



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

TractOracle: towards an anatomically-informed reward function for RL-based tractography

Antoine Théberge¹[0000-0002-3959-8316], Maxime Descoteaux¹[0000-0002-8191-2129], and Pierre-Marc Jodoin¹[0000-0002-6038-5753]

Université de Sherbrooke, Sherbrooke, Québec, Canada
antoine.theberge@usherbrooke.ca

Abstract. Reinforcement learning (RL)-based tractography is a competitive alternative to machine learning and classical tractography algorithms due to its high anatomical accuracy obtained without the need for any annotated data. However, the reward functions so far used to train RL agents do not encapsulate anatomical knowledge which causes agents to generate spurious false positive tracts. In this paper, we propose a new RL tractography system, *TractOracle*, which relies on a reward network trained for streamline classification. This network is used both as a reward function during training as well as a mean for stopping the tracking process early and thus reduce the number of false positive streamlines. This makes our system a unique method that evaluates and reconstructs WM streamlines at the same time. We report an improvement of true positive ratios by almost 20% and a reduction of 3x of false positive ratios on one dataset and an increase between 2x and 7x in the number true positive streamlines on another dataset.

Keywords: Tractography · Reinforcement Learning · Deep learning.

1 Introduction

Tractography is the process of virtually reconstructing the major white matter (WM) pathways of the human brain using diffusion magnetic resonance imaging (dMRI) [1]. Tractography has been successfully used in a number of applications such as presurgical planning [8], connectomics [28], tractometry [3] and disease progression modelling [22].

Despite many advances [6, 32, 10, 29], tractography still suffers from unresolved issues in the presence of complex fiber configurations [16, 24]. Tractography has been described as an *ill-posed problem* as it tries to infer global structure only from local information [16]. One common solution to this problem is to reconstruct overcomplete tractograms by seeding aggressively and then use a filtering tool, such as SIFT [27] or COMMIT [5] to remove unnecessary streamlines. However, this approach suffers from several limitations as it : (1) cannot recover *false negative* streamlines (i.e missing connections), (2) necessitates more data

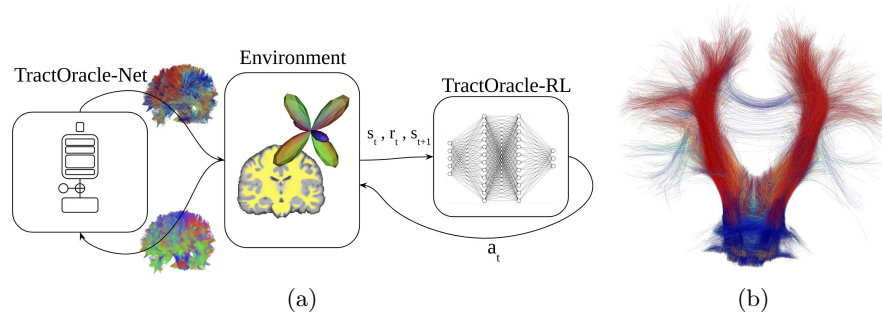


Fig. 1: (a) *TractOracle*: an RL system in which the environment sends streamlines to TractOracle-Net as they are being tracked to score their anatomical plausibility. The scores (illustrated as the colours of the streamlines) are used to reward the agent and stop the tracking process when streamlines diverge into an implausible shape. TractOracle-RL then uses the reward function to predict anatomically-informed tractograms. (b) TractOracle-Net scores along valid and invalid cortico-spinal tracts. Streamlines correctly terminating in the motor cortex get a high plausibility score (red); implausible streamlines diverging towards the corpus callosum get a low (blue) score.

handling and often (3) do not eliminate *false positive* streamlines completely as filtering algorithms are not immune to error [14].

Supervised machine learning approaches have been proposed to tackle this problem [2, 18, 20, 33] by training on curated sets of streamlines. However, datasets suitable for learning the tractography procedure are still few [21] and mostly limited to *in-silico* phantoms [4, 16].

Reinforcement learning (RL)-based tractography [31, 30] has been proposed as a way to learn tractography algorithms without reference streamlines. Instead, a reward function is used to guide the learning procedure to follow the orientation of local dMRI signal, as would typical tractography do. RL-based tractography agents tend to exhibit higher accuracy than classical tractography [30]. However, the proposed reward functions do not encapsulate anatomical priors and are thus prone to "reward hacking" issues [31, 25].

Here we propose a new RL tractography method trained to track and filter streamlines simultaneously. At the core of our system is a transformer network that predicts streamlines' anatomical score. This neural net acts as an oracle in our RL framework to reward streamlines based on the anatomical plausibility of their shape, guiding agent towards reconstructing streamlines with greater accuracy during training. Moreover, computing the plausibility score of a streamline while it is being tracked allows the system to stop the tracking when it diverges to an implausible path, preventing their reconstruction. This system is unique and exhibits state-of-the-art results in both tractogram segmentation and tractogram generation, outperforming competitive methods.

2 Preliminaries and proposed method

Tractography dMRI allows for virtual reconstruction of the WM pathways through a process called tractography. From an initial 3D point p_0 called a *seed*, a direction \mathbf{u}_0 of amplitude Δ is chosen according to a local model $v(p_0)$ (e.g. a fiber orientation distribution function, fODF [6]) of the diffusion signal at p_0 . Taking a step in $\Delta\mathbf{u}_0$ gives rise to p_1 , a new position in the WM. The process is repeated until a stopping criterion is met, i.e. exiting the WM mask. Formally:

$$p_{t+1} = p_t + \Delta\mathbf{u}_t, \quad \mathbf{u}_t \sim v(p_t). \quad (1)$$

The resulting ordered set of points $P = \{p_0, \dots, p_T\}$ with T being the length of the sequence is called a streamline, or tract. An ensemble of streamlines is called a tractogram. Seed points can be generated throughout the WM mask to ensure proper coverage of the WM, but the resulting streamlines may overrepresent major WM pathways [13]. Alternatively, seed points can be generated at the WM/grey matter (GM) interface, where axons are known to originate.

Reinforcement learning (RL) Tractography can be formulated as an RL learning problem [31]. In that case, the tractography system involves an agent π trained to predict tracking steps by interacting with its environment, i.e. 3D dMRI signal. RL uses the principles of the Markov Decision Process, describing the environment in terms of S the set of all possible states, A the set of all possible actions, $p(s_{t+1}|s_t, a_t)$ the transition probability between two consecutive states and a reward function $r(s_t, a_t)$ (r_t in short). An action $a_t \in A$ taken by the agent at state $s_t \in S$ leads to a new state $s_{t+1} \in S$ and a reward r_t given by the environment. By repeating the process of taking actions in states, a series of states and actions called a *trajectory* is generated. In tractography, a state s_t correspond to the dMRI signal in the vicinity of p_t , an action a_t corresponds to the translation vector \mathbf{u}_t between p_t and p_{t+1} , and a trajectory amounts to a WM streamline [31].

The goal of the agent is to optimize its policy π so as to maximize the expected discounted sum of future rewards, or return in short:

$$G_t = \sum_{k=t..T} [\gamma^k r(s_{t+k}, a_{t+k})] \quad (2)$$

where γ is a discount factor that prioritizes immediate rewards over distant ones. Central to RL policies are the concepts of value function $V_\pi(s_t)$ and Q function $Q_\pi(s_t, a_t)$. The value function estimates the expected return from a state s_t and subsequent trajectories under a policy π :

$$V_\pi(s) = \mathbb{E}_{s \sim \pi} [G_t | s_t = s]. \quad (3)$$

The Q function $Q_\pi(s_t, a_t)$ evaluates the expected return from taking an action a_t in state s_t and then following policy π , providing a direct measure for assessing the immediate and subsequent value of actions:

$$Q_\pi(s, a) = \mathbb{E}_\pi [r(s_t, a_t) + G_{t+1} | s_t = s, a_t = a]. \quad (4)$$

Previous RL tractography methods [30] are limited by their reward function $r(s_t, a_t)$ which considers only local information from the dMRI signal surrounding s_t . We instead reformulate the reward function to include a factor based on the anatomical plausibility of the reconstructed streamline from s_0 to s_t . This score is given by a neural network called *TractOracle-Net*.

TractOracle-Net TractOracle-Net is a transformer network [7] which takes streamlines as input and outputs scores related to their anatomical plausibility. Figure 1 illustrates how TractOracle-Net interacts with the environment to score streamlines. In the upcoming sections, it is represented as Ω_ψ where ψ are the network parameters.

Its input comprises the direction between the coordinates of a streamline that is resampled to a fixed number of points, to which a "SCORE" token is prepended. The transformer network consists of a 32-dimension embedding, positional encoding, 4 transformer encoder blocs each with 4 attention heads, and a linear and sigmoid layer to convert the SCORE token to a scalar between 0 and 1, for a total of 550k trainable parameters. TractOracle-Net is trained to perform regression using the mean-squared error as a loss function. For data augmentation, streamlines are randomly flipped, randomly cut (and then re-resampled to 128 points) and Gaussian noise is applied point-wise to the streamlines. We use 0.5 as a plausible/implausible classification threshold.

TractOracle-RL For TractOracle, the state s_t contains the fODFs at position p_t (i.e. the local diffusion orientation $v(p_t)$) to which we append the six surrounding fODFs and the previous 100 tracking directions [30]. The initial states s_0 are selected at the WM/GM interface, and the actions a_t produced by the agents are posed as the tracking directions \mathbf{u}_t , which the environment rescales to Δ to propagate streamlines.

The reward function is made of two terms: a local reward which accounts for how much the streamline segment at state s_t is aligned with the surrounding fODFs and its previous segment, and an anatomical score of the reconstructed streamline predicted by TractOracle-Net. Formally:

$$r_t = \underbrace{\left(\frac{|\max_{v(p_t)} \langle \overline{v(p_t)}, a_t \rangle| \cdot \langle a_t, a_{t-1} \rangle}{v(p_t)} \right)}_{local} + \alpha \underbrace{\mathbb{1}_{\Omega_\psi}(P_{0..t})}_{anatomical}, \quad (5)$$

with $\overline{v(p_t)}$ the maxima of $v(p_t)$ and $\mathbb{1}_{\Omega_\psi}$ is an indicator function of the score given by TractOracle-Net :

$$\mathbb{1}_{\Omega_\psi}(P_{0..t}) = \begin{cases} 1 & \text{if } \Omega_\psi(P_{0..t}) \geq 0.5 \text{ and } t = T \\ 0 & \text{else.} \end{cases} \quad (6)$$

Moreover, we propose a new stopping criterion based on the prediction of Ω_ψ : if the score given to a streamline after a number of steps t falls below 0.5,

the tracking stops. The stopping criteria become:

$$t = T \begin{cases} \text{if } \Omega_\psi(P_{0..t}) < 0.5 \text{ and } t > T_{min} \\ \text{if } \mathbb{1}_{WM}(p_t) < 0.1 \\ \text{if } (180/\pi)\langle \mathbf{u}_t, \mathbf{u}_{t-1} \rangle < \epsilon, \end{cases} \quad (7)$$

where T_{min} indicates the minimum number of steps before the criterion can be enforced, $\mathbb{1}_{WM}(p_t)$ is the value found via trilinear interpolation of the WM mask at position p_t and ϵ indicates the maximum angle (in degrees) between two streamline segments.

Implementation We set $\alpha = 10$, $T_{min} = 20$ and $\epsilon = 30$. We use the Soft-Actor Critic [12] RL algorithm to train the agents (c.f supplementary material 1 for an overview). The actor and twin critics are 3 layers fully-connected neural networks with a width of 1024. In our case, π outputs the mean and the standard deviation of a 3D Gaussian distribution which is sampled to obtain a_t , the tracking step. We set the learning rate and discount factor to 0.0005 and 0.95, respectively.

3 Experiments and results

3.1 Datasets and tools

ISMIRM2015 A synthetic dataset derived from global tractography on one subject of the Human Connectome Project [16, 17, 11]. We performed tracking using an ensemble of five algorithms [10, 6, 32], each generating 100k streamlines from both WM and WM/GM interface seeding. All other parameters were left to their respective default values. The resulting streamlines were segmented by the Tractometer [23] to obtain positive and negative examples.

BIL&GIN The Brain Imaging of Lateralization by the Groupe d’Imagerie Neurofonctionnelle (BIL&GIN) dataset is a publicly available dataset of 453 healthy adults participants. To train TractOracle-Net, we follow the experimental procedure as described in [15] to obtain the callosal fibres of 39 randomly selected subjects from the dataset.

TractoInferno Multi-site and multi-protocol dataset totalling 284 subjects. Because the reference bundles from [21] do not provide negative examples, we repeat the tracking procedure as presented in [21] to generate data for TractOracle-Net. We use Recobundles [9] to extract reference bundles and keep the unrecognized streamlines as negative examples, from which we select at random, per subject, 500k streamlines from recognized and unrecognized fibers respectively.

Table 1: Classification metrics on the ISMRM2015 dataset. FINTA results are reported from [15]. Best results per metric are highlighted in **bold**.

	Accuracy (%)	Sensitivity	Precision	F1-score
Recobundles	0.91	0.81	0.97	0.88
FINTA	0.91	0.91	0.91	0.91
TractOracle-Net	0.97	0.98	0.94	0.96

Table 2: Classification metrics (Mean \pm stdev) on 8 test subjects of BIL&GIN. Recobundles and FINTA results are reported from [15]. Best results are highlighted in **bold**.

	Accuracy	Sensitivity	Precision	F1-score
Recobundles	0.82 \pm 0.03	0.80 \pm 0.03	0.67 \pm 0.01	0.70 \pm 0.02
FINTA	0.91 \pm 0.01	0.91 \pm 0.01	0.78 \pm 0.01	0.83 \pm 0.01
TractOracle-Net	0.96 \pm 0.01	0.92 \pm 0.02	0.76 \pm 0.02	0.83 \pm 0.02

3.2 TractOracle-Net performance

To ensure that TractOracle-Net can properly guide the learning procedure of RL agents and validate the chosen architecture, we first need to ensure that it can classify streamlines based on their shape. We train two instances of TractOracle-Net on the ISMRM2015 and BIL&GIN datasets (c.f. section 3.1) for 50 and 200 epochs, respectively. We compare ourselves to the FINTA [14] streamline filtering algorithm and Recobundles [9] clustering algorithm. Recobundles uses the ground-truth ISMRM2015 bundles as atlas on the dataset.

Tables 1 and 2 report results for this experiment. Recobundles and FINTA obtain the highest precision on ISMRM2015 and BIL&GIN, respectively, but Recobundles has access to ground truth bundles as opposed to the other methods. Besides precision, TractOracle outperforms Recobundles and FINTA according to all other metrics on the ISMRM2015 and BIL&GIN datasets.

3.3 TractOracle *in-silico* performance

We then investigate the improvements brought by TractOracle-RL over other tractography methods by training and tracking on the ISMRM2015 dataset. All agents were trained for 1000 epochs, five times with different random seeds. We used the TractOracle-Net model trained on the ISMRM2015 dataset from the previous experiment. We compare the proposed method against some of the most widely used algorithms: the `sd_stream` and `ifod2` algorithms from MRTrix3 [32] as well as Track-to-Learn [30]. For all methods, we track five times at 20 seeds per voxels at the WM/GM interface using five different random seeds and report mean Tractometer metrics.

Table 3 reports Tractometer metrics for all agents considered in this experiment. As shown in red, TractOracle obtains by far the highest Valid Connection (VC) rate (22% better than the 2nd best method), as well as the lowest Invalid Connection (IC) (3x lower than the 2nd best) and No-connection (NC) rates

Table 3: Tractometer scores (mean \pm stddev) on ISMRM2015 for TractOracle-RL (Proposed), Track-to-Learn, and classical tractography algorithms. Scores in **bold** indicate the best method for each metric, scores in **red** indicate the method is severely outperforming the others.

	VC % \uparrow	VB (/21) \uparrow	IC % \downarrow	IB \downarrow
sd_stream	55.96 \pm 0.21	19.00 \pm 0.00	44.04 \pm 0.21	199.80 \pm 4.26
ifod2	31.53 \pm 0.20	19.00 \pm 0.00	68.47 \pm 0.20	281.00 \pm 4.00
Track-to-Learn	66.13 \pm 1.15	20.00 \pm 0.63	33.87 \pm 1.15	293.40 \pm 11.8
TractOracle-RL	88.05 \pm 0.35	19.33 \pm 0.47	11.95 \pm 0.35	195.67 \pm 4.99
	OL % \uparrow	OR % \downarrow	F1 % \uparrow	NC \downarrow
sd_stream	38.85 \pm 0.05	3.59 \pm 0.05	52.47 \pm 0.03	9.36 \pm 0.18
ifod2	48.70 \pm 0.11	8.18 \pm 0.19	59.10 \pm 0.07	12.45 \pm 0.08
Track-to-Learn	53.84 \pm 2.28	29.94 \pm 2.00	57.43 \pm 1.84	2.85 \pm 0.42
TractOracle-RL	48.43 \pm 0.64	17.68 \pm 1.00	57.00 \pm 0.47	0.73 \pm 0.12

(3x lower than the 2nd best). While the proposed method reduces the overlap (OL) of recovered bundles compared to Track-to-Learn, it also reduces the overreach (OR) preserving a similar F1 score, as well as greatly reducing the number of invalid bundles (IB) reconstructed. Compared to classical tractography, the proposed method has a similar number of VB and OL and F1 rates while having significantly higher VC rates and significantly lower IB numbers, IC, NC rates, the best scores ever reported on this dataset.

3.4 TractOracle *in-vivo* performance

After validating the *in-silico* performance of our agents, we turn to *in-vivo* subjects. We train TractOracle-Net on the reference streamlines of TractoInferno as described in section 3.1 for 10 epochs. We report an accuracy of 89.21% on the TractoInferno test split. We then train TractOracle-RL agents on the TractoInferno dataset and track on the test set, repeating the experimental procedure and reusing the same hyperparameters from the previous experiment.

We segment the tractograms using Recobundles to provide a qualitative comparison to reference bundles of the test set. We additionally quantitatively evaluate the anatomical accuracy of the reconstructed tractograms by classifying them using Recobundles [9], TractOracle-Net and extractor_flow [19], an automated set of anatomical rules derived from current anatomical knowledge, and report the number of streamlines that were classified as plausible by the tools. We do not report results from the TractoInferno pipeline [21] as the pipeline evaluates the volume of reconstructed bundles, where as our method is focused on connection accuracy. We argue the pipeline does not report relevant metrics for the goal of this work.

Figure 2 displays some bundles from subject 1006 of the dataset and as reconstructed by all methods (additional bundles are available in supplementary material B). We can observe that the proposed method reconstructs highly visually appealing bundles, with excellent fanning. Moreover, the proposed method

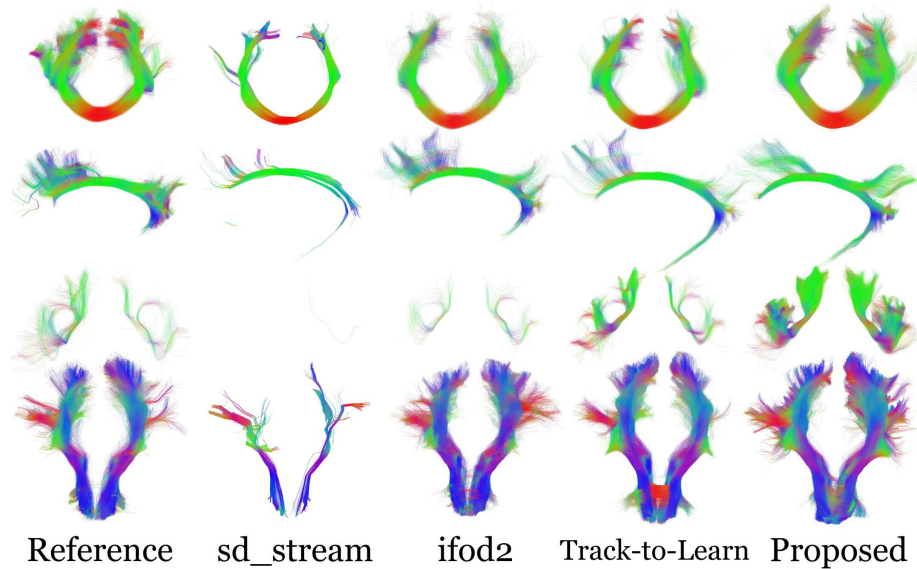


Fig. 2: Visualization of the occipital part of the corpus callosum (1st row), the left cingulum (2nd row), the left-right uncinate fasciculus (3rd row) and the left-right parieto-occipito pontine tracts (4th row) from subject 1006 of the Tractoinferno dataset (reference) and as reconstructed by all methods considered in this work.

Table 4: Total number of streamlines recovered by Recobundles [9], extractor_flow [19] and TractOracle-Net for all tracking algorithms considered.

	Recobundles \uparrow	extractor_flow \uparrow	TractOracle-Net \uparrow
sd_stream	2,713,507	7,698,181	28,901,678
ifod2	2,362,290	15,833,504	24,228,612
Track-to-Learn	9,309,681	46,135,363	85,447,863
TractOracle-RL	25,028,287	55,761,074	184,355,412

tends to reconstruct parts of the bundles that even the reference dataset misses, such as the "tail" of the cingulum, most of the uncinate fasciculus and some fanning of the parieto-occipito pontine tracts. sd_stream provides generally thin bundles, whereas ifod2 and Track-to-Learn produce similarly voluminous bundles.

Table 4 underlines that TractOracle produces highly anatomically accurate tractograms. Despite using the same number of seeding points for all methods, the number of plausible streamlines reported by TractOracle is drastically higher than other methods regardless of the classification tool used, going from an average improvement of 7x for sd_stream and ifod2, to 2x for Track-to-Learn.

4 Discussion and conclusion

In this work, we presented TractOracle, a highly accurate tractography system which evaluates and generates streamlines at the same time. From the results above, TractOracle-Net outperforms state-of-the-art streamline classification algorithms (c.f. tables 1, 2) while TractOracle-RL produces highly accurate (c.f. tables 3, 4) and voluminous (c.f. figure 2) tractograms. As of today, TractOracle is the most effective machine learning WM tracking method.

Future work should aim at improving the overlap between the reconstructed tractograms and reference bundles (c.f. table 3). Future work could also include further anatomical information via the use of Anatomically-Constrained Tractography [26]. Limitations and failure modes, including in the presence of pathology, need to be explored. But most importantly, future work should aim at giving the tracking agent global context of the diffusion volume to finally tackle the ill-posedness of tractography.

Acknowledgments. This work was funded by the Fonds de recherche du Québec - Nature et technologie.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

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. Benou, I., Riklin Raviv, T.: Deeptract: A probabilistic deep learning framework for white matter fiber tractography. In: *MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III* 22. pp. 626–635. Springer (2019)
3. Colby, J.B., Soderberg, L., Lebel, C., Dinov, I.D., Thompson, P.M., Sowell, E.R.: Along-tract statistics allow for enhanced tractography analysis. *NeuroImage* **59**(4), 3227–3242 (2012)
4. Côté, M.A., Girard, G., Boré, A., Garyfallidis, E., Houde, J.C., Descoteaux, M.: Tractometer: Towards validation of tractography pipelines. *Medical Image Analysis* **17**(7), 844–857 (2013), special Issue on the 2012 Conference on Medical Image Computing and Computer Assisted Intervention
5. 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)
6. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical q-ball imaging. *Magnetic Resonance in Medicine* **58**(3), 497–510 (2007)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
8. Essayed, W.I., Zhang, F., Unadkat, P., Cosgrove, G.R., Golby, A.J., O’Donnell, L.J.: White matter tractography for neurosurgical planning: A topography-based review of the current state of the art. *NeuroImage: Clinical* **15**, 659–672 (2017)
9. Garyfallidis, E., Côté, M.A., Rheault, F., Sidhu, J., Hau, J., Petit, L., Fortin, D., Cunanne, S., Descoteaux, M.: Recognition of white matter bundles using local and global streamline-based registration and clustering. *NeuroImage* **170**, 283–295 (2018), segmenting the Brain
10. Girard, G., Whittingstall, K., Deriche, R., Descoteaux, M.: Towards quantitative connectivity analysis: reducing tractography biases. *NeuroImage* **98**, 266–278 (2014)
11. Glasser, M.F. et al.: The human connectome project’s neuroimaging approach. *Nature Neuroscience* **19**(9), 1175–1187 (2016)
12. Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al.: Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018)
13. Jeurissen, B., Descoteaux, M., Mori, S., Leemans, A.: Diffusion MRI fiber tractography of the brain. *NMR in Biomedicine* **32**(4), e3785 (2019)
14. Jörgens, D., Descoteaux, M., Moreno, R.: Challenges for tractogram filtering. In: *Anisotropy Across Fields and Scales*. pp. 149–168. Springer International Publishing, Cham (2021)
15. Legarreta, J.H., Petit, L., Rheault, F., Theaud, G., Lemaire, C., Descoteaux, M., Jodoin, P.M.: Filtering in tractography using autoencoders (finta). *Medical Image Analysis* **72**, 102126 (2021)
16. Maier-Hein, K.H. et al.: The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications* **8**(1) (2017). <https://doi.org/10.1038/s41467-017-01285-x>, <http://dx.doi.org/10.1038/s41467-017-01285-x>

17. Neher, P.F., Reicht, I., van Bruggen, T., Goch, C., Reisert, M., Nolden, M., Zelzer, S., Meinzer, H.P., Stieltjes, B., Fritzsche, K.H.: Mitk diffusion imaging. *Methods of Information in Medicine* **51**(05), 441–448 (2012). <https://doi.org/10.3414/me11-02-0031>, <http://dx.doi.org/10.3414/ME11-02-0031>
18. Neher, P.F., Côté, M.A., Houde, J.C., Descoteaux, M., Maier-Hein, K.H.: Fiber tractography using machine learning. *Neuroimage* **158**, 417–429 (2017)
19. Petit, L., Ali, K.M., Rheault, F., Boré, A., Cremona, S., Corsini, F., De Benedictis, A., Descoteaux, M., Sarubbo, S.: The structural connectivity of the human angular gyrus as revealed by microdissection and diffusion tractography. *Brain Structure and Function* **228**(1), 103–120 (2023)
20. Poulin, P., Côté, M.A., Houde, J.C., Petit, L., Neher, P.F., Maier-Hein, K.H., Larochelle, H., Descoteaux, M.: Learn to track: deep learning for tractography. In: MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part I 20. pp. 540–547. Springer (2017)
21. Poulin, P., Theaud, G., Rheault, F., St-Onge, E., Bore, A., Renauld, E., de Beaumont, L., Guay, S., Jodoin, P.M., Descoteaux, M.: Tractoinferno—a large-scale, open-source, multi-site database for machine learning dmri tractography. *Scientific Data* **9**(1), 725 (2022)
22. Raj, A., Kuceyeski, A., Weiner, M.: A network diffusion model of disease progression in dementia. *Neuron* **73**(6), 1204–1215 (2012)
23. Renauld, E., Théberge, A., Petit, L., Houde, J.C., Descoteaux, M.: Validate your white matter tractography algorithms with a reappraised ismrm 2015 tractography challenge scoring system. *Scientific Reports* **13**(1), 2347 (2023)
24. Rheault, F., Poulin, P., Caron, A.V., St-Onge, E., Descoteaux, M.: Common misconceptions, hidden biases and modern challenges of dmri tractography. *Journal of neural engineering* **17**(1), 011001 (2020)
25. Skalse, J., Howe, N., Krashennnikov, D., Krueger, D.: Defining and characterizing reward gaming. In: *NeurIPS*. vol. 35, pp. 9460–9471 (2022)
26. Smith, R.E., Tournier, J.D., Calamante, F., Connelly, A.: Anatomically-constrained tractography: improved diffusion mri streamlines tractography through effective use of anatomical information. *Neuroimage* **62**(3), 1924–1938 (2012)
27. Smith, R.E., Tournier, J.D., Calamante, F., Connelly, A.: Sift: Spherical-deconvolution informed filtering of tractograms. *NeuroImage* **67**, 298–312 (2013)
28. Sotiropoulos, S.N., Zalesky, A.: Building connectomes using diffusion mri: why, how and but. *NMR in Biomedicine* **32**(4), e3752 (2019)
29. St-Onge, E., Daducci, A., Girard, G., Descoteaux, M.: Surface-enhanced tractography (set). *NeuroImage* **169**, 524–539 (2018)
30. Théberge, A., Desrosiers, C., Boré, A., Descoteaux, M., Jodoin, P.M.: What matters in reinforcement learning for tractography. *Medical Image Analysis* **93**, 103085 (2024)
31. Théberge, A., Desrosiers, C., Descoteaux, M., Jodoin, P.M.: Track-to-learn: A general framework for tractography with deep reinforcement learning. *Medical Image Analysis* **72**, 102093 (2021)
32. Tournier, J.D., Smith, R., Raffelt, D., Tabbara, R., Dhollander, T., Pietsch, M., Christiaens, D., Jeurissen, B., Yeh, C.H., Connelly, A.: Mrtrix3: A fast, flexible and open software framework for medical image processing and visualisation. *Neuroimage* **202**, 116137 (2019)
33. Wegmayr, V., Buhmann, J.M.: Entrack: Probabilistic spherical regression with entropy regularization for fiber tractography. *International Journal of Computer Vision* **129**(3), 656–680 (Nov 2020)