the resampling using a cubic B-spline interpolation and empirically compared different representations based on 12, 16 and 20 points. No meaningful performance differences were noticed. In Table 1 we report the results for streamlines resampled to 16 points.

To perform a fair comparison of the four models, we have configured all of them to have a similar number of parameters, approximately around 900K. While the architecture of DEC and sDEC models is reported in Figure 1, the configuration of the remaining models is:

bLSTM: $\text{MLP}(128) \rightarrow \text{LSTM}(256) \oplus \text{LSTM}^{-1}(256) \rightarrow \text{MLP}(256, 128) \rightarrow \text{FC}(2)$;
 PN: $\text{MLP}(64, 64, 64, 128, 1024) \rightarrow \text{MAX} \rightarrow \text{MLP}(512, 256, 40) \rightarrow \text{FC}(2)$,
 where FC is a Fully Connected layer, and MLP contains sequences of (FC, BatchNorm, ReLU). For the implementation of the models, we adopted the PyTorch library [15] with the PyTorch Geometric extension [4].

For all the experiments we considered a dataset of 20 tractograms from 20 different subjects. We partitioned the data into three sets: 12 tractograms for training the models, 4 tractograms for hyperparameters tuning and model selection, and 4 tractograms for testing. The training was designed as follows: 1000 epochs with evaluation step every 20 epochs; cross-entropy loss; Adam optimizer with default alfa and beta momentum (0.9, 0.99); initial learning rate

Fig. 3: Analysis of the error distribution for PN and sDEC models with respect to streamlines length and curvature: (A) accuracy, (B) false positive rate. (C) Sample of false negative streamlines, mistakenly classified as anatomically non-plausible.

of 10^{-3} multiplied by a factor of 0.7 every 90 epochs until a minimum value of $5 \cdot 10^{-5}$ is reached. In the training procedure, we used mini-batches composed of 16K streamlines randomly sampled from two subjects, 8K from each of them. A subject is sampled only once for each epoch. On the test set, we computed the following evaluation measures: accuracy, precision, recall, and DSC. In Table 1 the values refer to the average for each subject. Notice that the classification task is single-streamline, therefore the size of our test set is composed of ~ 3 million streamlines.

An additional experiment was designed to test the invariance of DEC and sensitivity of sDEC with respect to the order of points in a streamline. To this end, we tested the two models on a version of the test set where the sequence of points of streamlines were randomly permuted. In agreement with the working hypothesis, the accuracy for DEC remained the same, 94.3%, while there was a drop for sDEC to 30.0%.

A further analysis concerned the examination of the error distribution. We restricted the analysis to PN and sDEC only. Our interest was to characterize whether and how the error differs between the two models. We considered a couple of features in the analysis: length and curvature of streamlines. All streamlines were partitioned into three categories with respect to length (short $[0, 50]$ mm, medium $[50, 100]$ mm, long $[100, 300]$ mm) and curvature (straight $[0.0, 0.05]$, curved $[0.05, 0.10]$, very curved $[0.10, 0.20]$), where each of them contained at least a portion of 15% of streamlines. Combining the partitions we obtained 9 categories. In Figure 3A we depict how the accuracy score varies across the categories. The task becomes more difficult when both length and curvature are greater. A similar analysis was carried out focusing on the distribution of false positive error, i.e. streamlines mistakenly classified as anatomically plausible. The results are reported in Figure 3B. In this case, we observe a major difference between the two models in the most critical categories, namely long and curved

streamlines. In Figure 3C we illustrate a qualitative example of false negative errors, when streamlines should be recognized as anatomically plausible while they are classified as anatomically non-plausible. Visual inspection by an expert anatomist confirms that the classifier was indeed correct in those cases and that the labelling process was noisy.

5 Discussion and Conclusions

The results reported in Table 1 confirm that we may successfully approach the task of tractogram filtering as a supervised learning problem. The best accuracy achieved by sDEC is beyond 95% with low standard deviation across subjects. This performance is obtained considering only the representation of streamlines as a sequence of 3D points. The runtime application of sDEC model is very fast. We may filter a whole tractogram with $\sim 1M$ streamlines in less than one minute (46.2 sec using a gpu NVIDIA Titan Xp 12Gb).

We argue that the accuracy is underestimated and the true score may be even higher. If we look at the portion of false negative error, as reported in Figure 3C, we may agree that some of those streamline trajectories can be considered anatomically non-plausible, even though the true label states the contrary. Therefore, the computation of accuracy is biased by this apparent mismatch, which we can consider as part of the noise in the labelling process. Moreover, we may claim that the sDEC model has a good generalization capability and behaves accurately even in the presence of misleading labels.

The geometric deep learning models seem to provide only a small, even though meaningful (t-test with p-value < 0.001), improvement in terms of accuracy compared to bLSTM. Nevertheless PN and sDEC support higher flexibility in the streamlines representation, both as points clouds or graphs. The sDEC model seems the most appropriate for tractogram filtering because it is robust with respect to the order of points in the streamlines while preserving the highest accuracy in the discrimination of anatomical plausibility.

The analysis of error distribution allows a better understanding of the difference between PN and sDEC models. Around 40% of the misclassification error concerns different streamlines, indicating that the two models behave differently: in Figure 3A the most critical streamlines are those longer and more curved. Nevertheless, sDEC is more robust and the drop in accuracy is lower than PN. We may explain this difference with the property of sDEC to capture the relations among the points because it is considering also the edges of streamlines. When streamlines are long and curved, edges become more informative and provide a competitive advantage. As illustrated in Figure 3B, PN has the bias to classify long and curved streamlines as anatomically plausible doing more false positive error.

We believe that a fast and accurate filtering of tractograms, like the one supported by the sDEC model, is the premise for further investigations on how tasks like bundle segmentation may take advantage of the removal of artifactual streamlines.

References

1. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A.: In vivo fiber tractography using DT-MRI data. *Magnetic resonance in medicine* 44(4), 625–632 (2000)
2. Côté, M.A., Garyfallidis, E., Larochelle, H., Descoteaux, M.: Cleaning up the mess: tractography outlier removal using hierarchical QuickBundles clustering. *Proceedings of: International Society of Magnetic Resonance in Medicine (ISMRM)* (2015)
3. Daducci, A., Dal Palù, A., Lemkaddem, A., Thiran, J.P.: COMMIT: convex optimization modeling for microstructure informed tractography. *IEEE transactions on medical imaging* 34(1), 246–257 (2015)
4. Fey, M., Lenssen, J.E.: Fast Graph Representation Learning with PyTorch Geometric. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds* (2019)
5. Garyfallidis, E., Brett, M., Correia, M.M., Williams, G.B., Nimmo-Smith, I.: QuickBundles, a Method for Tractography Simplification. *Frontiers in neuroscience* 6 (2012), <http://dx.doi.org/10.3389/fnins.2012.00175>
6. Girard, G., Whittingstall, K., Deriche, R., Descoteaux, M.: Towards quantitative connectivity analysis: reducing tractography biases. *NeuroImage* 98, 266–278 (Sep 2014), <http://www.sciencedirect.com/science/article/pii/S1053811914003541>
7. Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M.: The minimal preprocessing pipelines for the Human Connectome Project. *NeuroImage* 80, 105–124 (Oct 2013), <http://dx.doi.org/10.1016/j.neuroimage.2013.04.127>
8. Gupta, V., Thomopoulos, S.I., Rashid, F.M., Thompson, P.M.: FiberNET: An Ensemble Deep Learning Framework for Clustering White Matter Fibers. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*. pp. 548–555. *Lecture Notes in Computer Science*, Springer International Publishing, Cham (2017)
9. Huang, Z., Xu, W., Yu, K.: Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015)
10. Laurent Petit, François Rheault, Maxime Descoteaux, Nathalie Tzourio-Mazoyer: Half of the streamlines built in a whole human brain tractogram is anatomically uninterpretable. (2019), <https://ww5.aievolution.com/hbm1901/index.cfm?do=abs.viewAbs&abs=3220>
11. Maier-Hein, K., et al.: Tractography-based connectomes are dominated by false-positive connections. *bioRxiv* pp. 084137+ (Nov 2016), <http://dx.doi.org/10.1101/084137>
12. Maier-Hein, K.H., et al.: The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications* 8(1) (Dec 2017), <http://www.nature.com/articles/s41467-017-01285-x>
13. Masci, J., Rodolà, E., Boscaini, D., Bronstein, M.M., Li, H.: Geometric deep learning. In: *SIGGRAPH ASIA 2016 Courses*. pp. 1–50. SA '16, Association for Computing Machinery, Macau (Nov 2016), <https://doi.org/10.1145/2988458.2988485>
14. O'Donnell, L.J., Westin, C.F.F.: Automatic tractography segmentation using a high-dimensional white matter atlas. In: *IEEE Trans. Med. Imag.* pp. 1562–1575 (2007), <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.163.7804>

15. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L.: PyTorch: An imperative style, high-performance deep learning library. In: *Advances in Neural Information Processing Systems*. pp. 8024–8035 (2019)
16. Pestilli, F., Yeatman, J.D., Rokem, A., Kay, K.N., Wandell, B.A.: Evaluation and statistical inference for human connectomes. *Nature Methods* 11(10), 1058–1063 (Sep 2014), <http://dx.doi.org/10.1038/nmeth.3098>
17. Presseau, C., Jodoin, P.M., Houde, J.C., Descoteaux, M.: A new compression format for fiber tracking datasets. *NeuroImage* 109, 73–83 (Apr 2015), <http://www.sciencedirect.com/science/article/pii/S1053811914010635>
18. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 652–660 (2017), http://openaccess.thecvf.com/content_cvpr_2017/html/Qi_PointNet_Deep_Learning_CVPR_2017_paper.html, arXiv: 1612.00593
19. Rheault, F., St-Onge, E., Sidhu, J., Maier-Hein, K., Tzourio-Mazoyer, N., Petit, L., Descoteaux, M.: Bundle-specific tractography with incorporated anatomical and orientational priors. *NeuroImage* 186, 382–398 (Feb 2019), <http://www.sciencedirect.com/science/article/pii/S1053811918320883>
20. Smith, R.E., Tournier, J.D.D., Calamante, F., Connelly, A.: SIFT: Spherical-deconvolution informed filtering of tractograms. *NeuroImage* 67, 298–312 (Feb 2013), <http://view.ncbi.nlm.nih.gov/pubmed/23238430>
21. Tournier, J.D., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35(4), 1459–1472 (May 2007), <http://dx.doi.org/10.1016/j.neuroimage.2007.02.016>
22. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An overview. *NeuroImage* 80, 62–79 (Oct 2013), <http://dx.doi.org/10.1016/j.neuroimage.2013.05.041>
23. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)* 38(5), 146 (2019)
24. Yeh, F.C., Panesar, S., Barrios, J., Fernandes, D., Abhinav, K., Meola, A., Fernandez-Miranda, J.C.: Automatic Removal of False Connections in Diffusion MRI Tractography Using Topology-Informed Pruning (TIP). *Neurotherapeutics* 16(1), 52–58 (Jan 2019), <https://doi.org/10.1007/s13311-018-0663-y>
25. Zhuang, J., Hrabe, J., Kangarlou, A., Xu, D., Bansal, R., Branch, C.A., Peterson, B.S.: Correction of eddy-current distortions in diffusion tensor images using the known directions and strengths of diffusion gradients. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* 24(5), 1188–1193 (2006)