The do's and don'ts of reinforcement learning for tractography

Antoine Théberge[1] ANTOINE.THEBERGE@USHERBROOKE.CA
Christian Desrosiers[2] CHRISTIAN.DESROSIERS@ETSMTL.CA
Pierre-Marc Jodoin[*1] PIERRE-MARC.JODOIN@USHERBROOKE.CA
Maxime Descoteaux[*1] MAXIME.DESCOTEAUX@USHERBROOKE.CA

Editors: Under Review for MIDL 2022

Abstract

Tractography is the process of virtually reconstructing the white matter structure of the brain in a non-invasive manner. To tackle the various known problems of tractography, deep learning has been proposed, but the lack of well curated diverse datasets makes neural networks incapable of generalizing well on unseen data. Recently, deep reinforcement learning (RL) has been shown to effectively learn the tractography procedure without reference streamlines. While the performances reported were competitive, the proposed framework is complex and little is known on the role and impact of its multiple parts. In this work, we thoroughly explore the different components of the proposed framework through seven experiments on two datasets and shed light on their impact. Our goal is to guide researchers eager to explore the possibilities of deep RL for tractography by exposing what works and what does not work with this category of approach. We find that directionality is crucial for the agents to learn the tracking procedure and that the input signal and the seeding strategy offer a trade-offs in connectivity vs. volume.

Keywords: Tractography, reinforcement learning

1. Introduction

Tractography has recently been posed as a reinforcement learning (RL) problem (Théberge et al., 2021) so as to leverage the expressiveness of machine learning without the need for hard-to-obtain reference streamlines. Agents trained with this procedure demonstrated competitive performance compared to their supervised and classical alternatives on *in-silico* and *in-vivo* datasets.

Track-to-Learn (Théberge et al., 2021) is the only deep RL tractography method available as of today. However, it is a complex framework with many moving parts. In their original paper, the authors report results for only one configuration of input signal, seeding strategy, etc. It is thus unclear from that only paper what the effects of each components truly are. **In this work, we explore the impacts of different instantiations of *Track-to-Learn* to assess the do's and don'ts of reinforcement learning for tractography.**

2. Experiments

For each experiment, we vary only one component of the *Track-to-Learn* framework. **(1)** we report baseline performance by training a learning agent with the same framework components originally defined. **(2)** We then investigate the usefulness of the retracking

[*] Equal contributions, Co-senior authors

procedure when seeding from the white matter (WM) by removing it. **(3)** We further investigate the effects of the seeding strategy by instead seeding from the gray-matter/white-matter interface (GM/WM). **(4)** we gauge the effects of the components of the input signal by **(4.1)** using the raw diffusion signal instead of fiber ODF (fODF) SH coefficients, **(4.2)** removing the WM mask, **(4.3)** including 0 previous directions, **(4.4)** including 2 previous directions.

Two datasets were used for training and testing the agents: the FiberCup (Fillard et al., 2011) and the ISMRM2015 (Maier-Hein et al., 2017) synthetic phantoms. Per experiment and dataset, five agents were trained for 1000 epochs using the SAC algorithm (Haarnoja et al., 2018). A grid-search on the learning rate and discount factor were done for each experiment and dataset. Actors and critics each use three-layer, 1024-neurons wide fully connected networks. The Adam optimizer was used to update weights. States s_t are the concatenation of the input signal in the voxel and its six neighbours at the head of the streamline and the last n directions, depending on the experiment. We use the same reward function as in (Théberge et al., 2021):

$$r_t = |\max_{\boldsymbol{p}_i} \langle \boldsymbol{p}_i, \boldsymbol{a}_t \rangle| \cdot \langle \boldsymbol{a}_t, \boldsymbol{u}_{-1} \rangle \tag{1}$$

where $\boldsymbol{p}_i$ are the fODF peaks at the head of the streamline, $\boldsymbol{a}_t$ is the 3D vector produced by the agents, $\boldsymbol{u}_{-1}$ is the last streamline segment, $\boldsymbol{p}_i, \boldsymbol{a}_t, \boldsymbol{u}_{-1}$ are unit vectors. Tracking and seeding parameters were the same as in (Théberge et al., 2021). The Tractometer (Côté et al., 2013) was used to measure valid connections ratio (VC), number of valid bundles (VB), invalid connections ratio (IC), overlap ratio (OL) and no-connections ratio (NC). We compare the results of our agents to the Particle Filtering Tractography (PFT) classical algorithm (Girard et al., 2014) using both WM and GM/WM interface seeding.

3. Results

Table 1 reports mean scores extracted by the Tractometer from five tractograms per agent for all experiments and the PFT algorithms on the two datasets. First, we can observe that the RL agents in experiment 1 produce more accurate reconstructions (in terms of VC) than their classical PFT counterpart. The RL agents also produced fewer invalid (IC) and broken (NC) streamlines. The classical algorithm, however, produced fuller bundles (OL) and a comparable or higher number of bundles (VB). We can observe that some components are essential to the framework: removing the retracking procedure negatively impacted the reconstructed tractograms. Removing all previous directions prevented the agents from learning the tractography procedure entirely, but including only two directions did not significantly impact the performances. Some components in the framework offered a tradeoff in metrics: removing the WM mask or using the raw diffusion signal produced fuller tractograms, at the expense of valid connections. Seeding from the interface greatly improved the VC rate and reduced the NC and IC rates, but lowered overlap. Empirically, we also found that interface seeding greatly improved tracking time and stabilized training.

4. Conclusion

Results presented in this work reveal that reinforcement learning for tractography is a truly competitive ML-based alternative to classical algorithms. It also underlines the impacts of some of the *Track-to-Learn* components: directionality plays a crucial role in the learning

Table 1: Scores computed by the Tractometer. **Bold** indicates per dataset best results when compared to experiment 1. ↑ / ↓ indicates if higher or lower is better.

		VC % ↑	VB ↑	IC % ↓	OL % ↑	NC % ↓
FiberCup	PFT WM	27.9	7.0	13.7	99.4	58.3
	PFT Interf.	29.1	7.0	11.46	97.8	59.3
	1 Baseline	82.7	**7.0**	7.11	87.7	10.2
	2 No retrack	10.7	**7.0**	**4.4**	75.2	85.0
	3 Interf. seed.	**87.2**	**7.0**	10.2	72.7	**2.6**
	4.1 Raw diff.	76.3	**7.0**	12.6	**89.2**	11.1
	4.2 No WM	83.7	**7.0**	6.5	87.4	9.8
	4.3 0 dirs.	8.1	4.0	18.1	21.2	73.8
	4.4 2 dirs.	78.9	**7.0**	12.6	88.6	8.5
ISMRM2015	PFT WM	61.0	24.0	30.1	78.5	8.9
	PFT Interf.	54.1	23.2	29.0	40.0	17.0
	1 Baseline	69.5	**23.2**	23.0	52.7	7.5
	2 No retrack	42.5	23.0	41.7	39.7	15.8
	3 Interf. seed.	**74.7**	23.0	**18.3**	31.6	**7.0**
	4.1 Raw diff.	68.3	**23.2**	22.7	52.6	8.9
	4.2 No WM	68.1	**23.2**	24.5	**54.4**	7.4
	4.3 0 dirs.	7.5	12.4	49.5	2.3	43.0
	4.4 2 dirs.	70.6	23.0	21.8	52.9	7.6

procedure, but including two directions instead of four suffices. The retracking method was also shown to be crucially important. Varying the input signal and seeding strategy may allow the user to prioritize either volume reconstruction or connectivity.

References

Côté et al. Tractometer: Towards validation of tractography pipelines. *Medical Image Analysis*, 17(7):844–857, 2013. ISSN 1361-8415. doi: 10.1016/j.media.2013.03.009.

Fillard et al. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage*, 56(1):220–234, 2011. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2011.01.032.

Girard et al. Towards quantitative connectivity analysis: reducing tractography biases. *Neuroimage*, 98:266–278, 2014.

Haarnoja et al. Soft actor-critic algorithms and applications. *ArXiv*, abs/1812.05905, 2018.

Maier-Hein et al. The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8(1):1–13, 2017. ISSN 2041-1723. doi: 10.1038/s41467-017-01285-x.

Théberge et al. Track-to-learn: A general framework for tractography with deep reinforcement learning. *Medical Image Analysis*, page 102093, 2021.