



HAL
open science

Filtering in tractography using autoencoders (FINTA)

Jon Haitz Legarreta, Laurent Petit, François Rheault, Guillaume Theaud,
Carl Lemaire, Maxime Descoteaux, Pierre-Marc Jodoin

► **To cite this version:**

Jon Haitz Legarreta, Laurent Petit, François Rheault, Guillaume Theaud, Carl Lemaire, et al.. Filtering in tractography using autoencoders (FINTA). *Medical Image Analysis*, 2021, 72, pp.102126. 10.1016/j.media.2021.102126 . hal-03346557

HAL Id: hal-03346557

<https://hal.science/hal-03346557v1>

Submitted on 16 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FILTERING IN TRACTOGRAPHY USING AUTOENCODERS (FINTA)

PREPRINT. PUBLISHED IN MEDICAL IMAGE ANALYSIS 72 (2021) 102126
DOI: <https://doi.org/10.1016/j.media.2021.102126>

Jon Haitz Legarreta

Sherbrooke Connectivity Imaging Laboratory (SCIL)
Videos & Images Theory and Analytics Laboratory (VITAL)
Department of Computer Science
Université de Sherbrooke, Canada

Laurent Petit

Groupe d'Imagerie Neurofonctionnelle (GIN)
Univ. Bordeaux, CNRS, CEA, IMN, UMR 5293, France

François Rheault

Sherbrooke Connectivity Imaging Laboratory (SCIL)
Department of Computer Science
Université de Sherbrooke, Canada

Guillaume Theaud

Sherbrooke Connectivity Imaging Laboratory (SCIL)
Department of Computer Science
Université de Sherbrooke, Canada

Carl Lemaire

Centre de Calcul Scientifique
Université de Sherbrooke, Canada

Maxime Descoteaux*

Sherbrooke Connectivity Imaging Laboratory (SCIL)
Department of Computer Science
Université de Sherbrooke, Canada

Pierre-Marc Jodoin*

Videos & Images Theory and Analytics Laboratory (VITAL)
Department of Computer Science
Université de Sherbrooke, Canada

*Co-senior author. These authors contributed equally.

July 31, 2021

ABSTRACT

Current brain white matter fiber tracking techniques show a number of problems, including: generating large proportions of streamlines that do not accurately describe the underlying anatomy; extracting streamlines that are not supported by the underlying diffusion signal; and under-representing some fiber populations, among others. In this paper, we describe a novel autoencoder-based learning method to filter streamlines from diffusion MRI tractography, and hence, to obtain more reliable tractograms. Our method, dubbed FINTA (*Filtering in Tractography using Autoencoders*) uses raw, unlabeled tractograms to train the autoencoder, and to learn a robust representation of brain streamlines. Such an embedding is then used to filter undesired streamline samples using a nearest neighbor algorithm. Our experiments on both synthetic and in vivo human brain diffusion MRI tractography data obtain accuracy scores exceeding the 90% threshold on the test set. Results reveal that FINTA has a superior filtering performance compared to conventional, anatomy-based methods, and the RecoBundles state-

of-the-art method. Additionally, we demonstrate that FINTA can be applied to partial tractograms without requiring changes to the framework. We also show that the proposed method generalizes well across different tracking methods and datasets, and shortens significantly the computation time for large (> 1 M streamlines) tractograms. Together, this work brings forward a new deep learning framework in tractography based on autoencoders, which offers a flexible and powerful method for white matter filtering and bundling that could enhance tractometry and connectivity analyses.

Keywords Representation Learning · Autoencoder · diffusion MRI · Tractography · Filtering

1 Introduction

Tractography uses local orientation information from diffusion Magnetic Resonance Imaging (dMRI) data to delineate the brain white matter fiber pathways, and hence, to provide information about its structural connectivity. These connectivity pathways in tractograms are composed of streamlines virtually representing fascicles of white matter fibers. Tractography methods face a number of challenges when propagating streamlines, including avoiding the early termination of the tracking procedure (Smith et al., 2012; Girard et al., 2014), providing streamlines between gray matter regions that are known to be connected while avoiding spurious streamlines, ensuring the full occupancy of the white matter volume by the streamlines (Rheault et al., 2019), and a complete gray matter surface coverage when streamlines reach the cortex (St-Onge et al., 2018). As a result, and despite the efforts to address these issues, white matter tracking procedures are known to produce a disproportionately large amount of invalid streamlines (referred to as “implausibles” in this work, as opposed to “plausible” streamlines). The category of invalid streamlines spans a broad group of streamlines that are known not to satisfy accepted anatomical constraints. These include streamlines that contain loops or sharp bends; streamlines that stop in non-gray matter tissues, such as the cerebrospinal fluid (CSF); streamlines that prematurely stop in the white matter; or streamlines describing trajectories between gray matter regions that are not connected structurally. An illustration of plausible and implausible streamlines can be seen in figure 1. According to the ISMRM 2015 Tractography Challenge benchmark (Maier-Hein et al., 2017), tracking methods produce a non-negligible proportion of invalid streamlines as a trade-off between sensitivity and specificity. The study reported that non-existing, false positive connections were present in as many as 64% of the submissions. Hence, improved filtering strategies are required to detect and remove anatomically unfeasible streamlines and mitigate some of the limitations derived from current streamline propagation methods.

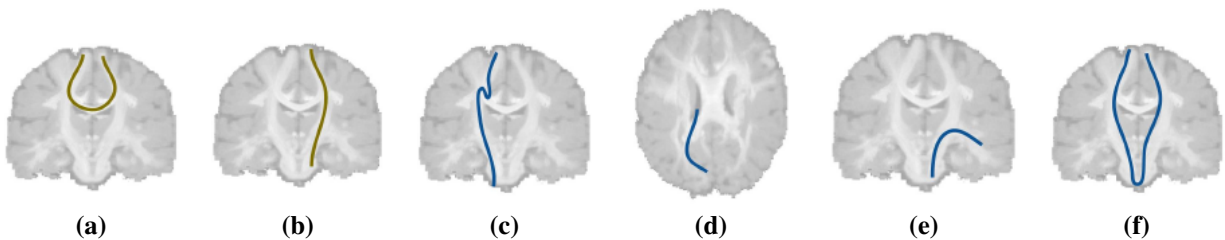


Figure 1: Schematic representation of plausible and implausible streamlines: (a) plausible streamlines belonging to the corpus callosum; (b) plausible streamlines belonging to the corticospinal tract; (c) implausible streamline containing a loop or a sharp bend; (d) implausible streamline stopped at the cerebrospinal fluid; (e) implausible streamline stopped in the white matter (or hitting a white matter or atlas boundary); (f): implausible streamline describing an invalid bundle.

Filtering implausible streamlines generated by tracking methods is a necessary step to avoid biologically unrealistic connectivity patterns and biases in reconstructing the brain structural connectome from diffusion MRI data (Sotiropoulos and Zalesky, 2019; Yeh et al., 2019) (see section 1.1.1). The quality of the tractograms influence various aspects of the connectome analysis, including the connectivity density, the clustering degree or the existence of connections themselves, among others (Yeh et al., 2016). Hence, tractogram filtering, in its various forms, is essential towards providing a reliable imaging evidence using diffusion MRI of how the white matter fascicles are arranged in the brain.

Filtering in tractography can also be regarded as a procedure that separates streamlines that are useful to serve a certain task from those that are not. As such, some filtering methods involve a clustering operation which may be used to extract a predefined set of bundles (see section 1.1.2). While connectomics applications (Sotiropoulos and Zalesky, 2019) require every anatomically plausible fiber to be recovered, other applications often need a more specialized set of fibers. For example, in the human *versus* macaque study by Takemura et al. (2017), only streamlines from the occipital white matter are considered. In Chenot et al. (2019) only fibers from the pyramidal tract are retained, while in Bullock et al. (2019) the brain fibers connecting the dorsal and ventral posterior cortex are the ones of interest.

Streamline filtering can be seen as an application of choice for deep learning classification methods. Unfortunately, neural networks are not easily applicable to brain tractography, mostly because of the fundamental difficulty of building a labeled dataset. In fact, finding irreproachable ground truth streamlines is a very difficult endeavour, with significant differences in reproducibility measures even among well-trained expert neuroanatomists, as shown in Rheault et al. (2020a). Also, depending on the task at hand, the labeling of the streamlines might vary considerably, and hence may require time-consuming manual verification and/or relabeling. Classification neural networks are trained to predict a fixed number of classes, and thus cannot be used to predict a different set of classes without being retrained on a newly labeled set of data.

This work proposes a deep autoencoder for streamline-based tractography filtering. The method, named FINTA (*Filtering in Tractography using Autoencoders*), is trained on non-curved, raw tractograms. Once training is over, the learned latent space is a low-dimensional representation of the input streamlines where similar streamlines are located next to each other. The filtering process is then carried out by projecting to the latent space examples of reference streamlines. Then, the to-be-filtered streamlines are projected to the latent space and labeled according to a nearest neighbor test. We demonstrate the filtering ability of FINTA on both synthetic and human brain datasets. Additional quantitative evaluations show its superior performance compared to both anatomy-based tractography filtering methods and the state-of-the-art RecoBundles (Garyfallidis et al., 2018; Rheault, 2020) streamline similarity filtering method. Similarly, FINTA is linear in terms of the streamline count at test time, allowing for a faster tractogram filtering.

1.1 Related work

1.1.1 Filtering plausible *versus* implausible streamlines

Diffusion MRI tractography methods generate tractograms that may contain several million candidate streamlines representing white matter fiber populations. Yet, a large proportion of such streamlines are anatomically implausible, meaning that they do not comply with neuroanatomical constraints attributed to fiber populations. A regular tractogram including only anatomically plausible streamlines can contain in the order of 500 000 to 3 000 000 streamlines (Presseau et al., 2015). Hence, having an automatic procedure to filter out implausible streamlines is critical. Tractography filtering is currently approached based on five main patterns: (i) streamline geometry features; (ii) region-of-interest-driven streamline inclusion and exclusion; (iii) clustering; (iv) connectivity-based; and (v) diffusion signal mapping or forward models (Jeurissen et al., 2019; Sotiropoulos and Zalesky, 2019; Jörgens et al., 2020). The first category of criteria is based on anatomical features of individual streamlines, such as unfeasible streamline length or local curvature indices. Streamline inclusion and exclusion criteria, adopted, for example, in Catani et al. (2002); Wakana et al. (2007); Yendiki et al. (2011); Wassermann et al. (2016); Warrington et al. (2020), incorporate white matter and/or tissue local and connectivity constraints for streamline traversal (Schilling et al., 2020). Such principles are usually concatenated into a sequence of restrictions to form a succession of step-wise refinement rules (Maier-Hein et al., 2017; Sarubbo et al., 2019; Jeurissen et al., 2019; Jörgens et al., 2020). Clustering approaches, such as QuickBundles (Garyfallidis et al., 2012), use a similarity measure to remove non-meaningful data issued by a tracking method. These methods are based on the assumption that a properly defined distance measure is able to split apart the relevant streamlines from the rest. These methods usually require a representative set of streamlines of the population that is being studied to be built so that they can be used as reference models to exclude the extraneous streamlines. Connectivity-based approaches, such as the one presented by Wang et al. (2018), propose to remove undesired streamlines usually imposing some regularization constraint on the tractogram-derived connectivity matrices. The forward models, such as SIFT/SIFT2 (Smith et al., 2013, 2015), LiFE (Pestilli et al., 2014) and COMMIT/COMMIT2 (Daducci et al., 2015; Schiavi et al., 2020), posit that only a subset of streamlines in a whole-brain tractogram are essential or relevant to explain the underlying diffusion signal. SIFT/SIFT2 is based on modifying the local fiber Orientation Distribution Function (fODF) based on the diffusion signal, whereas LiFE and COMMIT/COMMIT2 use specific local models to weigh the contribution of the signal to the streamline representation. Anatomical *a priori* constraints can be incorporated to improve the accuracy of the reconstructed pathways. Some of these methods have large computational demands in terms of the time required to filter the tractogram of a single subject (in the order of hours) (O'Donnell et al., 2019).

A number of works use deep learning-based methods for tractography filtering. de Lucena (2018) used an autoencoder to learn features about the input tractogram, and then used a residual network to regress on those features. However, their results were only limited to the uncinate fascicle and no comparison to other filtering methods was provided. More recently Astolfi et al. (2020) proposed a geometrical deep learning method to filter tractograms. The method implements a supervised geometric model which builds on streamline graph structures.

1.1.2 Streamline clustering

White matter streamline clustering (also referred to as *bundling*) allows to assign every streamline in a tractogram into anatomically coherent bundles (Maddah et al., 2008; Li et al., 2010; Wassermann et al., 2010; Guevara et al., 2011;

Garyfallidis et al., 2012; Jin et al., 2014; Wassermann et al., 2016; Siless et al., 2018; Sharmin et al., 2018; Garyfallidis et al., 2018). It is a procedure where the streamlines of interest need to be identified from the rest of the streamlines (the rejection class). The task of attributing a streamline to a bundle needs to address the fundamental question of how streamlines are similar to, or different from, each other. One of the challenges lies in finding the appropriate measure to quantify the affinity of streamlines, the other major difficulty being computing such affinities or distances in a reasonable time.

Probabilistic model-based methods (e.g. Dayan et al. (2018)) or traditional machine learning-based frameworks (e.g. Brun et al. (2004); Wang and Shi (2018); Kumar et al. (2019); Bertò et al. (2020)) have also been proposed to solve the clustering task. However, some of the findings in these works are limited to clustering a single bundle, require separate learning models for each bundle, are still affected by the streamline-wise distance computational complexity, or suffer from the trade-off of the granularity of the arrangement.

Some works (Brun et al., 2004; O’Donnell et al., 2006; O’Donnell and Westin, 2007; Sydnor et al., 2018) have investigated the feasibility of feature spaces to perform the streamline clustering, and have highlighted the ability of such embeddings to overcome some of the limitations of region-of-interest-based methods.

A number of different strategies have been developed recently to address the streamline clustering task using deep learning networks, including stacked, bi-directional siamese networks (FS2Net) (Patil et al., 2017), regular convolutional neural networks using different streamline re-parameterizations (FiberNet/FiberNet2.0) (Gupta et al., 2017, 2018), or custom feature vectors (TRAFIC) (Lam et al., 2018), (DeepWMA) (Zhang et al., 2020), ensemble models using fully connected neural networks (Ugurlu et al., 2018), or representation learning methods (Zhong et al., 2020) based on recurrent autoencoders. Yet, many of these methods still work in a supervised fashion, and thus do not allow to generalize to a variable number of classes, require a distinctly trained network for each bundle, or present results limited to a reduced number of bundles.

As opposed to the previous deep learning approaches, a few other ones (Wasserthal et al. (2018); Pomiecko et al. (2019); Li et al. (2020)) provide voxel-wise segmentations representing the locations traversed by a given bundle’s streamlines. Most of these approaches are based on a U-Net (Ronneberger et al., 2015) or another form of classification convolutional neural network using a specialized feature vector or parameterization. The major drawback of voxel-based bundle segmentation methods is the need of a separate model (including the need to train it) for each bundle, and hence, a variable number of models depending on the target classes.

In this work, we describe a novel unsupervised learning framework based on deep autoencoders to filter streamlines in diffusion MRI tractography. The proposed autoencoder is trained in an unsupervised fashion, and an optimal filtering threshold is computed in the learned representation space using a set of labeled streamlines. We demonstrate its performance on both synthetic and in vivo human brain tractography data, including partial tractograms. We provide evidence that the described deep learning setting can be applied to filter plausible and implausible streamlines as well as to multi-label classification problems without requiring to be retrained.

2 Material and methods

Autoencoders (AEs) are deep neural networks that have the ability to learn an efficient representation of the data in an unsupervised fashion using a dimensionality reduction approach (Hinton and Salakhutdinov, 2006). AEs are trained to reconstruct an input signal through a series of compressing operations (known as the *encoder*) followed by a series of decompression operations (referred to as the *decoder*). As shown in figure 2, between the encoder and the decoder there is the so-called *latent space*, where each point is an encoded representation of an input data sample. In this work, the input data is a *streamline*, which is a sequence of three-dimensional points. A streamline is an ordered sequence of points $s = \{c_1, c_2, \dots, c_{n_s}\}$, $c_i \in \mathbb{R}^3$ that represents a package of similarly oriented fibers describing a neural pathway within the brain white matter. A tractogram representing a set of M streamlines is mathematically expressed as $T = \{s_1, s_2, \dots, s_M\}$.

AEs have two major advantages over other types of neural networks. First, they are trained on raw unlabeled data, which is a major asset in the context of brain tractography. Second, since the latent representation is obtained by a series of compressing operations, two neighboring points in the latent space correspond to input data instances that share similar characteristics (Painchaud et al., 2019). Thus, the encoded representation can be used to redefine the notion of inter-data distance. In the context of tractography, instead of measuring the distance between two streamlines with a geometric distance function as is usually the case (Siless et al., 2013), one can project the streamlines into the latent space and measure their Euclidean distance.

Our filtering approach, shown in figure 2, can be summarized by the following steps:

1. Train an autoencoder with a large, raw, uncurated tractogram.
2. From the training tractogram, label the streamlines of interest. These streamlines could be examples of anatomically plausible fibers, streamlines representing fibers belonging to predefined sets of bundles, or any sorts of streamlines that are of interest for some application. These streamlines are labeled as *positive streamlines*.
3. Project the *positive* streamlines into the latent space with the encoder network. The positive latent vectors are illustrated by the golden circles in figure 2. Note that in case of more than two classes (as for a multi-label bundling operation) streamlines of several classes could also be projected into the latent space.
4. Given a new tractogram that ought to be filtered, project its streamlines into the latent space (cf. the three bold circles in figure 2(b)) and label them according to the distance to their nearest neighbor.

Hence, FINTA's filtering process actually takes place in the latent space. The nearest neighbor discrimination employs a latent space distance cut-off value, herein named *filtering threshold*, relative to a (labeled) reference set of streamlines to discriminate the uncurated tractogram's streamlines. The filtering threshold represents, thus, the minimum distance at which a streamline is considered to be implausible, and its value is computed on the separate reference set of streamlines. Section 3 provides details about the criterion adopted to compute the optimal filtering threshold.

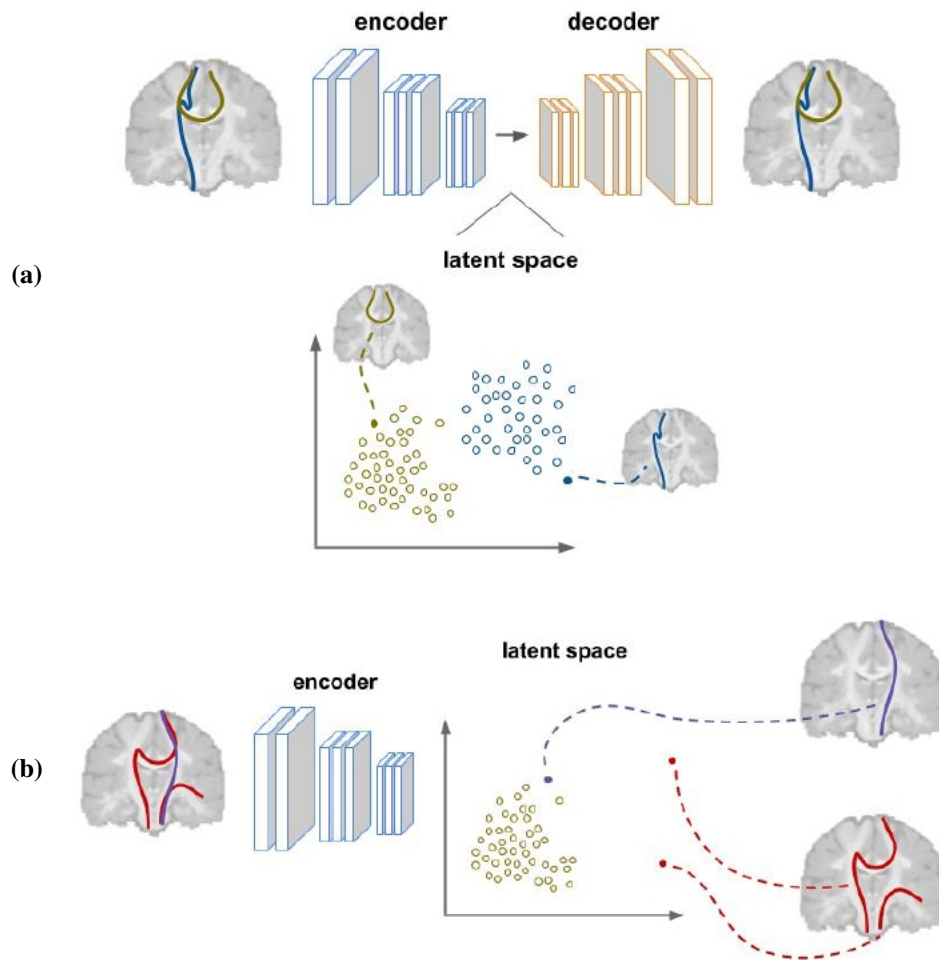


Figure 2: Schematic representation of tractography filtering with FINTA: (a) training time; (b) test time. Although 1D convolutions are used in the proposed framework, convolutional blocks are depicted using a 2D shape for illustrative purposes.

The proposed AE is a fully convolutional neural network whose overall structure is summarized in table 1. Our convolutional AE accepts a streamline at its input. Since raw streamlines have a different number of vertices (long streamlines have more vertices than shorter ones), we re-sample the streamlines so that they all have an equal number

of 256 vertices. The latent space length was fixed at a value of 32 for all of our experiments. The latent space dimensionality choice was made after finding an appropriate dimensionality on the ‘‘Fiber Cup’’ dataset (see section 2.1.1) through a Bayesian search procedure. The value was found to be 16 for the dataset; the next power of two (32) that fit into the dimensionality reduction scheme was then used as a safety margin to account for the more complex, human brain tractography configurations.

Table 1: FINTA autoencoder structure. The encoder uses strides of size 2, and the decoder uses strides of size 1. The decoder’s upsampling stages use an upsampling factor of 2. The autoencoder uses ReLU activations throughout its convolutional layers.

Part	Type	Features	Size
input	-	-	3; 256
encoder	1D conv	32	256
	1D conv	64	128
	1D conv	128	64
	1D conv	256	32
	1D conv	512	16
	1D conv	1024	8
latent space	fully connected	32	
decoder	upsampling + 1D conv	1024	8
	upsampling + 1D conv	512	16
	upsampling + 1D conv	256	32
	upsampling + 1D conv	128	64
	upsampling + 1D conv	64	128
	upsampling + 1D conv	32	256
output	-	-	3; 256

The autoencoder was trained with an Adam optimizer (Kingma and Ba, 2015) with a mean squared-error loss. The hyper-parameters were adjusted using a Bayesian search method. The learning rate of the optimizer was fixed to a value of 6.68×10^{-4} , and weight decaying regularization with a 0.13 valued parameter.

2.1 Data

Experiments were carried out using the “Fiber Cup” synthetic data (Fillard et al., 2011; Côté et al., 2013), the ISMRM 2015 Tractography Challenge human-based synthetic dataset (Maier-Hein et al., 2015), as well as the BIL&GIN (Mazoyer et al., 2016; Chenot et al., 2019) and Human Connectome Project (HCP) (Glasser et al., 2016) in vivo human subject brain datasets.

2.1.1 “Fiber Cup” synthetic data

A synthetic dataset re-generated (using Fiberfox (Neher et al., 2014)) from the “Fiber Cup” phantom (Fillard et al., 2011; Côté et al., 2013) was used as a simplified test set mimicking the human brain white matter bundle anatomy. The “Fiber Cup” phantom contains 7 bundles, and the ground truth consists of a total of 7833 streamlines. The raw diffusion data were generated using 30 gradient directions, a diffusion gradient strength of $1000s/mm^2$, and a $3mm$ isotropic spatial resolution for a $64 \times 64 \times 3$ volume. Figure 3(a) shows the ground truth tractogram for the synthetic “Fiber Cup” phantom.

2.1.2 ISMRM 2015 Tractography Challenge human-based synthetic data

The ISMRM 2015 Tractography Challenge dataset of Maier-Hein et al. (2015) was used in a further step towards a more complex and realistic scenario. The dataset consists of a clinical-style, realistic single subject tractogram (approximately 200 000 ground truth streamlines) with 25 ground truth fiber bundles, raw diffusion-weighted data and a structural T1-weighted MRI volume generated using Fiberfox (Neher et al., 2014). The raw dMRI data were generated at a $2mm$ isotropic spatial resolution, with 32 gradient directions, and a b -value of $1000s/mm^2$. Figure 3(b) shows the sagittal left view of the ground truth tractogram for the ISMRM 2015 Tractography Challenge dataset.

2.1.3 BIL&GIN human data

A sample of 39 subjects of the BIL&GIN (*Brain Imaging of Lateralization by the Groupe d’Imagerie Neurofonctionnelle*) human brain dataset (Mazoyer et al., 2016; Chenot et al., 2019) was used to gauge the performance of FINTA. Acquisitions were done with a Philips Achieva 3 Tesla MR scanner using 21 non-colinear diffusion gradient directions and a diffusion gradient encoding strength of $1000s/mm^2$ with four averages and an isotropic spatial resolution of $2mm$. In this work, only the corpus callosum was considered, and within the bundle, homotopic streamlines lying within any of 26 pairs of gyral-based segments were considered as plausible streamlines. Additional details about the gyral-based callosal segments used for the extraction of the streamlines, shown in figure 3(c), are provided in section A.5.

2.1.4 HCP human data

A subject from the Human Connectome Project (HCP) dataset (Glasser et al., 2016) was used to perform generalization experiments. The data had been acquired using an HCP scanner equipped with high-end hardware enabling diffusion encoding gradient strengths of $100 mT/m$ at $1.25mm$ isotropic spatial resolution along 270 gradient directions. The available data had been pre-processed for head motion, eddy currents and susceptibility distortions correction.

3 Experimental

The same experiment design was implemented for the datasets. Noise-free diffusion data were used for the “Fiber Cup” and ISMRM 2015 Tractography Challenge datasets, and no denoising procedure was required for the BIL&GIN dataset. A Constrained Spherical Deconvolution (CSD) (Tournier et al., 2007) method using a spherical harmonics order of 6 was used to extract the fiber ODFs from the raw dMRI data. Local probabilistic whole brain tractography was performed using white matter seeding (single seed per voxel; $0.5mm$ step size; 30° maximum aperture angle; default values for other parameters) for the “Fiber Cup” and the ISMRM 2015 Tractography Challenge datasets to generate the sets for the autoencoder. The extracted streamlines were processed and scored to obtain the (anatomically) implausible streamline and (anatomically) plausible streamline sets. The scoring procedure for the “Fiber Cup” labelled as plausible streamline instances obeying constrained length ($20 - 200mm$) and shape features (bends $< 335^\circ$), as well as mask (streamlines contained within the known bundle masks) and connectivity (streamlines connect the known regions) criteria. The implausible set, thus, comprised streamlines exceeding the given length threshold; streamlines whose trajectories described sharp bends; streamlines that did not reach the corresponding terminal regions; and streamlines connecting regions that were not actually connected in the ground truth (see examples of such trajectories in figure 1). The ISMRM 2015 Tractography Challenge dataset was scored according to the procedure in Maier-Hein et al. (2017),

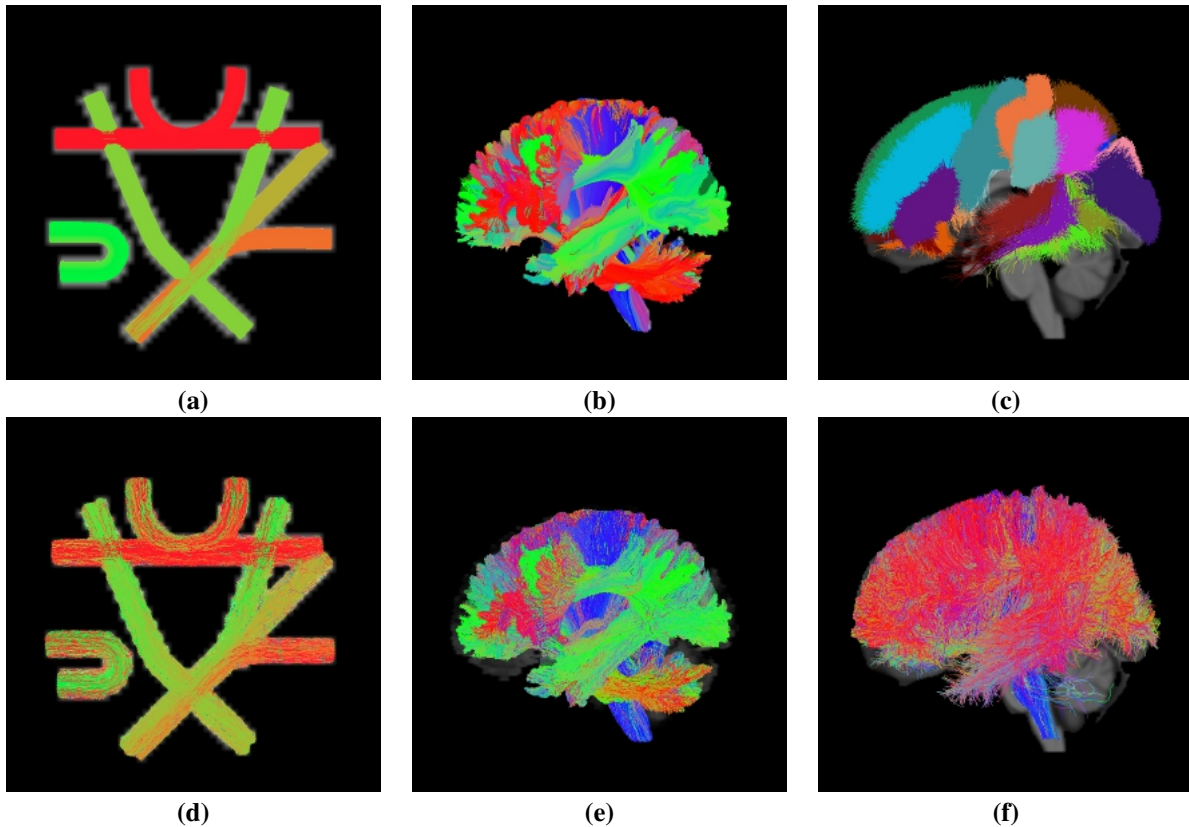


Figure 3: Tractograms corresponding to the considered datasets. Upper row: (a) ground truth tractogram from the “Fiber Cup” dataset; (b) ground truth tractogram from the ISMRM 2015 Tractography Challenge dataset; (c) sagittal left view of the gyrus-wise homotopic callosal streamline models corresponding to the BIL&GIN dataset. Lower row: test sets corresponding to the tracking performed on each of the three datasets: (d) view of the “Fiber Cup” test set containing 4963 streamlines; (e) sagittal left view of ISMRM 2015 Tractography Challenge test set containing 63615 streamlines; (f) sagittal left view of a randomly picked test subject corresponding to the BIL&GIN dataset and containing 55160 streamlines.

based on the streamline clustering strategy proposed in Garyfallidis et al. (2012). The terms *Tractometer_basic* and *Tractometer_adv* are used, respectively, to refer to the described scoring methods in section A.7. For the ISMRM 2015 Tractography Challenge dataset, the anterior commissure (CA) and the posterior commissure (CP) bundles were left out of the experiments since the tracking algorithm was unable to extract them or the extracted bundles did not meet the scoring system criteria. Hence, 23 bundles were used in the experiments.

Probabilistic tractography was performed on the BIL&GIN human brain dataset using the Particle Filtering Tracking (PFT) method (Girard et al., 2014) with white matter seeding (10 seeds per voxel; $0.5mm$ step size; 20° maximum aperture angle; default values for other parameters). The resulting tractograms were registered (Avants et al., 2011) to the MNI template common space using the MNI152 2009 standard-space T1-weighted average structural $1mm$ isotropic resolution template image (Fonov et al., 2011). For the experiments in this work, only the callosal subset of the generated tractograms was used. The plausible streamlines within this subset consisted of the streamlines within 26 gyral-based homotopic callosal segment pairs (see section A.5 for further details about their definition). The implausible set comprised, as for the “Fiber Cup” and ISMRM 2015 Tractography Challenge datasets, streamlines that did not reach the cortical surface, describing sharp bends, having a single endpoint in the callosal region, having endpoints in non-white matter tissues, or connecting both hemispheres through non-callosal segments or tissues (e.g. the brainstem). The procedure was supervised by a neuroanatomist. The procedure yielded, on average, 41838 streamlines per-subject (standard deviation: 5648) for the considered callosal dataset. The lower row in figure 3 shows the test set tractograms, including both plausible and implausible streamlines, corresponding to all three datasets.

The ratio of implausible streamlines over the resulting tractograms was in-line with the findings in previous studies: 79% for the “Fiber Cup” dataset (similar to the figures reported by Côté et al. (2013) for probabilistic methods); 46%

for the ISMRM 2015 Tractography Challenge dataset (see the valid connection ratios across submissions reported in Maier-Hein et al. (2017)); and 89% on average (standard deviation: 2%) for the callosal BIL&GIN dataset. All streamlines were resampled to contain an equal number of points (256) and their head-to-tail orientations rearranged to make the head vertices closest with respect to the origin for each dataset. The obtained streamlines were split randomly into training/validation and test sets using an 80%/20% ratio. The proposed convolutional autoencoder was trained in an unsupervised manner using the training split of both the plausible and implausible streamlines issued by the probabilistic tracking methods. A separate training was performed for each dataset. The filtering threshold was set by finding the optimal point in the receiver operating characteristic (ROC) curves evenly rewarding true positives and penalizing false positives.

3.1 Performance measure definitions

FINTA's performance was graded according to the following measures:

- *Accuracy*: proportion of correct predictions (true positives and true negatives) over the total number of streamlines.
- *Sensitivity*: proportion of relevant instances that are predicted as positives (true positives) among all positive streamlines in the data.
- *Precision*: proportion of relevant instances that are predicted as true positives among all retrieved (predicted) positive streamlines.
- *F1-score*: harmonic mean of precision and sensitivity.

3.2 Comparison to other filtering methods

The quantitative comparison was performed on the callosal BIL&GIN dataset. FINTA was compared to both anatomy-based tractography filtering methods and two variants (single-bundle and multi-bundle) of RecoBundles (Garyfallidis et al., 2018; Rheault, 2020), a state-of-the-art method applicable to partial data. The anatomy-based filtering methods were applied successively, each one filtering the previous step's output, as used in several works such as Maier-Hein et al. (2017); Sarubbo et al. (2019). The purpose of including the results of each step of this filtering pipeline (baseline filtering methods #1 to #4) was to show the successive improvement of the classification performance scores by using anatomical heuristics. The models of plausible streamlines used for the RecoBundles experiments were built according to Rheault (2020), and based on the 26 gyral-based homotopic callosal segment pairs (see section A.5). The multi-bundle RecoBundles experiment considered each gyrus-wise homotopic callosal streamline model as a separate, parameterizable entity, whereas the single-bundle RecoBundles method considered their concatenation as a single target group. This yielded a total of 6 tractography filtering methods as baselines:

1. filtering based on streamline length (named *length* within this text);
2. filtering with loop removal (*no_loops*);
3. filtering with early stops in cerebrospinal fluid (*no_end_in_csf*);
4. filtering with early stops in the white matter (*end_in_atlas*);
5. single-bundle RecoBundles (*recobundles_single*);
6. multi-bundle RecoBundles (*recobundles_multi*).

In addition to the measures mentioned in section 3.1, the balanced accuracy was computed on the predictions on the callosal BIL&GIN dataset as the mean of the sensitivity on the plausible and implausible classes. Similarly, the subscripts m and w are used (see A.4) to distinguish the unweighted (or macro) and weighted measures. The last terms account for the class imbalance between the plausible and implausible streamline classes.

The goodness of the retrieved white matter anatomy was measured in terms of the ability to preserve homotopic streamlines across the gyral segments of interest. A filtered tractogram was considered to preserve successfully the underlying white matter anatomy if at least one homotopic streamline was present in each gyral-based homotopic callosal segment pair. Thus, the anatomical goodness measure, named *valid gyrus-wise rate* (VGR), is defined for a given subject as the ratio between the number of segments (n_{VG}) containing at least one homotopic streamline after filtering and the total number of segments containing homotopic streamlines (N_{VG}) in the underlying ground truth, i.e. n_{VG}/N_{VG} . Note that N_{VG} might be different across subjects. The measure is bounded in the $[0, 1]$ range. Note also that the preserved streamlines were recovered from the positives predicted by each method using the anatomical criteria defined for each segment (see section A.5).

To estimate the task performance uniquely, a *success rate* (SR) measure was computed as the unweighted average of the classification measures and the mentioned valid gyrus-wise rate.

3.3 Generalization

In order to test the generalization ability of FINTA, we performed an additional local deterministic tracking using an ensemble tractography protocol (Takemura et al., 2016) on the ISMRM 2015 Tractography Challenge dataset and a global tracking (Reisert et al., 2011) on the Human Connectome Project subject. The ensemble local deterministic tractography was performed using white matter seeding (single seed per voxel), $\{0.1, 0.5, 1.0\}mm$ step sizes, and $\{10, 20, 30, 40\}^\circ$ maximum aperture angles. The tractograms resulting from all permutations of parameters were merged into a single tractogram file. The global tracking used 900 000 000 iterations, a particle length of $1mm$, a particle width of $0.1mm$, and a particle weight of 0.002. Diffusion data was tracked six (6) times and the resulting whole-brain tractograms were combined into a single set. In each case 500 000 streamlines were randomly sampled and scored using the same method as described in the experimental setting earlier. The scoring yielded a plausible-to-implausible ratio of 57/43% on the ensemble local deterministic tracking data and 6/94% on the global tracking data. For each set the performance measures described in section 3.1 were computed using two different filtering threshold values (i.e. the latent space minimum distance at which a streamline is considered to be implausible):

- Computed on the ISMRM 2015 Tractography Challenge dataset for the local probabilistic setting (*Generalized thr*): i.e. the reference streamline set and filtering threshold remained unaltered relative to the ISMRM 2015 Tractography Challenge local probabilistic experiment.
- Computed on a subset of each new dataset (*Specific thr*): i.e. the reference streamline sets were re-defined as new dataset-specific subsets to provide a filtering threshold value.

In either case, the autoencoder remained unaltered (i.e. was not retrained) with respect to the training done on the ISMRM 2015 Tractography Challenge local probabilistic streamlines.

3.4 Time requirements

The filtering clock time of FINTA was measured on each of the datasets. For quantitative comparison purposes, tractograms of different sizes were generated from the ISMRM 2015 Tractography Challenge data using the same procedure described earlier. In total, six (6) different tractograms (generated using a local probabilistic method) were generated containing 20 000, 40 000, 100 000, 200 000, 600 000, and 1 000 000 streamlines, respectively. The time required to filter each of these tractograms was measured for both FINTA and RecoBundles three (3) times, and the mean and standard deviation values were computed for each tractogram and method. In all cases, solely the time required for filtering was measured, excluding I/O operation time. All time tests were performed on a conventional desktop machine (Intel(R) Xeon(R) W-2133 CPU @ 3.60GHz 6 core processor; 16 GB RAM; NVIDIA GeForce GTX 1080 Ti 12 GB graphics card).

4 Results

4.1 Tractogram latent space

The autoencoder’s latent space allows to re-interpret the streamline pair-wise distance. For illustrative purposes, a number of randomly picked plausible streamlines were assigned to their corresponding bundle (or gyrus-wise homotopic callosal streamline group for BIL&GIN), and together with some implausible streamlines, they were projected into the latent space of a trained autoencoder. Figure 4 shows a view of the latent spaces corresponding to each of the three datasets. The views were generated by applying a t-SNE dimensionality reduction (van der Maaten and Hinton, 2008) to the autoencoder’s latent space vectors. The figure suggests that streamlines belonging to the same bundle (or streamline group) lie at a short distance from each other whereas implausible streamlines are uniformly distributed. Similarly, it provides evidence supporting that the autoencoder setting can be generalized and give way to multi-label streamline classification tasks such as bundling.

Furthermore, the application of mathematical tools in the latent space is straightforward. Figure 5(a) shows the interpolation principle applied to the latent space. Interpolating in the latent space allows to explore the space through a series of intermediate representations between any number of input streamlines of choice. Figure 5(b) and (c) show streamlines interpolated in the latent space between streamline pairs corresponding to two distinct sets of bundles of the “Fiber Cup” dataset. This example stresses the underlying significance of the latent space streamline pair-wise distance. Our proposed algorithm relies on the underlying hypothesis that the Euclidean distance between the latent

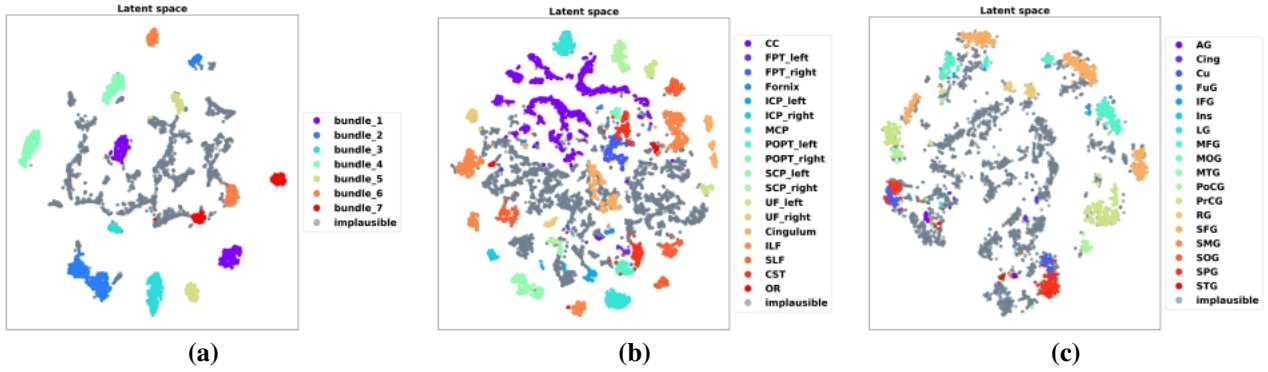


Figure 4: Latent space views using a t-SNE dimensionality reduction. (a) “Fiber Cup” data; (b) ISMRM 2015 Tractography Challenge data; (c) Callosal BIL&GIN human brain data. The number of projected streamlines was limited in each case for visualization purposes. The callosal BIL&GIN data corresponds to a randomly picked subject which was missing all streamlines in 8 of the gyral-based segments.

representation of two streamlines is a good proxy to measure their structural similarity. Thus, the fact that a linear interpolation between two latent vectors leads to a smooth transition between their reconstructed streamlines underlines that the learned manifolds are smooth and that the latent Euclidean distance is a good similarity metric.

The reconstruction of the streamlines at the output of the autoencoder involves decoding the samples in the latent space. Section A.1 shows the reconstructed tractograms corresponding to the “Fiber Cup” test set streamlines.

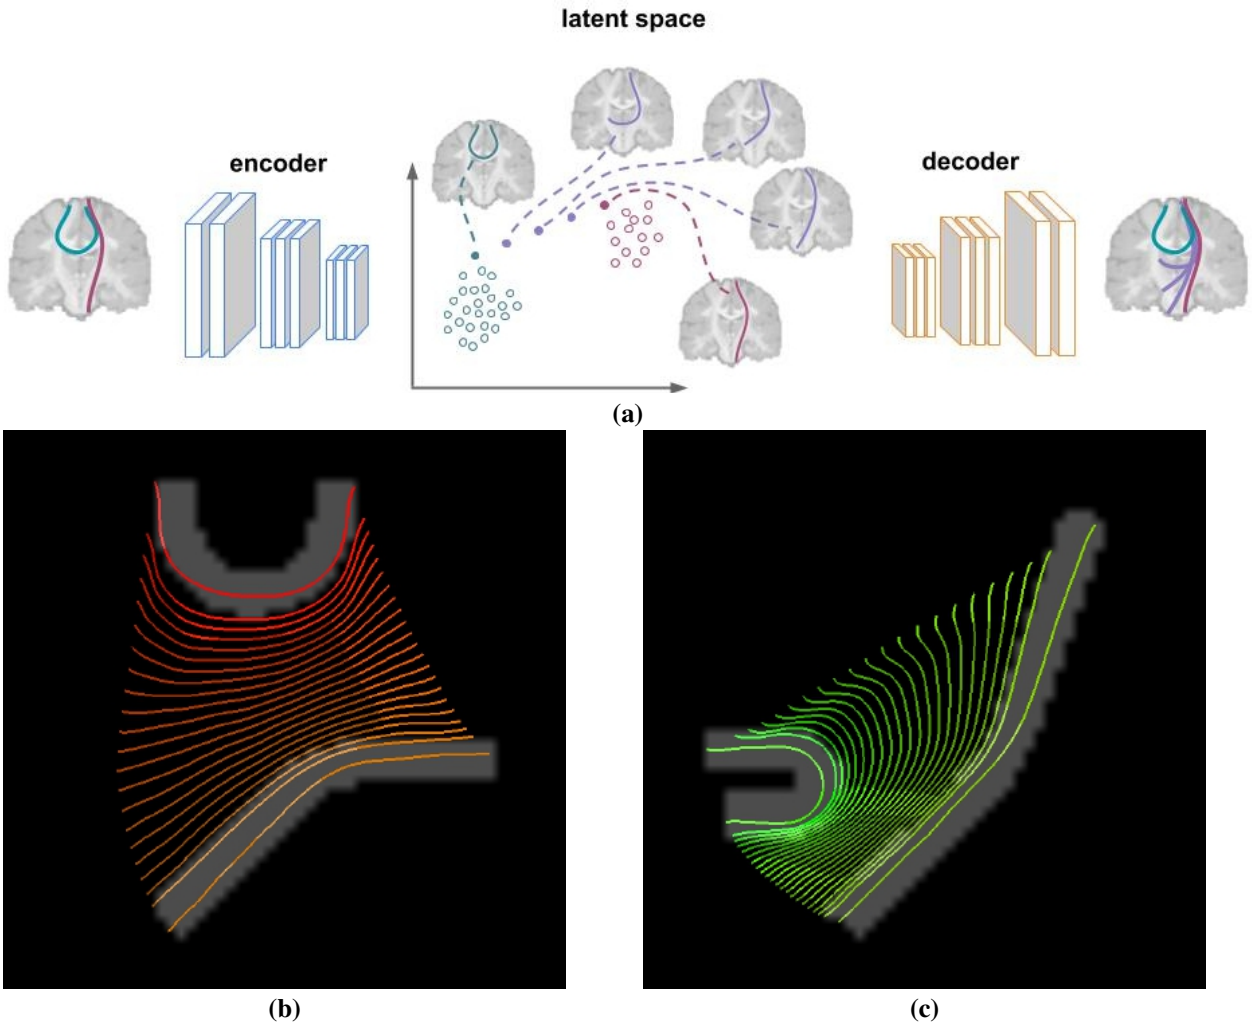


Figure 5: Latent space interpolation. (a) Schematic representation of latent space streamlines interpolation. The streamline representation in the latent space allows to choose a pair of encoded streamlines and fill the space between them through a series of interpolated streamlines. (b) and (c) “Fiber Cup” latent space streamline interpolation: streamlines interpolated between the instances at the boundaries of each series: (b) bundles 1 and 5; (c) bundles 4 and 7 (see the bundle numbering in Côté et al. (2013)). Only the structural shapes of the bundles used for the interpolation are shown for the sake of clearness. The streamlines are shown with a larger diameter with respect to the one used for the rest of the figures.

4.2 “Fiber Cup” synthetic data

Figure 6 (a) and (c) show the nearest neighbor latent space distance distribution histograms and within-class mean and variance for the “Fiber Cup” dataset. The figure illustrates the filtering principle of FINTA based on the boundary between the inliers and outliers: the plausible streamline within-class nearest neighbor distance is low (close to 0 for the “Fiber Cup”) and shows a reduced variance. As a contrast, the implausible streamlines show a much higher mean distance to their nearest neighbor streamline, as well as a higher variance. Note that the distance measure is computed considering the plausible streamlines from the train set as the reference.

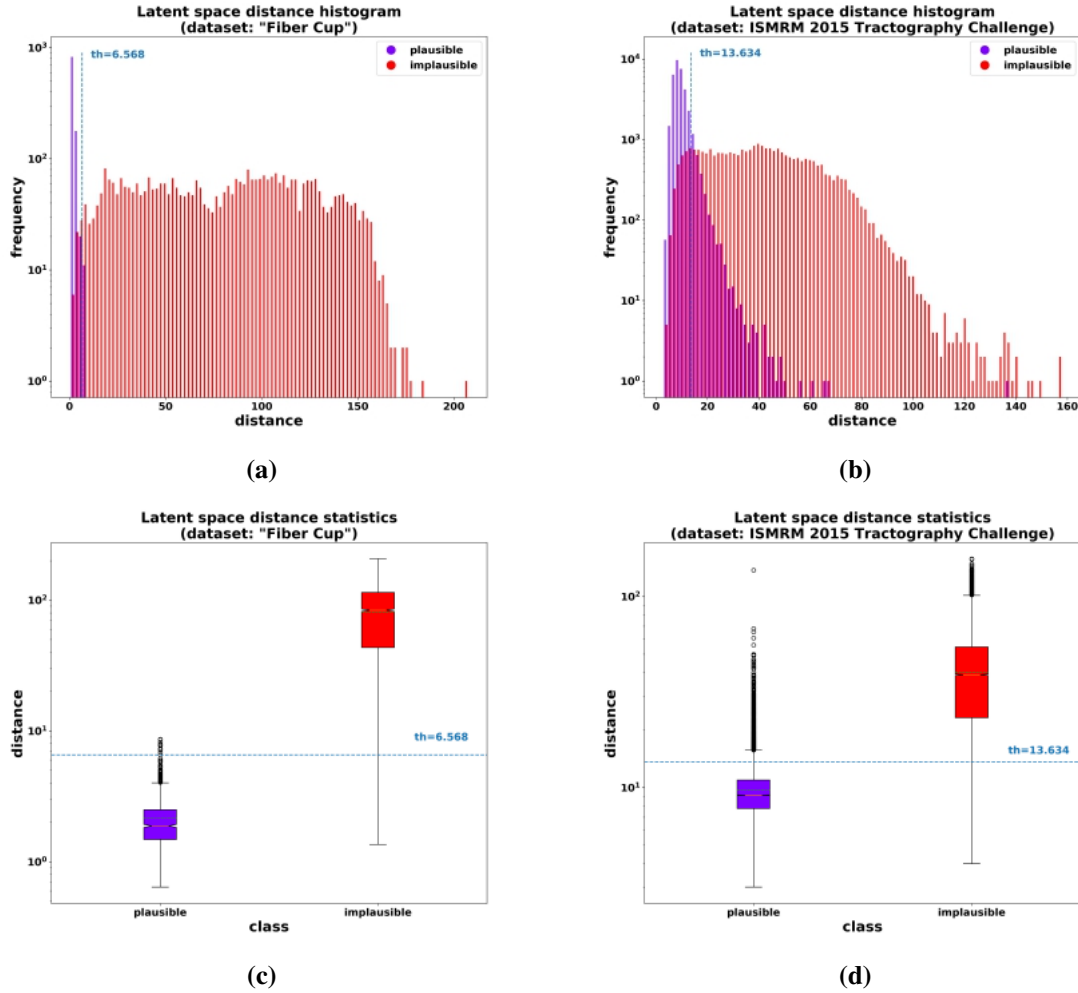


Figure 6: Nearest neighbor latent space distance distribution: histogram (a) and within-class mean and variance (c) corresponding to the “Fiber Cup” dataset test set streamlines; histogram (b) and within-class mean and variance (d) corresponding to the ISMRM 2015 Tractography Challenge. The dashed lines indicate the filtering threshold distance in each case. Note that the vertical axes are in logarithmic scale.

As quantitatively shown in table 2, the filtering distance threshold obtained by means of the autoencoder-based strategy allows to filter out the undesired implausible streamlines with a high degree of success.

Table 2: Filtering performance for the “Fiber Cup” and the ISMRM 2015 Tractography Challenge datasets.

Measure	Dataset	
	“Fiber Cup”	ISMRM 2015 Tractography Challenge
Accuracy	0.99	0.91
Sensitivity	0.99	0.91
Precision	0.97	0.91
F1-score	0.98	0.91

The predicted positive streamlines for the “Fiber Cup” are shown in figure 7. These are the streamlines considered as plausibles by FINTA from the test set, and hence include both the true positives and the false positives. Qualitatively, looking at the tractogram bundle color code information, and in agreement with the quantitative results, the implausibles seem to have been substantially filtered out, implying a decreased influence of the false positives.

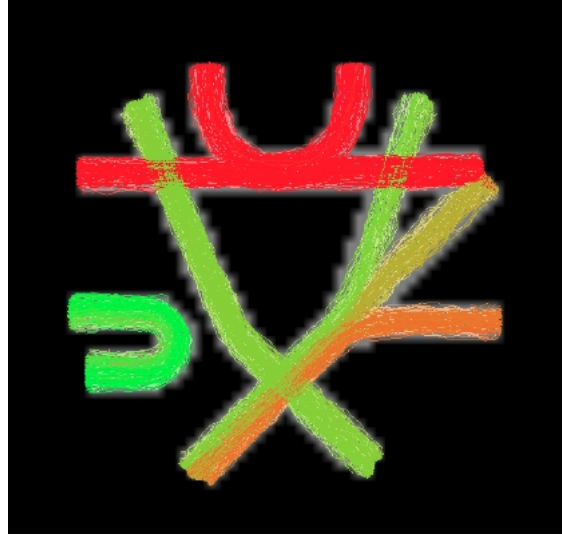


Figure 7: Predicted positives on the test set corresponding to the “Fiber Cup” dataset. Note that, in contrast to the input dataset shown in figure 3(d), the filtered tractogram shows a bundle color code information that matches better the “Fiber Cup” ground truth shown in figure 3(a), implying that implausible streamlines have been eliminated to a large extent.

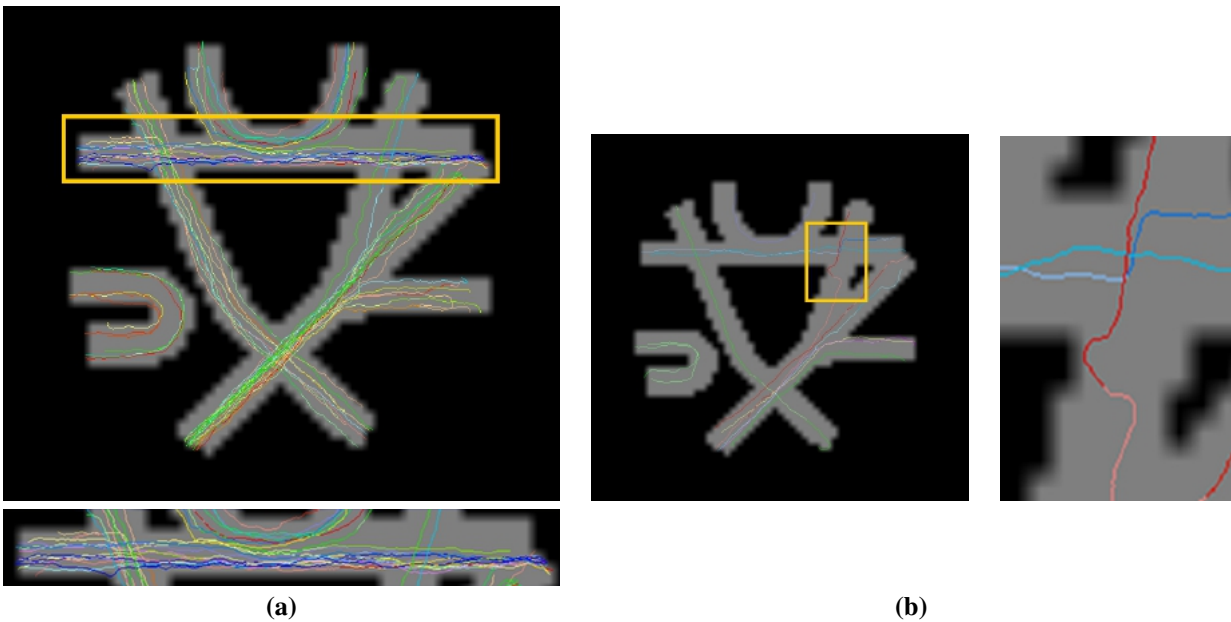


Figure 8: False positive and false negative streamlines issued by FINTA on the “Fiber Cup” dataset test set. (a) false positives: streamlines labeled as plausible by FINTA but considered implausible by the scoring method; (b) false negatives: streamlines labeled as being implausible by FINTA but considered plausible by the scoring method.

Figures 8(a) and (b) show, respectively, the false positive and false negative streamlines as detected by FINTA for the “Fiber Cup” dataset, and color-coded according to the latent space nearest neighbor distance. From the figures and the corresponding highlighted areas it can be drawn that, at least for part of the streamlines, FINTA has identified that these streamlines were misclassified in the reference set: in figure 8(a) a large part of the highlighted streamlines seem to be close to an anatomically plausible streamline within the bundle at issue. Many of these false positives are labeled so because they end 1 or 2 voxels away from the terminal regions (which would correspond to the cortical surface or gray matter border on human subject brain data). Conversely, part of the highlighted streamlines in figure 8(b) describe

trajectories that could have been identified as anatomically implausible due to erratic bends. This illustrates how close from the ground truth our results are.

Training on the “Fiber Cup” dataset took approximately 25 min. FINTA takes around 0.2s to filter the “Fiber Cup” test tractogram.

4.3 ISMRM 2015 Tractography Challenge human-based synthetic data

The nearest neighbor latent space distance distribution histograms and within-class mean and variance for the ISMRM 2015 Tractography Challenge dataset are shown in figure 6(b) and (d). As it is the case with the “Fiber Cup” dataset, although a larger degree of overlap between classes is observed, our autoencoder-based filtering strategy allows to obtain a threshold to filter out the undesired implausible streamlines with a degree of success of approximately 91% over all measures (see table 2). Here again, the latent space similarity metric shows the same behavior in both experiments: implausible streamlines have a larger distance to the reference nearest neighbor.

For qualitative visualization purposes, figure 9 shows the streamline-wise classification tractograms for some selected cases corresponding to the ISMRM 2015 Tractography Challenge dataset test set, color-coded according to their nearest neighbor distance in the latent space. In particular, plausible predictions corresponding to the the corticospinal tract (CST) and the optic radiation (OR) bundles are shown, together with an example corresponding to implausible streamlines. As it can be drawn from the scale bar, the latent space dissimilarity distances for the true positive samples (a) to (f) are below the distance threshold determined shown in figure 6(b, d) (13.634) while the distances for the implausible predictions shown are above that threshold. The filtered whole brain tractogram is shown in A.2.

Training on the ISMRM 2015 Tractography Challenge dataset lasted around 5 h. FINTA takes around 12s to filter the corresponding test set tractogram.

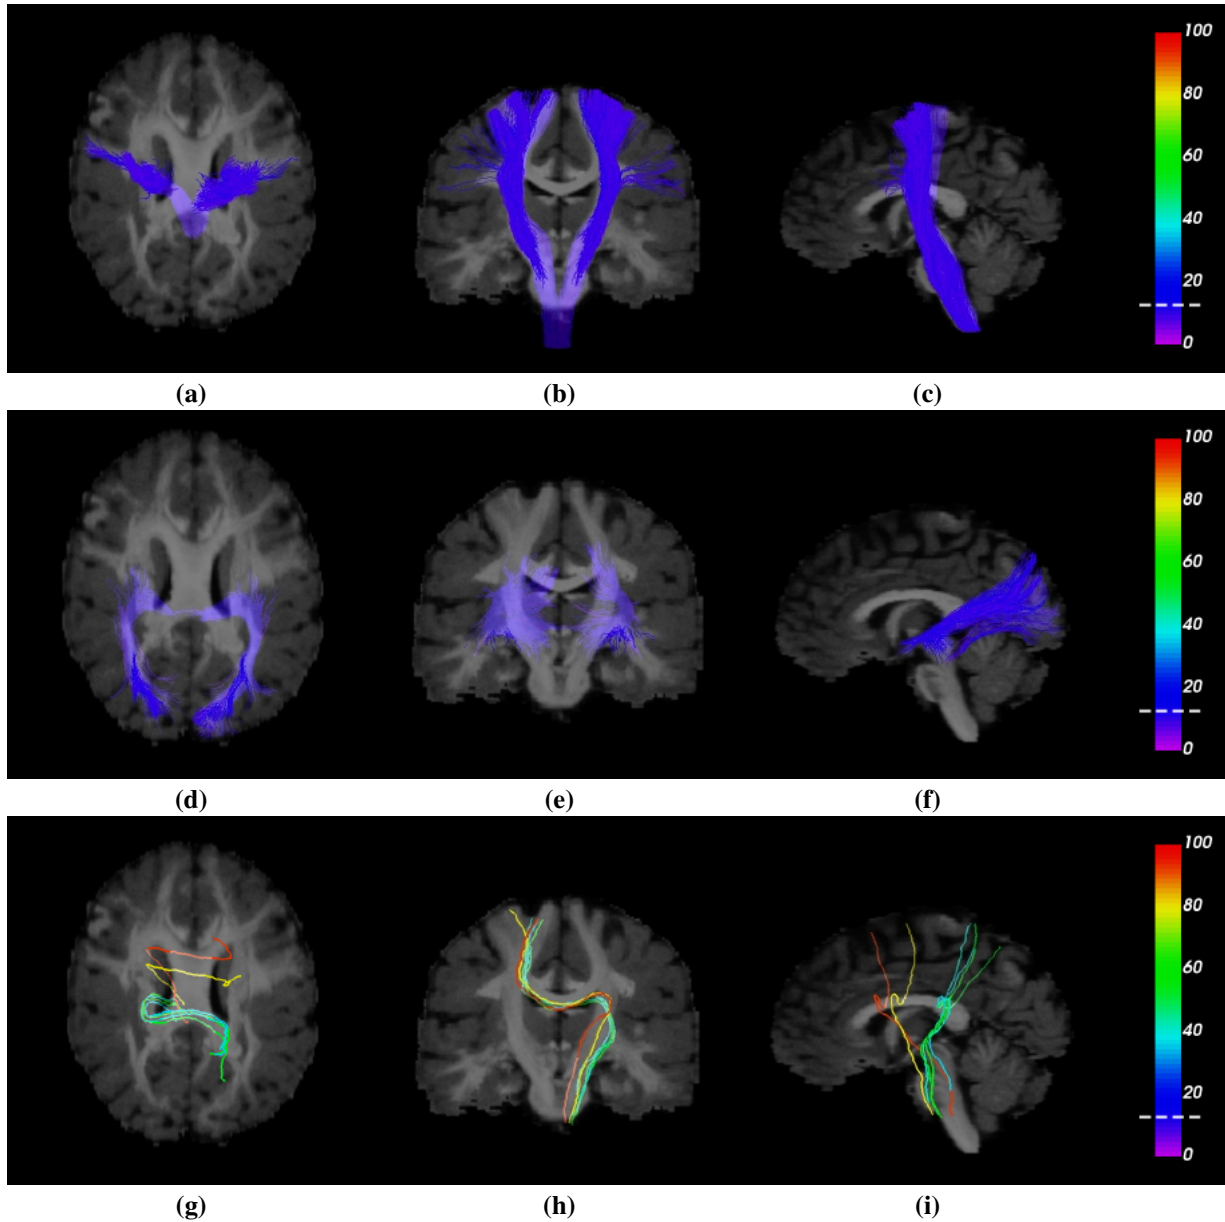


Figure 9: Streamline classification predictions corresponding to some cases of the ISMRM 2015 Tractography Challenge dataset test set: (a), (b), (c) corticospinal tract (CST); (d), (e), (f) optic radiation (OR); (g), (h), (i) example corresponding to implausible streamlines. All axial superior, coronal anterior and sagittal left views. Streamlines are color-coded according to their nearest neighbor distance in the latent space. The dashed white lines on the scale bars mark the filtering threshold. The streamlines corresponding to the implausible class are shown with a larger diameter with respect to the one used for the rest of the figures.

4.4 BIL&GIN human data

FINTA’s filtering ability on in vivo human data is consistent with the results on synthetic datasets. The quantitative results for the comparison to other filtering methods are shown in table 3, and plot in figure 10¹. As it can be drawn from these results, FINTA achieves a better classification performance than its competitors. Similarly, compared to any of the baseline methods, it is observed that FINTA shows a reduced variability in all measures consistently, with the sole exception of the valid gyrus-wise rate (VGR). Supplementary results of the comparison to the baseline filtering methods on the callosal BIL&GIN dataset are reported in the appendix sections A.3 and A.4.

The recovery of the existing anatomy, measured by the valid gyrus-wise rate, requires to be interpreted with care. Although results show that the anatomy-based baseline methods show a superior ability in this aspect, this comes at the expense of significantly lower scores for the rest of the measures (see section 5.2 for a discussion). Yet, FINTA still achieves an average valid gyrus-wise rate of 80%, performing coherently across all measures.

Table 3: Callosal BIL&GIN dataset macro results. Macro mean (standard deviation) values over test subjects. The highest mean score is marked in bold face.

	Method	Accuracy _m	Sensitivity _m	Precision _m	F1-score _m	VGR	SR _m
#1	length	0.14 (0.02)	0.51 (0.0)	0.56 (0.01)	0.13 (0.01)	0.98 (0.03)	0.46
#2	no_loops	0.19 (0.02)	0.54 (0.01)	0.56 (0.01)	0.19 (0.02)	0.97 (0.03)	0.49
#3	no_end_in_csf	0.21 (0.02)	0.55 (0.01)	0.57 (0.01)	0.21 (0.01)	0.97 (0.03)	0.5
#4	end_in_atlas	0.71 (0.02)	0.67 (0.05)	0.58 (0.02)	0.58 (0.03)	0.96 (0.04)	0.70
#5	recobundles_single	0.83 (0.02)	0.85 (0.03)	0.69 (0.02)	0.73 (0.02)	0.87 (0.1)	0.79
#6	recobundles_multi	0.82 (0.03)	0.8 (0.03)	0.67 (0.01)	0.7 (0.02)	0.96 (0.04)	0.79
#7	FINTA	0.91 (0.01)	0.91 (0.01)	0.78 (0.01)	0.83 (0.01)	0.8 (0.09)	0.84

Qualitatively, figure 11 shows that, compared to FINTA, the single-bundle RecoBundles method misses streamlines in the anterior, ventral part of the frontal lobe, and includes false positives in the occipital lobe (pointed regions).

Table 4 reports a gyral-wise analysis on the callosal BIL&GIN dataset of the sensitivity measure on the predicted positives for FINTA. The reported data correspond to the averaged values over the test subjects. Given that the callosal BIL&GIN dataset gyral-based segments allow solely to extract homotopic streamlines, only the sensitivity is reported (see the appendix section A.6 for further details). The gyral-based Hippo and TPole segments are not reported since none of the test subjects contain plausible streamlines in such segments. Similarly, the mean streamline count being rounded to the nearest integer, some segments contain 0 streamlines on average (IOG, Ins, LFOG, MFOG, and PHG) and hence are not reported. Also, note that some others have a very low (as low as 1) streamline count, and thus missing a few streamlines bears a significant impact on the results. That is the case, for example, for the FuG segment, where 2

¹The reader might notice that methods #1 to #4 in table 3 do not have a perfect score of 1.0 for the VGR measure. This is because such pipeline: (i) involves the use of steps and thresholding values that might differ from those used by the annotation tool; (ii) is affected by the effects of the registration; and (iii) uses an atlas that is slightly different from that of the annotator.

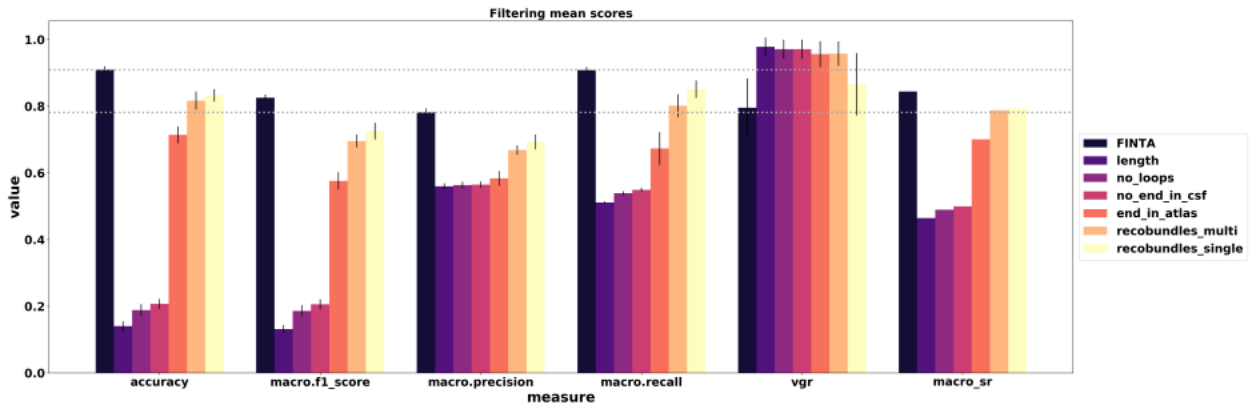


Figure 10: Classification performance measures on the callosal BIL&GIN dataset. For each of the compared filtering methods, the mean and standard deviation of the measures are presented. The horizontal dotted lines indicate FINTA’s minimum and maximum scores across the measures.

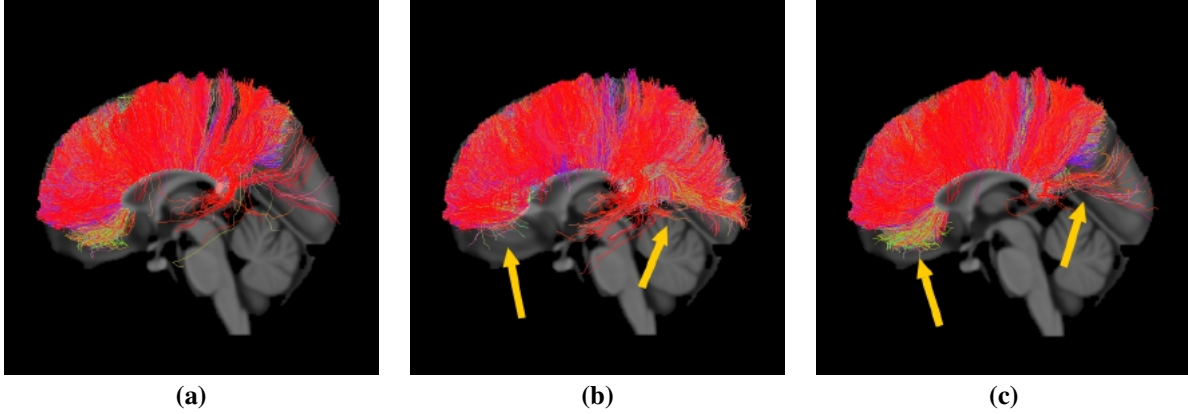


Figure 11: Predicted positives on a randomly picked test subject corresponding to the callosal BIL&GIN dataset. Predicted positives represent the passband of the filtering process. (a) reference streamlines (homotopic connections); (b) single-bundle RecoBundles; (c) FINTA. All sagittal left views. The pointed regions show areas where FINTA’s filtered tractogram is qualitatively closer to the reference tractogram compared to the single-bundle RecoBundles baseline method.

of the test subjects have both 2 streamlines, and another 2 have both 1 streamline, the remaining 4 lacking streamlines in that segment. As seen in the table, the sensitivity values for these are noticeably lower than for the segments containing a high streamline density, such as the MFG, PrCu, or SFG, all of the latter being above 86%.

Table 4: Gyral segment-wise callosal BIL&GIN dataset sensitivity of FINTA. Mean (standard deviation; [min, max]) values over test subjects for FINTA. The *Count* column contains the mean streamline count for each segment across the test subjects.

Segment	Count	Sensitivity	Segment	Count	Sensitivity
AG	14	0.14 (0.14) [0, 0.4]	PoCG	122	0.98 (0.03) [0.91, 1]
Cing	23	0.81 (0.12) [0.63, 0.96]	PrCG	177	0.98 (0.01) [0.95, 1]
Cu	6	0.08 (0.11) [0, 0.3]	PrCu	978	0.88 (0.05) [0.77, 0.94]
FuG	1	0 (0) [0, 0]	RG	87	0.99 (0.01) [0.98, 1]
IFG	12	0.5 (0.23) [0, 0.74]	SFG	3428	0.93 (0.01) [0.9, 0.95]
ITG	1	0 (0) [0, 0]	SMG	10	0.64 (0.18) [0.45, 1]
LG	4	0.2 (0.28) [0, 0.75]	SOG	23	0.52 (0.14) [0.25, 0.71]
MFG	274	0.87 (0.28) [0.77, 0.93]	SPG	142	0.67 (0.22) [0.23, 0.91]
MOG	9	0.21 (0.17) [0, 0.41]	STG	5	0.1 (0.17) [0, 0.5]
MTG	4	0 (0) [0, 0]			

FINTA was trained for 24 h on the callosal BIL&GIN dataset. FINTA takes around 21s, on average, to filter a callosal tractogram corresponding to a BIL&GIN test subject. The single-bundle RecoBundles baseline method takes, on average, 66s, in our experiments.

4.5 Generalization

Table 5 shows the generalization ability of FINTA according to the experimental setting described in section 3.3. Results reveal that the filtering performance decreases when using the generalized threshold value with respect to the results reported in table 2. Computing a dataset-specific threshold notably improves the filtering ability of the method.

4.6 Time requirements

Figure 12 compares the time required by FINTA and RecoBundles to filter tractograms of different sizes corresponding to the ISMRM 2015 Tractography Challenge data. For each method, the curves account for the mean filtering time (standard deviation values are below 10s in all cases). FINTA requires less time than RecoBundles to filter a tractogram of the same size. Also, it follows from the curves that FINTA is linear in terms of the streamline count in a tractogram, whereas RecoBundles requires a time that tends to be several order of magnitudes longer than FINTA as the streamline count increases.

Table 5: Filtering performance ability on generalization.

Measure	ISMRM 2015 Tractography Challenge		Human Connectome Project	
	Deterministic tracking		Global tracking	
	Generalized thr	Specific thr	Generalized thr	Specific thr
Accuracy	0.8	0.94	0.95	0.93
Sensitivity	0.81	0.94	0.67	0.92
Precision	0.82	0.94	0.85	0.72
F1-score	0.8	0.94	0.73	0.78

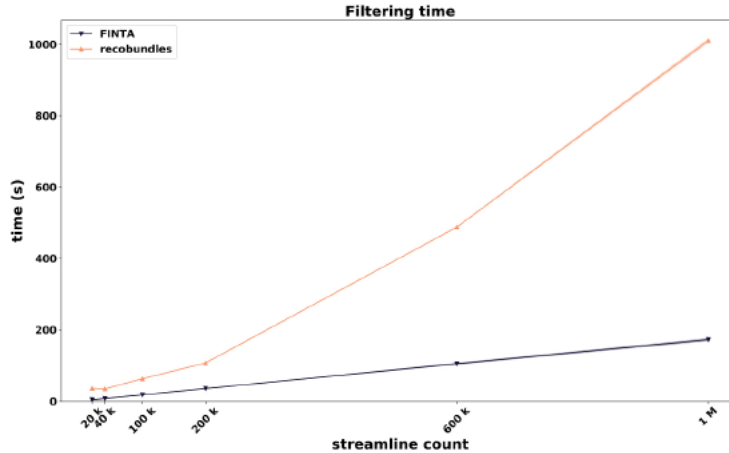


Figure 12: Filtering time requirements for FINTA and RecoBundles. The curves demonstrate that FINTA is fast, having linear time complexity with the streamline count, while RecoBundles requires a significantly longer time to filter a given tractogram. Note that due to the vertical scale and reduced standard deviation values, the latter are hardly noticeable around the mean value. Similarly, note that the streamline counts are expressed with SI prefixes and engineering notation.

5 Discussion

The results across the three different datasets show that FINTA is a state-of-the-art alternative for tractography filtering. The latent space views suggest that streamlines belonging to the same class lie close to each other in the latent space, supporting the hypothesis that the autoencoder successfully clusters such streamlines in an unsupervised manner. The proposed filtering method achieves good performance statistics on both synthetic and in vivo human brain data, and consistently ranks higher in the considered measures compared to the baseline methods.

5.1 Performance on synthetic datasets

Although less convoluted than a tractogram derived from a human brain dataset, the “Fiber Cup” dataset still presents some challenges to tractography filtering methods, and hence it offers a good test bed to show the behavior of a given method. Our results show that FINTA provides almost perfect performances on it. Additionally, the dataset allows to gain further insight into the behavior of a given filtering method. Our findings illustrated in figure 8 suggest that: (1) false positives and false negatives generated by FINTA are very close to true positives and true negatives; and (2) in some cases, FINTA can reveal examples of streamlines that were misclassified by the scoring method. Such inconsistencies with respect to the reference highlight the ability our proposed method has to exploit the learned features from the streamlines. At the same time, it is a demonstration of some potential limitations of conventional filtering approaches, which employ voxel-based or anatomical criteria, and pair-wise streamline dissimilarity measures that parallel those employed by existing scoring methods.

The slight decrease in performance registered by FINTA on the ISMRM 2015 Tractography Challenge (also seen in the human in vivo callosal BIL&GIN dataset) compared to the “Fiber Cup” is explained by the complexity of the former, where a larger variety of bundle and streamline configurations are present. Nonetheless, with a mean accuracy of 91%, our method still gets an excellent score. Together, the results show that there is a good agreement between

the performance on the ISMRM 2015 Tractography Challenge and callosal BIL&GIN datasets, suggesting that these results are fully descriptive of the power of the method on human in vivo tractogram settings.

5.2 Performance on human data

Although the anatomy-based tractography filtering methods (methods #1 to #4) were expected to record lower scores than their competitors, they are still an important part of routine tractography filtering pipelines prior to downstream tasks, such as connectivity studies (Jeurissen et al., 2019; Yeh et al., 2020). No additional steps, such as obtaining the gyral-based homotopic streamlines, were added to avoid reproducing the scoring method providing the reference.

Classical filtering methods obtain a better valid gyrus-wise rate than FINTA at the cost of including an excessive number of false positives among the kept streamlines. Penalizing more severely missing groups would have had some impact on the reported success rates without harm to the rest of the filtering performance measures. Similarly, the superior ability of the RecoBundles bundle recognition methods to recover the existing anatomy may be explained by the use of gyrus-wise homotopic callosal streamline models specifically built for that purpose. The single-bundle RecoBundles obtained slightly better scores than the multi-bundle version except for the valid gyrus-wise rate measure. This might be explained by a greed of the multi-bundle version to assimilate streamlines to every model of the corpus callosum gyral-based segment pairs, leading it to an excessive number of false positives. Compared to FINTA, the improved valid gyrus-wise rate of the multi-bundle RecoBundles baseline comes at the cost of potentially recognizing bundles that are not contained in the data (i.e. invalid bundles).

Note that the valid gyrus-wise rate does not account for the potential invalid (i.e. not existing in the underlying anatomy) groups that some of the methods may recover. It must be noted that the 26 pairs of gyral-based callosal segments (see the appendix section A.5) are representative of the entire BIL&GIN population, but subjects may be missing streamlines on an entire segment. Furthermore, the streamline count on some segments may be as low as a single streamline for some subjects, which may impact considerably the valid gyrus-wise rate measure.

In this work, true positives and false positives were evenly weighted to find the optimal filtering threshold. We found that maximizing only the accuracy penalized excessively the least numerous class in a class imbalance situation. Likewise, the results in figure 10 show that this strategy leads to a high specificity, and allows to keep false positives at low rates in FINTA. Meanwhile, the rest of methods suffer from an excessive number of false positives. Note that this effect is observed consistently in the measures where false positives are involved (accuracy, F1-score, and precision). This finding is relevant for connectome studies, where keeping a good specificity is essential (Zalesky et al., 2016).

The reduced variability in the measures shown by FINTA suggests that the method is less prone to be influenced by the local, potentially noisy streamline trajectories. Conventional methods' ability to filter streamlines is determined by local decisions based on either local (e.g. curvature, regions of interest, endpoint location) or non-local (e.g. length) features, and thus a single streamline coordinate not complying with the given criterion (e.g. falling outside a structural mask) is enough to filter out the streamline. This reaffirms the ability of FINTA to encode those features of the input data which are most relevant for an accurate reconstruction.

The corpus callosum spans a large part of the brain volume in the anterior-posterior axis across both hemispheres; it is one of the largest fiber systems in the white matter, both in terms of the absolute number of streamlines and the occupied brain volume. Compared to some other white matter systems, it presents a less compact configuration. As a consequence, it is challenging to precisely delineate the corpus callosum homotopies in the reference set (De Benedictis et al., 2016). Similarly, as previously noted, the 26 gyral segment pairs used in this work differ greatly in the streamline density they hold, and thus bear an impact on the gyrus-wise sensitivity scores shown in table 4. Yet, FINTA records high classification performance scores, and shows a reduced variability in the valid gyrus-wise rate measure.

The performance assessment on human data was restricted by the nature of the callosal BIL&GIN dataset and its corresponding ground truth (see section A.5 for the relevant aspects). This constrained the baseline filtering methods, excluding some state-of-the-art techniques such as COMMIT2 or SIFT2. However, the baselines used are commonly found in current white matter analysis practice.

5.3 General considerations

As opposed to the conventional filtering methods that assume that the estimated streamline population at every voxel should be proportionally supported by the diffusion data (which is the case of COMMIT/COMMIT2, LiFE, or SIFT/SIFT2), our autoencoder works on tracking data only. This makes it immune against well-known domain adaptation issues (Perone et al., 2019). Furthermore, compared to a quadratic or super-quadratic complexity of some methods, FINTA is linear in terms of the streamline count at test time as demonstrated by the results in figure 12.

Compared to other deep learning approaches that may be specifically oriented to classification tasks (e.g. regular classification convolutional neural networks), FINTA offers the benefit of using an unsupervised learning approach. Thus, it does not depend on the number of classes in the input data as the output of the network does not look for maximizing the probability of a given class among the possible ones. As such, the method is better suited to the reality of tractography, where there is a limited knowledge about the ground truth (Yeh et al., 2016; Sotiropoulos and Zalesky, 2019), and where the analysis lends itself to different organizational levels, and thus different classification degrees. As such, our method naturally adapts to new sets of classes without the need of having to retrain the network.

A potential side-benefit of the proposed autoencoder approach is that the reconstructed streamlines describe a locally smooth trajectory. Although not explicitly used or exploited in the present work, succeeding pipeline or visualization processes may benefit from it, providing a less complex or more realistic long-range apparent fiber trajectory representation.

Downstream tractography tasks benefit from filtering strategies applied at earlier stages when done in a coherent manner. One of such tasks is the structural connectivity analysis (or connectomics). Any proposed filtering strategy to reject implausible streamlines is expected to preserve the existing connectivity. Estimating the accuracy of a connectome yielded by a tractography and filtering pipeline can only be done faithfully on synthetic data, and yet many confounding factors (e.g. parcellation and discretization effects), including those from upstream steps (such as the seeding method), are involved. Conceding such limitations to an analysis, we provide evidence in section A.7 that FINTA is able to regenerate the connectome to a high degree of agreement with the ground truth on the synthetic datasets used in this work.

5.4 Generalization, streamline diversity and balance effects

FINTA’s reported global performances stress the fact that the proposed framework is robust across a varying spectrum of settings. The proposed autoencoder-based method assumes that, provided that the implausible streamlines have overall distinctive features from the plausible ones, they will be cast to different regions in the latent space, and hence the filtering framework will be able to separate them, regardless of their location and arrangement in the native space or within-fascicle balance. Similarly, as highlighted in section 4.1, the autoencoder shows the ability to disentangle a varying number of bundles or streamline groups.

The results in section 4.5 show that FINTA’s optimal filtering threshold computation framework obtains best results when tailoring its value to the dataset at issue. Otherwise, a trade-off value can be proposed by finding a threshold on data from a variety of tracking settings (e.g. probabilistic, deterministic, global, etc.) and datasets. It is worthwhile noting that the autoencoder was not retrained across datasets for the generalization experiments. These results show how immune against domain adaptation problems FINTA is.

The datasets used in this work varied greatly in the overall plausible *versus* implausible streamline density: the “Fiber Cup” dataset had an overall plausible/implausible ratio of 21/79%, whereas the ISMRM 2015 Tractography Challenge settings had 54/46% (probabilistic) and 57/43% (deterministic) respectively. Meanwhile, the callosal BIL&GIN dataset had an 11/89% ratio, and the Human Connectome Project global tracking data showed a 6/94% ratio.

It can be hypothesized that the unsupervised optimization training phase and the sensitivity *versus* specificity trade-off used to set the global filtering threshold might have a larger impact on the performance on the least numerous class. Particularly, this effect is visible in the global tracking setting of the Human Connectome Project dataset (table 5), where the performance might be influenced by the misclassifications in the plausible streamlines. Nevertheless, even in the unfavorable conditions of the callosal BIL&GIN dataset, the weighted performance scores reported in tables 6 and 8 for the callosal BIL&GIN dataset show that FINTA successfully filters implausible streamlines when accounting for the class imbalance.

5.5 How good is the ground truth in tractography?

Inter-subject variability, and the lack of a “gold standard” to reliably determine bundles on in vivo human brain data (Rheault et al., 2020b) add a degree of uncertainty to tractography and downstream tasks. Our work demonstrates the ability of a 1D convolutional autoencoder to reveal uncertainties inherent to current state-of-the-art tractography filtering methods. Our filtering method is a powerful tool to uncover disparities and biases in current data scoring practices that can otherwise go unnoticed. The insight gained from such a model can be useful to build upon existing knowledge in tractography. Re-defining some of the current tractography filtering approaches may be essential when dealing with human data, where consensus on the best predictors of the underlying white matter anatomy is still hard to achieve. Together, this also highlights the need of a larger set of validated, shared data to converge towards a more quantitative tractography approach.

5.6 Future work

The proposed architecture requires all streamlines to have an equal number of points. A limitation of this work is that this number is larger than the minimal number of points required to express a streamline in regular tractography pipelines (Presseau et al., 2015). Further work is needed to determine a more efficient characterization of a streamline in terms of a representation learning framework. Following from this, the proposed framework would also benefit from investigating alternatives to avoid the streamline flipping pre-processing step. Note that we did not study the effect of the number of parameters of the network, and only focused on the optimization of the reconstruction accuracy. Additional investigations are required to quantify the impact of the network size on the reconstructed streamlines and filtering performance.

We did not investigate the streamline or tractography features the autoencoder-based approach focuses on. Further work is required to gain more insight on the data aspects that the model is able to capture or tends to discard. This is still an active area of research in deep learning, where many efforts are being put to explain how a model behaves in terms of explainable concepts related to the data and its features. This would allow, among others, to estimate reliably how a given ground truth labeling constraint may affect FINTA’s performance.

Additional work includes confirming the robustness of the method against particular anatomical conditions. This involves, for example, measuring the ability of our method to generalize on infant brain tractography data or on brain data presenting a condition, such as a tumor, potentially affecting the tractogram distribution, without having to re-train the system. These experiments so could lead to further classification categories or to the need of a conditional mechanism attached to the autoencoder.

The current findings on the effects of the proposed filtering method in downstream tractography tasks have a limited scope. FINTA records low false positive rates, which is paramount to connectome analyses (Zalesky et al., 2016), and successfully reconstructs the connectome on the synthetic datasets. However, misses in sparse streamline fascicles penalize its performance on the callosal BIL&GIN human brain *in vivo* dataset. Further work would be required to improve FINTA’s ability to preserve the white matter fascicle support in complex tractography settings by possibly incorporating additional a priori information into the autoencoder learning framework.

6 Conclusions

We have proposed an unsupervised, deep learning-based dimensionality reduction method for dMRI tractography filtering. We have dubbed our method FINTA, *Filtering in Tractography using Autoencoders*. We have shown that, working with a simple architecture, FINTA is able to robustly learn the structure of streamlines in tractography, and have applied such a framework to discriminate between an input tractogram’s “positive” and “negative” streamlines. These categories may target plausible and implausible streamlines or a subset of bundles *versus* the rest of the tractogram. For the plausible *versus* implausible setting, once the neural network has been trained on unlabeled tractograms, simply labeling a subset of the streamlines of interest to obtain the partitioning threshold suffices to use this value to filter new tractograms using FINTA. We have demonstrated that the method allows to successfully and reliably filter tractograms in this context with accuracies above the 90% bar, obtaining improved scores compared to state-of-the-art filtering methods. To this end, we have proposed a means for redefining the notion of inter-streamline distance which can be similarly used for any streamline filtering tasks without the need of having to re-train the network. Additionally, we have shown that the proposed convolutional neural network autoencoder generalizes effectively under a number of varying tractography settings, such as partial tractograms or different tracking methods across datasets.

The proposed FINTA framework offers a fast, robust, data-driven approach for tractography filtering. Its unique characteristics unfold the potential of autoencoders to improve the accuracy and reliability of downstream tractometry and connectomics derivatives. Thus, it may assist neuroanatomists in better describing the white matter anatomy using tractography, and hence bring a significant positive impact to research in neuroscience. FINTA is expected to be integrated into a specialized diffusion MRI analysis open source toolkit.

Acknowledgments

This work has been partially supported by the Centre d’Imagerie Médicale de l’Université de Sherbrooke (CIMUS); the Axe d’Imagerie Médicale (AIM) of the Centre de Recherche du CHUS (CRCHUS); the Réseau de Bio-Imagerie du Québec (RBIQ)/Quebec Bio-imaging Network (QBIN) (FRSQ - Réseaux de recherche thématiques File: 35450); and the Samuel de Champlain 2019-2020 Program of the Conseil Franco-Québécois de Coopération Universitaire (CFQCU). This research was enabled in part by support provided by Calcul Québec (www.calculquebec.ca/en/) and Compute Canada (www.computecanada.ca). We also thank the research chair in Neuroinformatics of the Université

de Sherbrooke. Data were provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

References

- P. Astolfi, R. Verhagen, L. Petit, E. Olivetti, J. Masci, D. Boscaini, and P. Avesani. Tractogram filtering of anatomically non-plausible fibers with geometric deep learning. *arXiv*, 03 2020.
- B. B. Avants, N. J. Tustison, G. Song, P. A. Cook, A. Klein, and J. C. Gee. A reproducible evaluation of ants similarity metric performance in brain image registration. *NeuroImage*, 54(3):2033–2044, 02 2011.
- G. Bertò, D. Bullock, P. Astolfi, S. Hayashi, L. Zigiotta, L. Annicchiarico, F. Corsini, A. D. Benedictis, S. Sarubbo, F. Pestilli, P. Avesani, and E. Olivetti. Classifyber, a robust streamline-based linear classifier for white matter bundle segmentation. *bioRxiv*, 02 2020.
- A. Brun, H. Knutsson, H.-J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber traces using normalized cuts. 3216: 368–375, 09 2004.
- D. Bullock, L. K. Hiromasa Takemura, Cesar F. Caiafa, B. McPherson, B. Caron, and F. Pestilli. Associative white matter connecting the dorsal and ventral posterior human cortex. *Brain structure and Function*, 224(8):2631–2660, 11 2019.
- M. Catani, R. J. Howard, S. Pajevic, and D. K. Jones. Virtual in vivo interactive dissection of white matter fasciculi in the human brain. *NeuroImage*, 17(1):77–94, 09 2002.
- Y.-P. Chao, K.-H. Cho, C.-H. Yeh, K.-H. Chou, J.-H. Chen, and C.-P. Lin. Probabilistic topography of human corpus callosum using cytoarchitectural parcellation and high angular resolution diffusion imaging tractography. *Human Brain Mapping*, 30(10):3172–3187, 02 2009.
- Q. Chenot, N. Tzourio-Mazoyer, F. Rheault, M. Descoteaux, F. Crivello, L. Zago, E. Mellet, G. Jobard, M. Joliot, B. Mazoyer, and L. Petit. A population-based atlas of the human pyramidal tract in 410 healthy participants. *Brain Structure and Function*, 224(2):599–612, 03 2019.
- M.-A. Côté, G. Girard, A. Boré, E. Garyfallidis, J.-C. Houde, and M. Descoteaux. Tractometer: Towards validation of tractography pipelines. *Medical Image Analysis*, 17(7):844–857, 10 2013. Special Issue on the 2012 Conference on Medical Image Computing and Computer Assisted Intervention.
- A. Daducci, A. D. Palù, A. Lemkaddem, and J.-P. Thiran. COMMIT: Convex optimization modeling for microstructure informed tractography. *IEEE Transactions on Medical Imaging*, 34(1):246–257, 01 2015.
- M. Dayan, V.-M. Katsageorgiou, L. Doderò, V. Murino, and D. Sona. Unsupervised detection of white matter fiber bundles with stochastic neural networks. In *25th IEEE International Conference on Image Processing (ICIP)*, pages 3513–3517, Athens, Greece, 10 2018.
- A. De Benedictis, L. Petit, M. Descoteaux, C. E. Marras, M. Barbareschi, F. Corsini, M. Dallabona, F. Chioffi, and S. Sarubbo. New insights in the homotopic and heterotopic connectivity of the frontal portion of the human corpus callosum revealed by microdissection and diffusion tractography. *Human Brain Mapping*, 37(12):4718–4735, 12 2016.
- O. A. S. de Lucena. Deep learning for brain analysis in MR imaging. Master’s thesis, Universidade Estadual de Campinas, 09 2018.
- S. Deslauriers-Gauthier, M. Zucchelli, M. Frigo, and R. Deriche. A unified framework for multimodal structure–function mapping based on eigenmodes. *Medical Image Analysis*, 66:101799, 12 2020.
- P. Fillard, M. Descoteaux, A. Goh, S. Gouttard, B. Jeurissen, J. Malcolm, A. Ramirez-Manzanares, M. Reisert, K. Sakaie, F. Tensaouti, T. Yo, J.-F. Mangin, and C. Poupon. Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom. *NeuroImage*, 56(1):220–234, 05 2011.
- V. S. Fonov, A. C. Evans, K. Botteron, C. R. Almlí, R. C. McKinstry, and D. L. Collins. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage*, 54(1):313–327, 2011.
- E. Garyfallidis, M. Brett, M. M. Correia, G. B. Williams, and I. Nimmo-Smith. QuickBundles, a method for tractography simplification. *Frontiers in Neuroscience*, 6:175, 12 2012.
- E. Garyfallidis, M.-A. Côté, F. Rheault, J. Sidhu, J. Hau, L. Petit, D. Fortin, S. Cunanne, and M. Descoteaux. Recognition of white matter bundles using local and global streamline-based registration and clustering. *NeuroImage*, 170:283–295, 04 2018. Segmenting the Brain.

- G. Girard, K. Whittingstall, R. Deriche, and M. Descoteaux. Towards quantitative connectivity analysis: reducing tractography biases. *NeuroImage*, 98:266–278, 09 2014.
- M. F. Glasser, S. M. Smith, D. S. Marcus, J. L. R. Andersson, E. J. Auerbach, T. E. J. Behrens, T. S. Coalson, M. P. Harms, M. Jenkinson, S. Moeller, E. C. Robinson, S. N. Sotiropoulos, J. Xu, E. Yacoub, K. Ugurbil, and D. C. V. Essen. The human connectome project’s neuroimaging approach. *Nature Neuroscience*, 19(9):1175–1187, 09 2016.
- P. Guevara, C. Poupon, D. Rivière, Y. Cointepas, M. Descoteaux, B. Thirion, and J.-F. Mangin. Robust clustering of massive tractography datasets. *NeuroImage*, 54(3):1975–1993, 02 2011.
- V. Gupta, S. I. Thomopoulos, F. M. Rashid, and P. M. Thompson. FiberNET: An ensemble deep learning framework for clustering white matter fibers. In M. Descoteaux, L. Maier-Hein, A. Franz, P. Jannin, L. D. Collins, and S. Duchesne, editors, *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 548–555, Québec, Canada, 09 2017. Springer International Publishing.
- V. Gupta, S. I. Thomopoulos, C. K. Corbin, F. M. Rashid, and P. M. Thompson. FIBERNET 2.0: An automatic neural network based tool for clustering white matter fibers in the brain. In *IEEE 15th International Symposium on Biomedical Imaging (ISBI)*, pages 708–711, Washington, DC, USA, 04 2018.
- G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786): 504–507, 07 28, 2006.
- S. Hofer and J. Frahm. Topography of the human corpus callosum revisited - comprehensive fiber tractography using diffusion tensor magnetic resonance imaging. *NeuroImage*, 32(3):989–994, 2006.
- B. Jeurissen, M. Descoteaux, S. Mori, and A. Leemans. Diffusion MRI fiber tractography of the brain. *NMR in Biomedicine*, 32(4):e3785, 04 2019. e3785 NBM-17-0045.R2.
- Y. Jin, Y. Shi, L. Zhan, B. A. Gutman, G. I. de Zubicaray, K. L. McMahon, M. J. Wright, A. W. Toga, and P. M. Thompson. Automatic clustering of white matter fibers in brain diffusion MRI with an application to genetics. *NeuroImage*, 100:75–90, 2014.
- D. Jörgens, M. Descoteaux, and R. Moreno. Challenges for tractogram filtering. In E. Özarlan, T. Schultz, E. Zhang, and A. Fuster, editors, *Anisotropy Across Fields and Scales*, Mathematics and Visualization. Springer, 2020. in press.
- D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *3rd International International Conference for Learning Representations (ICLR)*, San Diego, CA, USA, 05 2015.
- K. Kumar, K. Siddiqi, and C. Desrosiers. White matter fiber analysis using kernel dictionary learning and sparsity priors. *Pattern Recognition*, 95:83–95, 11 2019.
- P. D. N. Lam, G. Belhomme, J. Ferrall, B. Patterson, M. Styner, and J. C. Prieto. TRAFIC: Fiber tract classification using deep learning. In *Proceedings of The International Society for Optical Engineering (SPIE)*, volume 10574, page 1057412. The International society for optics and photonics (SPIE), 03 2018.
- B. Li, M. de Groot, R. M. E. Steketee, R. Meijboom, M. Smits, M. W. Vernooij, M. A. Ikram, J. Liu, W. J. Niessen, and E. E. Bron. Neuro4Neuro: A neural network approach for neural tract segmentation using large-scale population-based diffusion imaging. *NeuroImage*, 218:116993, 09 2020.
- H. Li, Z. Xue, L. Guo, T. Liu, J. Hunter, and S. T. Wong. A hybrid approach to automatic clustering of white matter fibers. *NeuroImage*, 49(2):1249–1258, 2010.
- M. Maddah, W. E. L. Grimson, S. K. Warfield, and W. M. Wells. A unified framework for clustering and quantitative analysis of white matter fiber tracts. *Medical Image Analysis*, 12(2):191–202, 2008.
- K. H. Maier-Hein, P. Neher, J.-C. Houde, E. Caruyer, A. Daducci, T. Dyrby, B. Stieltjes, and M. Descoteaux. Tractography challenge ISMRM 2015 data, 05 2015. URL <https://doi.org/10.5281/zenodo.572345>.
- K. H. Maier-Hein, P. F. Neher, J.-C. Houde, M.-A. Côté, E. Garyfallidis, J. Zhong, M. Chamberland, F.-C. Yeh, Y. C. Lin, Q. Ji, W. E. Reddick, J. O. Glass, D. Q. Chen, Y. Feng, C. Gao, Y. Wu, J. Ma, H. Renjie, Q. Li, C.-F. Westin, S. Deslauriers-Gauthier, J. O. O. González, M. Paquette, S. St-Jean, G. Girard, F. M. Rheault, J. Sidhu, C. M. W. Tax, F. Guo, H. Y. Mesri, S. Dávid, M. Froeling, A. M. Heemskerk, A. Leemans, A. Boré, B. Pinsard, C. Bedetti, M. Desrosiers, S. Brambati, J. Doyon, A. Sarica, R. Vasta, A. Cerasa, A. Quattrone, J. Yeatman, A. R. Khan, W. Hodges, S. Alexander, D. Romascano, M. Baraković, A. Auría, Ó. Esteban, A. Lemkaddem, J.-P. Thiran, H. E. Cetinul, B. L. Odry, B. Mailhe, M. Nadar, F. Pizzagalli, G. Prasad, J. E. Villalón-Reina, J. Galvis, P. M. Thompson, F. D. S. Requejo, P. L. Laguna, L. M. Lacerda, R. Barrett, F. Dell’Acqua, M. Catani, L. Petit, E. Caruyer, A. Daducci, T. B. Dyrby, T. Holland-Letz, C. C. Hilgetag, B. Stieltjes, and M. Descoteaux. The challenge of mapping the human connectome based on diffusion tractography. *Nature Communications*, 8(1349):1–13, 11 2017.

- B. Mazoyer, E. Mellet, G. Perchey, L. Zago, F. Crivello, G. Jobard, N. Delcroix, M. Vigneau, G. Leroux, L. Petit, M. Joliot, and N. Tzourio-Mazoyer. BIL&GIN: A neuroimaging, cognitive, behavioral, and genetic database for the study of human brain lateralization. *NeuroImage*, 124:1225–1231, 01 2016. Sharing the wealth: Brain Imaging Repositories in 2015.
- P. F. Neher, F. B. Laun, B. Stieltjes, and K. H. Maier-Hein. Fiberfox: Facilitating the creation of realistic white matter software phantoms. *Magnetic Resonance in Medicine*, 72(5):1460–1470, 2014.
- L. J. O’Donnell and C.-F. Westin. Automatic tractography segmentation using a high-dimensional white matter atlas. *IEEE Transactions on Medical Imaging*, 26(11):1562–1575, 11 2007.
- L. J. O’Donnell, M. R. Kubicki, M. E. Shenton, M. H. Dreusicke, W. E. L. Grimson, and C.-F. Westin. A method for clustering white matter fiber tracts. *American Journal of Neuroradiology*, 27(5):1032–1036, 05 2006.
- L. J. O’Donnell, A. Daducci, D. Wassermann, and C. Lenglet. Advances in computational and statistical diffusion MRI. *NMR in Biomedicine*, 32(4):e3805, 04 2019.
- N. Painchaud, Y. Skandarani, T. Judge, O. Bernard, A. Lalande, and P.-M. Jodoin. Cardiac MRI segmentation with strong anatomical guarantees. In D. Shen, T. Liu, T. M. Peters, L. H. Staib, C. Essert, S. Zhou, P. Yap, and A. Khan, editors, *Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 - 22nd International Conference, Shenzhen, China, October 13-17, 2019, Proceedings*, volume 11765 of *Lecture Notes in Computer Science*, pages 632–640. Medical Image Computing and Computer Assisted Intervention - MICCAI 2019 - 22nd International Conference, Springer, 10 2019.
- S. M. Patil, A. Nigam, A. Bhavsar, and C. Chattopadhyay. Siamese LSTM based fiber structural similarity network (fs2net) for rotation invariant brain tractography segmentation. arXiv, 12 2017. abs/1712.09792.
- C. S. Perone, P. Ballester, R. C. Barros, and J. Cohen-Adad. Unsupervised domain adaptation for medical imaging segmentation with self-ensembling. *NeuroImage*, 194:1–11, 07 2019.
- F. Pestilli, J. D. Yeatman, A. Rokem, K. N. Kay, and B. A. Wandell. Evaluation and statistical inference for human connectomes. *Nature Methods*, 11(10):1058–1063, 09 2014.
- K. Pomiecko, C. Sestili, K. Fissell, S. Pathak, D. Okonkwo, and W. Schneider. 3D convolutional neural network segmentation of white matter tract masks from MR diffusion anisotropy maps. In *16th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, Venice, Italy, 04 2019. Institute of Electrical and Electronics Engineers (IEEE).
- C. Presseau, P.-M. Jodoin, J.-C. Houde, and M. Descoteaux. A new compression format for fiber tracking datasets. *NeuroImage*, 109:73–83, 04 2015.
- M. Reisert, I. Mader, C. Anastasopoulos, M. Weigel, S. Schnell, and V. Kiselev. Global fiber reconstruction becomes practical. *NeuroImage*, 54(2):955–962, 01 2011.
- F. Rheault, E. St-Onge, J. Sidhu, K. H. Maier-Hein, N. Tzourio-Mazoyer, L. Petit, and M. Descoteaux. Bundle-specific tractography with incorporated anatomical and orientational priors. *NeuroImage*, 186:382–398, 02 2019.
- F. Rheault, A. De Benedictis, A. Daducci, C. Maffei, C. M. W. Tax, D. Romascano, E. Caverzasi, F. C. Morency, F. Corrivetti, F. Pestilli, G. Girard, G. Theaud, I. Zemmoura, J. Hau, K. Glavin, K. M. Jordan, K. Pomiecko, M. Chamberland, M. Baraković, N. Goyette, P. Poulin, Q. Chenot, S. S. Panesar, S. Sarubbo, L. Petit, and M. Descoteaux. Tractostorm: The what, why, and how of tractography dissection reproducibility. *Human Brain Mapping*, 41(7):1859–1874, 01 2020a.
- F. Rheault, P. Poulin, A. V. Caron, E. St-Onge, and M. Descoteaux. Common misconceptions, hidden biases and modern challenges of dMRI tractography. *Journal of Neural Engineering*, 17(1):011001, 02 2020b.
- F. M. Rheault. *Analyse et reconstruction de faisceaux de la matière blanche*. PhD thesis, Université de Sherbrooke, 08 2020. URL <http://hdl.handle.net/11143/17255>. Chapter 4: Segmentation automatique de faisceaux par approches multi-atlas avec fusion d’étiquettes.
- O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, Cham, 2015. Springer International Publishing.
- S. Sarubbo, L. Petit, A. De Benedictis, F. Chioffi, M. Pfito, and T. B. Dyrby. Uncovering the inferior fronto-occipital fascicle and its topological organization in non-human primates: the missing connection for language evolution. *Brain Structure and Function*, 224(4):1553–1567, 03 2019. Supplementary material 1.
- S. Schiavi, M. Ocampo-Pineda, M. Baraković, L. Petit, M. Descoteaux, J.-P. Thiran, and A. Daducci. A new method for accurate in vivo mapping of human brain connections using microstructural and anatomical information. *Science Advances*, 6(31):1–11, 07 2020.

- K. G. Schilling, L. Petit, F. Rheault, S. Remedios, C. Pierpaoli, A. W. Anderson, B. A. Landman, and M. Descoteaux. Brain connections derived from diffusion MRI tractography can be highly anatomically accurate-if we know where white matter pathways start, where they end, and where they do not go. *Brain Structure and Function*, 225(8): 2387–2402, 08 2020.
- N. Sharmin, E. Olivetti, and P. Avesani. White matter tract segmentation as multiple linear assignment problems. *Frontiers in Neuroscience*, 11:754, 02 2018.
- V. Siless, S. Medina, G. Varoquaux, and B. Thirion. A comparison of metrics and algorithms for fiber clustering. In *2013 International Workshop on Pattern Recognition in Neuroimaging (PRNI)*, pages 190–193, Philadelphia, PA, USA, 06 2013.
- V. Siless, K. Chang, B. Fischl, and A. Yendiki. AnatomiciCuts: Hierarchical clustering of tractography streamlines based on anatomical similarity. *NeuroImage*, 166:32–45, 02 2018.
- R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly. Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information. *NeuroImage*, 62(3):1924–1938, 09 2012.
- R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly. SIFT: Spherical-deconvolution informed filtering of tractograms. *NeuroImage*, 67:298–312, 02 2013.
- R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly. SIFT2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography. *NeuroImage*, 119:338–351, 10 2015.
- S. N. Sotiropoulos and A. Zalesky. Building connectomes using diffusion MRI: why, how and but. *NMR in Biomedicine*, 32(4):e3752, 04 2019. e3752 NBM-16-0139.R1.
- E. St-Onge, A. Daducci, G. Girard, and M. Descoteaux. Surface-enhanced tractography (SET). *NeuroImage*, 169: 524–539, 04 2018.
- V. Sydnor, A. M. Rivas-Grajales, A. Lyall, F. Zhang, S. Bouix, S. Karmacharya, M. E. Shenton, C.-F. Westin, N. Makris, D. Wassermann, L. J. O’Donnell, and M. R. Kubicki. A comparison of three fiber tract delineation methods and their impact on white matter analysis. *NeuroImage*, 178:318–331, 09 2018.
- H. Takemura, C. F. Caiafa, B. A. Wandell, and F. Pestilli. Ensemble tractography. *PLOS Computational Biology*, 12(2): 1–22, 02 2016.
- H. Takemura, F. Pestilli, K. S. Weiner, G. A. Keliris, S. M. Landi, J. Sliwa, F. Q. Ye, M. A. Barnett, D. A. Leopold, W. A. Freiwald, N. K. Logothetis, and B. A. Wandell. Occipital white matter tracts in human and macaque. *Cerebral Cortex*, 27(6):3346–3359, 03 2017.
- J.-D. Tournier, F. Calamante, and A. Connelly. Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage*, 35(4):1459–1472, 05 2007.
- D. Ugurlu, Z. Firat, U. Türe, and G. Unal. Supervised classification of white matter fibers based on neighborhood fiber orientation distributions using an ensemble of neural networks. In E. Bonet-Carne, F. Grussu, L. Ning, F. Sepelband, and C. M. W. Tax, editors, *Computational Diffusion MRI (CDMRI) Medical Image Computing and Computer Assisted Intervention (MICCAI) Workshop*, Mathematics and Visualization, pages 143–154, Granada, Spain, 09 2018. Medical Image Computing and Computer Assisted Intervention (MICCAI) Society, Springer.
- L. van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 11 2008.
- S. Wakana, A. Caprihan, M. M. Panzenboeck, J. H. Fallon, M. Perry, R. L. Gollub, K. Hua, J. Zhang, H. Jiang, P. Dubey, A. Blitz, P. van Zijl, and S. Mori. Reproducibility of quantitative tractography methods applied to cerebral white matter. *NeuroImage*, 36(3):630–644, 07 2007.
- J. Wang and Y. Shi. GIFE: Efficient and robust group-wise isometric fiber embedding. In G. Wu, I. Reikik, M. D. Schirmer, A. W. Chung, and B. Munsell, editors, *Connectomics in Neuroimaging (CNI) Medical Image Computing and Computer Assisted Intervention (MICCAI) Second International Workshop*, volume 11083 of *Lecture Notes in Computer Science*, pages 20–28, Granada, Spain, 09 2018. Medical Image Computing and Computer Assisted Intervention (MICCAI) Society, Springer.
- J. Wang, D. B. Aydogan, R. Varma, A. W. Toga, and Y. Shi. Modeling topographic regularity in structural brain connectivity with application to tractogram filtering. *NeuroImage*, 183:87–98, 12 2018.
- S. Warrington, K. L. Bryant, A. A. Khrapitchev, J. Sallet, M. Charquero-Ballester, G. Douaud, S. Jbabdi, R. B. Mars, and S. N. Sotiropoulos. Xtract - standardised protocols for automated tractography in the human and macaque brain. *NeuroImage*, 217:116923, 08 2020.

- D. Wassermann, L. Bloy, E. Kanterakis, R. Verma, and R. Deriche. Unsupervised white matter fiber clustering and tract probability map generation: Applications of a gaussian process framework for white matter fibers. *NeuroImage*, 51(1):228–241, 05 2010.
- D. Wassermann, N. Makris, Y. Rathi, M. E. Shenton, R. Kikinis, M. Kubicki, and C.-F. Westin. The white matter query language: a novel approach for describing human white matter anatomy. *Brain Structure and Function*, 221(9):4705–4721, 12 2016.
- J. Wasserthal, P. F. Neher, and K. H. Maier-Hein. TractSeg - fast and accurate white matter tract segmentation. *NeuroImage*, 183:239–253, 12 2018.
- C.-H. Yeh, R. E. Smith, X. Liang, F. Calamante, and A. Connelly. Correction for diffusion MRI fibre tracking biases: The consequences for structural connectomic metrics. *NeuroImage*, 142:150–162, 11 2016.
- C.-H. Yeh, R. E. Smith, T. Dhollander, F. Calamante, and A. Connelly. Connectomes from streamlines tractography: Assigning streamlines to brain parcellations is not trivial but highly consequential. *NeuroImage*, 199:160–171, 10 2019.
- C.-H. Yeh, D. K. Jones, X. Liang, M. Descoteaux, and A. Connelly. Mapping structural connectivity using diffusion MRI: challenges and opportunities. *Journal of Magnetic Resonance Imaging*, 53(6):1666–1682, 06 2020.
- A. Yendiki, P. Panneck, P. Srinivasan, A. Stevens, L. Zöllei, J. Augustinack, R. Wang, D. Salat, S. Ehrlich, T. E. Behrens, S. Jbabdi, R. L. Gollub, and B. Fischl. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. *Frontiers in Neuroinformatics*, 5:23, 10 2011.
- A. Zalesky, A. Fornito, L. Cocchi, L. L. Gollo, M. P. van den Heuvel, and M. Breakspear. Connectome sensitivity or specificity: which is more important? *NeuroImage*, 142:407–420, 11 2016.
- F. Zhang, S. Cetin Karayumak, N. Hoffmann, Y. Rathi, A. J. Golby, and L. J. O’Donnell. Deep white matter analysis (DeepWMA): Fast and consistent tractography segmentation. *Medical Image Analysis*, 65:101761, 10 2020.
- Y. Zhang, J. Zhang, K. Oishi, A. V. Faria, H. Jiang, X. Li, K. Akhter, P. Rosa-Neto, G. B. Pike, A. C. Evans, A. W. Toga, R. Woods, J. C. Mazziotta, M. I. Miller, P. C. M. van Zijl, and S. Mori. Atlas-guided tract reconstruction for automated and comprehensive examination of the white matter anatomy. *NeuroImage*, 52(4):1289–1301, 10 2010.
- S. Zhong, Z. Chen, and G. Egan. Auto-encoded latent representations of white matter streamlines. In *28th Virtual Conference & Exhibition of the International Society for Magnetic Resonance in Medicine (ISMRM)*. International Society for Magnetic Resonance in Medicine, 08 2020. Abstract #0850.

A Appendix

A.1 Reconstruction

Figure 13 demonstrates the result of the autoencoder reconstruction corresponding to a subset of the “Fiber Cup” dataset test set streamlines. The considered classes (plausibles and implausibles) are shown separately according to the FINTA filtering procedure. The figures show that the plausible streamlines match the expected bundle-wise color code information from the synthetic ground truth streamlines (see figure 3(a)). The implausible set can be seen to contain streamlines whose tracking had prematurely stopped within the tissue of interest, as well as streamlines connecting regions that are not connected in the synthetic ground truth. Additionally, and as anticipated by the autoencoder theory, it can be seen that the reconstructions are a smoothed version of the input tractogram streamlines in figure 3(d) (also visible in figure 5).

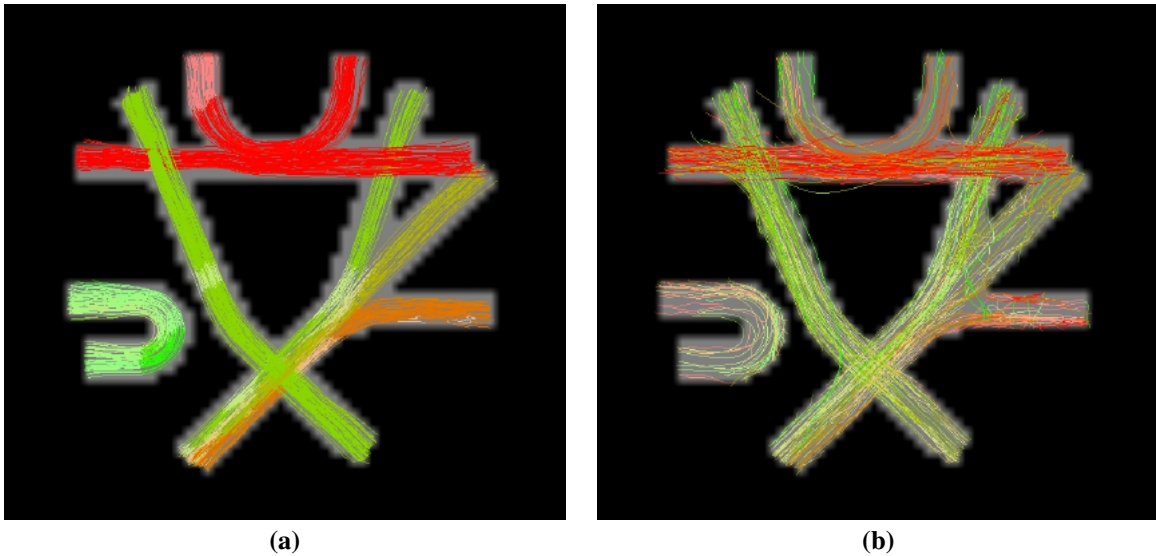


Figure 13: “Fiber Cup” test set filtered streamlines as reconstructed by the autoencoder: (a) plausible streamlines; (b) implausible streamlines. The number of streamlines displayed is lower than the actual number in the test set for illustrative purposes.

A.2 Predicted positives on the ISMRM 2015 Tractography Challenge human-based synthetic data

Figure 14 shows the predicted positive streamlines on the ISMRM 2015 Tractography Challenge dataset test tractogram. A visual inspection of the middle cerebellar peduncle (pointed bundle) shows that the filtered tractogram shows an orientation information that matches better the dataset ground truth shown in figure 3(b) in comparison to the raw, unfiltered tractogram in figure 3(e).

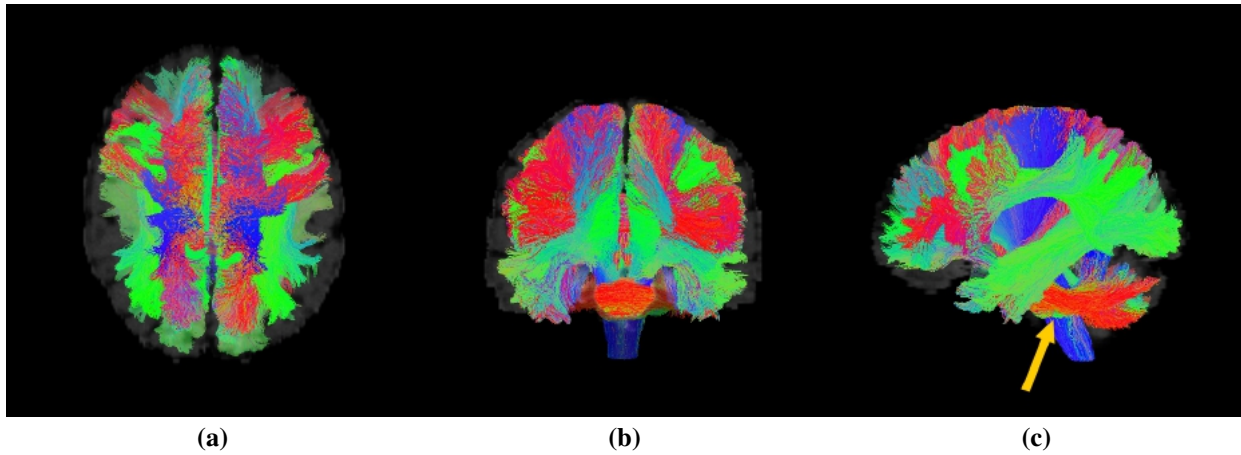


Figure 14: Predicted positives on the test set corresponding to the ISMRM 2015 Tractography Challenge dataset. (a) Axial superior view; (b) coronal anterior view; (c) sagittal left view. Although due to the streamline density the visualization can be difficult to interpret, the middle cerebellar peduncle (pointed bundle) clearly shows that the filtered tractogram has an orientation information that matches better the dataset ground truth shown in figure 3(b) in comparison to the raw, unfiltered tractogram in figure 3(e).

A.3 BIL&GIN human data filtering supplementary results

Figure 15 shows the ROC curves corresponding to the callosal BIL&GIN dataset subjects used to compute the filtering threshold. As it can be concluded from the curves, FINTA is consistent at filtering the tractograms across subjects. At the same time, results are in agreement with the experiments performed on the synthetic datasets.

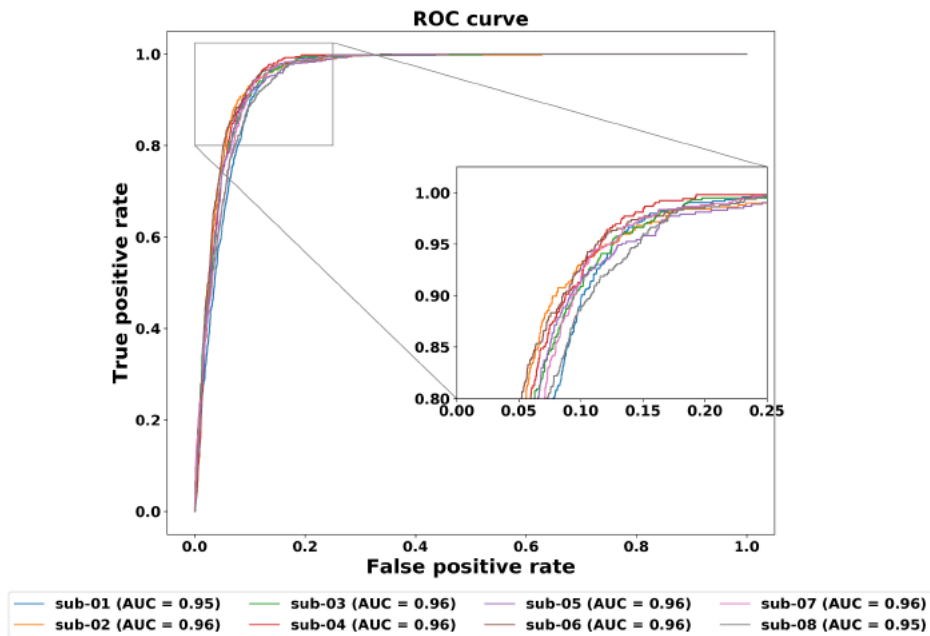


Figure 15: Callosal BIL&GIN dataset receiver operating characteristic (ROC) curves. The ROC curve shows the true positive rate (TPR) against the false positive rate (FPR) at various threshold settings. The figure shows the curves for the subjects used to compute the filtering threshold. The area under the curve (AUC) is indicated for each subject.

A.4 BIL&GIN human data performance supplementary results

Tables 6, 7, and 8 present supplementary results of the comparison to the baseline filtering methods on the callosal BIL&GIN human data: table 6 shows the weighted mean and standard deviation values, and tables 7 and 8 show the minimum and maximum values on the macro and weighted mean values. The weighted measures still show that FINTA provides an improved performance compared to the baseline methods. As mentioned, the latter yield an excessively large amount of false positives, and thus their valid gyrus-wise rates are generally higher than the values scored by FINTA. Nevertheless, FINTA still scores values over 0.68, with a maximum value of 0.94 across subjects.

Table 6: Callosal BIL&GIN dataset weighted results. Weighted mean (standard deviation) values over test subjects. The highest mean score is marked in bold face.

	Method	Balanced Accuracy	Sensitivity _w	Precision _w	F1-score _w	VGR	SR _w
#1	length	0.51 (0.0)	0.14 (0.02)	0.89 (0.01)	0.07 (0.01)	0.98 (0.03)	0.52
#2	no_loops	0.54 (0.01)	0.19 (0.02)	0.9 (0.01)	0.16 (0.02)	0.97 (0.03)	0.55
#3	no_end_in_csf	0.55 (0.01)	0.21 (0.02)	0.9 (0.01)	0.19 (0.02)	0.97 (0.03)	0.56
#4	end_in_atlas	0.67 (0.05)	0.71 (0.02)	0.85 (0.03)	0.76 (0.02)	0.96 (0.04)	0.79
#5	recobundles_single	0.85 (0.03)	0.83 (0.02)	0.91 (0.01)	0.86 (0.02)	0.87 (0.1)	0.86
#6	recobundles_multi	0.8 (0.03)	0.82 (0.03)	0.89 (0.02)	0.84 (0.03)	0.96 (0.04)	0.86
#7	FINTA	0.91 (0.01)	0.91 (0.01)	0.94 (0.01)	0.92 (0.01)	0.8 (0.09)	0.89

Table 7: Callosal BIL&GIN dataset macro minimum and maximum values. Macro mean [minimum, maximum] values over test subjects.

	Method	Accuracy _m	Sensitivity _m	Precision _m	F1-score _m	VGR
#1	length	[0.12, 0.16]	[0.5, 0.52]	[0.55, 0.57]	[0.11, 0.15]	[0.94, 1]
#2	no_loops	[0.15, 0.21]	[0.53, 0.55]	[0.55, 0.58]	[0.15, 0.21]	[0.94, 1]
#3	no_end_in_csf	[0.18, 0.24]	[0.54, 0.56]	[0.55, 0.57]	[0.18, 0.24]	[0.94, 1]
#4	end_in_atlas	[0.68, 0.75]	[0.56, 0.74]	[0.55, 0.62]	[0.53, 0.62]	[0.88, 1]
#5	recobundles_single	[0.81, 0.87]	[0.79, 0.88]	[0.65, 0.72]	[0.68, 0.75]	[0.72, 1]
#6	recobundles_multi	[0.78, 0.86]	[0.75, 0.85]	[0.64, 0.69]	[0.67, 0.72]	[0.89, 1]
#7	FINTA	[0.89, 0.91]	[0.89, 0.92]	[0.76, 0.81]	[0.81, 0.85]	[0.68, 0.94]

Table 8: Callosal BIL&GIN dataset weighted minimum and maximum values. Weighted mean [minimum, maximum] values over test subjects. The valid gyrus-wise rate (VGR) is the same as the one reported in table 7.

	Method	Balanced Accuracy	Sensitivity _w	Precision _w	F1-score _w	VGR
#1	length	[0.5, 0.52]	[0.12, 0.16]	[0.87, 0.92]	[0.05, 0.08]	[0.94, 1]
#2	no_loops	[0.53, 0.55]	[0.15, 0.21]	[0.88, 0.92]	[0.13, 0.18]	[0.94, 1]
#3	no_end_in_csf	[0.54, 0.56]	[0.18, 0.24]	[0.88, 0.92]	[0.16, 0.22]	[0.94, 1]
#4	end_in_atlas	[0.56, 0.74]	[0.68, 0.75]	[0.79, 0.88]	[0.73, 0.79]	[0.88, 1]
#5	recobundles_single	[0.79, 0.88]	[0.81, 0.87]	[0.89, 0.93]	[0.84, 0.89]	[0.72, 1]
#6	recobundles_multi	[0.75, 0.85]	[0.78, 0.86]	[0.86, 0.93]	[0.8, 0.88]	[0.89, 1]
#7	FINTA	[0.89, 0.93]	[0.89, 0.93]	[0.93, 0.95]	[0.9, 0.93]	[0.68, 0.94]

A.5 The callosal BIL&GIN dataset

The callosal BIL&GIN dataset was composed of the streamlines passing through the corpus callosum in 39 participants of the BIL&GIN dataset. For each subject, callosal streamlines were extracted from a whole brain tractogram computed with the probabilistic setting of the Particle Filtering Tracking (PFT) method by Girard et al. ((2014)) as previously described in Chenot et al. ((2019)). Each individual callosal tractogram was composed of anatomically plausible and implausible streamlines. The former were composed of gyral-based homotopic callosal streamlines, any other type of callosal streamline being considered as anatomically implausible. Twenty-six (26) different types of gyral-based homotopic streamlines were considered depending on the homotopic gyri they belonged to, namely the 26 gyral regions of interest of the JHU template ((Zhang et al., 2010)). Each gyral ROI was composed of both cortical and superficial white matter parts of the gyrus as defined in the JHU template excepted for the insula, hippocampus and parahippocampal gyrus that have no superficial white matter part. The 26 gyral-based ROIs were therefore: Angular Gyrus (AG); Cingulum (Cing); Cuneus (Cu); Fusiform Gyrus (FuG); Hippocampus (Hippo); Inferior Frontal Gyrus (IFG); Inferior Occipital Gyrus (IOG); Inferior Temporal Gyrus (ITG); Insula (Ins); Lateral Fronto-Orbital

Gyrus (LFOG); Lingual Gyrus (LG); Middle Frontal Gyrus (MFG); Middle Fronto-Orbital Gyrus (MFOG); Middle Occipital Gyrus (MOG); Middle Temporal Gyrus (MTG); Parahippocampal Gyrus (PHG); Post-Central Gyrus (PoCG); Pre-Central Gyrus (PrCG); Pre-Cuneus (PrCu); Rectus Gyrus (RG); Superior Frontal Gyrus (SFG); Supramarginal Gyrus (SMG); Superior Occipital Gyrus (SOG); Superior Parietal Gyrus (SPG); Superior Temporal Gyrus (STG); and Temporal Pole (TPole). The gyrus-wise homotopic callosal streamline models, used in the RecoBundles filtering experiments, were built using the aforementioned procedure and according to the method described in Rheault ((2020)) over the entire BIL&GIN population (411 subjects). Figure 16 provides an overview of the considered models.

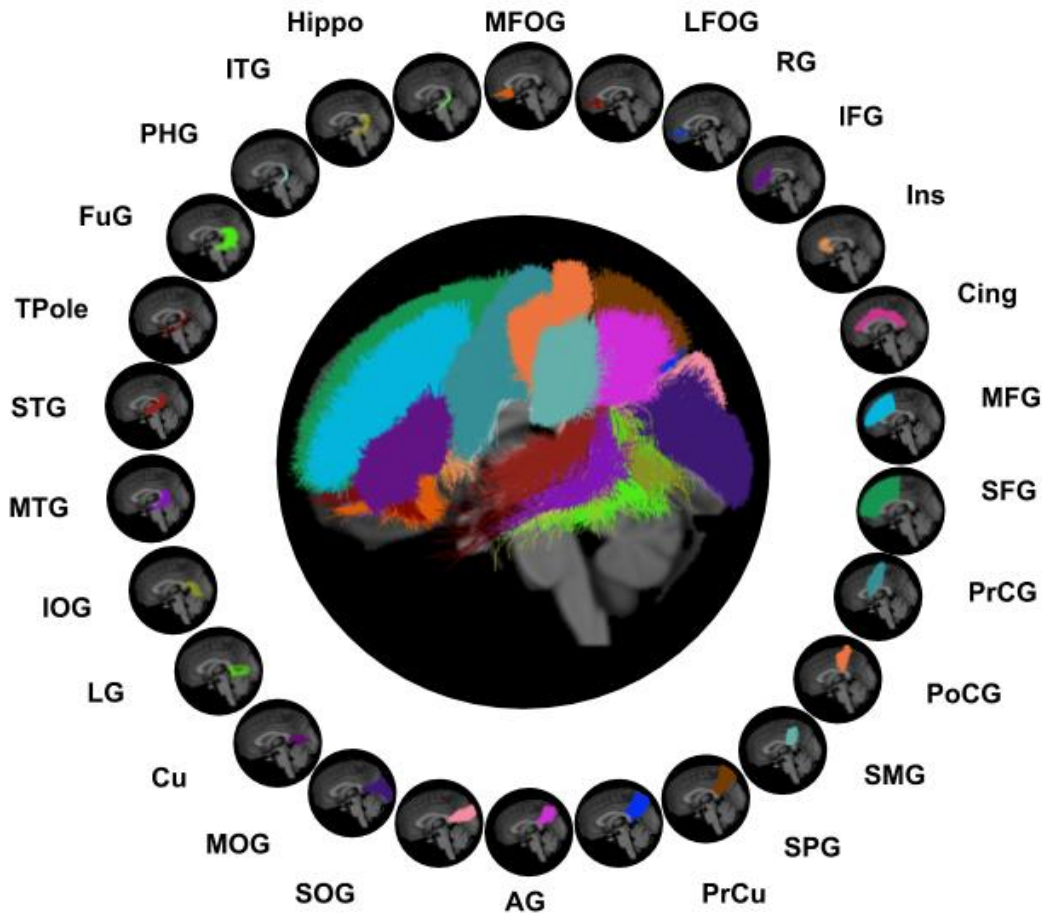


Figure 16: Callosal BIL&GIN dataset gyrus-wise homotopic streamline models. All sagittal left views.

As confirmed by diffusion-weighted tractography studies, the general research on the anatomy of the corpus callosum has traditionally been limited to its homotopic feature ((Hofer and Frahm, 2006; Chao et al., 2009)). Although recent developments ((De Benedictis et al., 2016)) have proposed solutions to reconstruct the most lateral homotopic connectivity, and, to a lesser extent, some heterotopic callosal connectivity in the frontal cortex, these still require further investigation with anatomical validation from dissection to be considered for the whole brain. Hence, in this work we focus on the most reliable representation of the callosal streamlines at the whole-brain level, namely the homotopic ones.

A.6 Callosal BIL&GIN gyrus-wise analysis

The gyral segment pairs defined in section A.5 for the callosal BIL&GIN dataset solely allow the extraction of homotopic streamlines. Hence, by definition, no implausible streamline can be recovered using the constraints established by these criteria. As a consequence, the only measure that can be reported for the predictions over the gyral segments is the sensitivity (see table 4). The mean number of segments with a non-zero streamline count in the reference data was 17 for the test subjects. Similarly, there was a non-negligible number of segments whose streamline density was very low compared to the most heavily populated ones. Test subjects averaged 5287 (standard deviation: 1450) homotopic

streamlines across all segments, and the number of segments that had a mean streamline count larger than 264 (5% of the mean homotopic streamline count) was 3: MFG; PrCu; and SFG.

A.7 Connectivity analysis

Additional experiments were performed to quantify the effect of FINTA into downstream tractography tasks. Structural connectivity matrices in terms of the streamline count were computed and analyzed for the “Fiber Cup” and ISMRM 2015 Tractography Challenge datasets. The following measures were computed on the binarized connectivity matrices to quantitatively analyze and compare them:

- *Density*: occupation index of the given connectivity matrix, an empty matrix having a value of 0, and 1 corresponding to a fully occupied matrix.
- *Chi-square*: defined as the squared element-wise differences across the matrices divided by their addition as in St-Onge et al. ((2018)). Lower is better.
- *Total variation*: computed as the absolute difference between matrices and normalized to the $[0, 1]$ range. Lower is better.
- *Correlation*: as defined by Deslauriers-Gauthier et al. ((2020)) and modified to be normalized to the $[0, 1]$ range. Higher is better.

Considered together, the presented results in terms of the density of the matrices and the distances across the matrices show that the connectomes issued by FINTA are populated as expected.

A.7.1 “Fiber Cup” synthetic dataset

Figure 17 shows the defined regions used to compute the connectivity on the “Fiber Cup” dataset. Figure 18 groups the log-scale and binarized streamline count connectivity matrices. It detaches from the figures that FINTA is in agreement with the *Tractometer_basic*-scored data, which is at the same time, close to the ground truth.

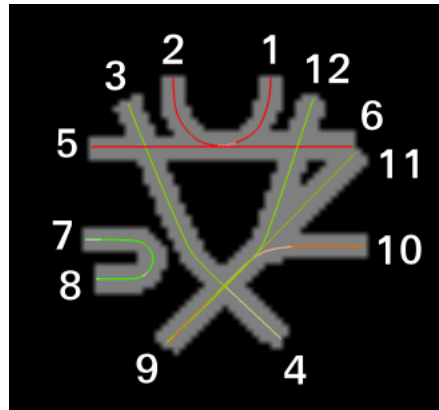


Figure 17: “Fiber Cup” dataset connectivity region definition.

Table 9 displays the density of the binarized connectivity matrices. The results reveal that FINTA is able to remove the undesired connections from the unfiltered dataset, and follows closely the value expected from the ground truth.

Table 9: Binarized streamline count connectivity matrix density for the “Fiber Cup” dataset. Note that the data corresponds to the test data (except for the ground truth).

Data	Ground truth	Unfiltered	Tractometer_basic	FINTA
Density	0.12	0.5	0.14	0.14

Table 10 shows that FINTA is identically correlated to the ground truth when compared to the scored data (i.e. *Tractometer_basic*), meaning that the proposed filtering method matches the expectations from the scoring.

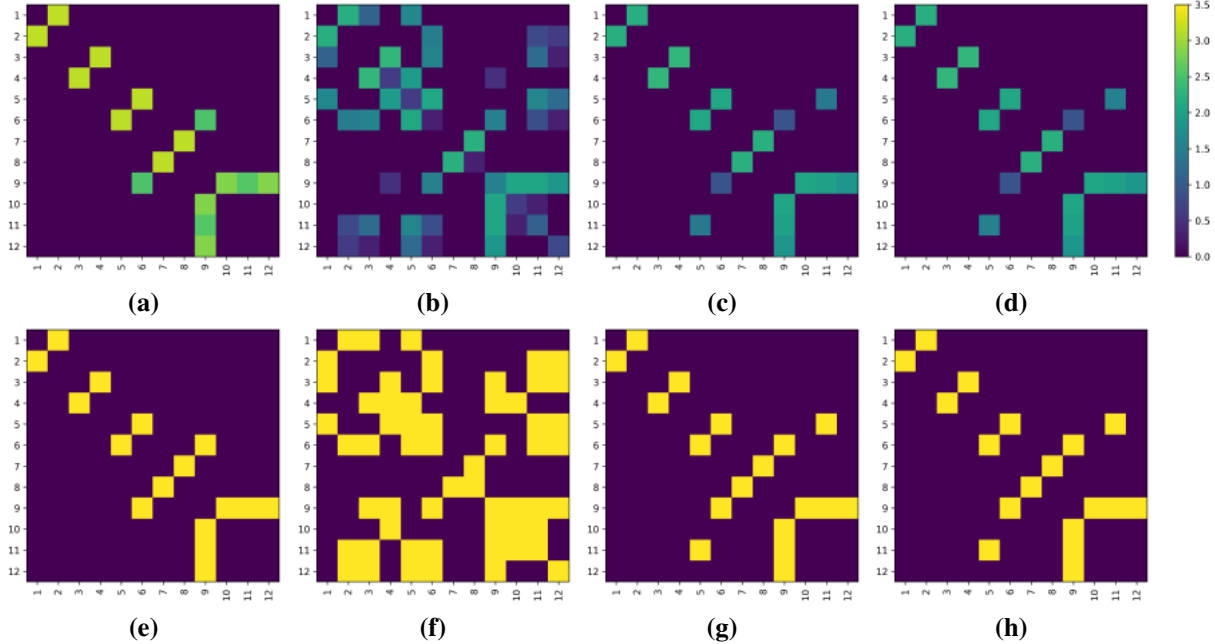


Figure 18: Streamline count connectivity matrices for the “Fiber Cup”. Upper row: log-scale; lower row: binarized. (a) and (e): ground truth; (b) and (f): unfiltered tractogram; (c) and (g): *Tractometer_basic*-scored tractogram; (d) and (h): FINTA-filtered tractogram.

Table 10: Distance measures on the binarized connectivity matrices for the “Fiber Cup” dataset.

	Chi-square	Total variation	Correlation
Ground truth vs. Unfiltered	21.0	0.29	0.48
Ground truth vs. Tractometer_basic	1.0	0.01	0.93
Ground truth vs. FINTA	1.0	0.01	0.93
Tractometer_basic vs. FINTA	0	0	1.0

A.7.2 ISMRM 2015 Tractography Challenge human-based synthetic data

The connectivity analysis was performed on the local probabilistic tracking data (see section 3) using the 112 endpoint regions corresponding to those originally used in the Challenge data preparation. Figure 19 shows the log-scale and binarized streamline count connectivity matrices. A qualitative assessment evidences FINTA’s ability to coherently dismiss the undesired associations from the connectome of a raw, unfiltered tractogram.

The values in table 11 evidence that the connectivity matrix density corresponding to the tractogram filtered with FINTA matches closely the one obtained for the *Tractometer_adv*-scored data, being both slightly above the ground truth value.

Table 11: Binarized streamline count connectivity matrix density for the local probabilistic tracking ISMRM 2015 Tractography Challenge dataset. Note that the data corresponds to the test data (except for the ground truth).

Data	Ground truth	Unfiltered	Tractometer_adv	FINTA
Density	0.17	0.34	0.21	0.2

As it is the case with the “Fiber Cup” dataset, table 12 shows that FINTA exhibits a correlation similar to the *Tractometer_adv*-scored data with respect to the ground truth on the realistic whole-brain synthetic dataset considered. Overall these results reveal that FINTA successfully filtered the tractogram without introducing significant changes to the structural connectivity.

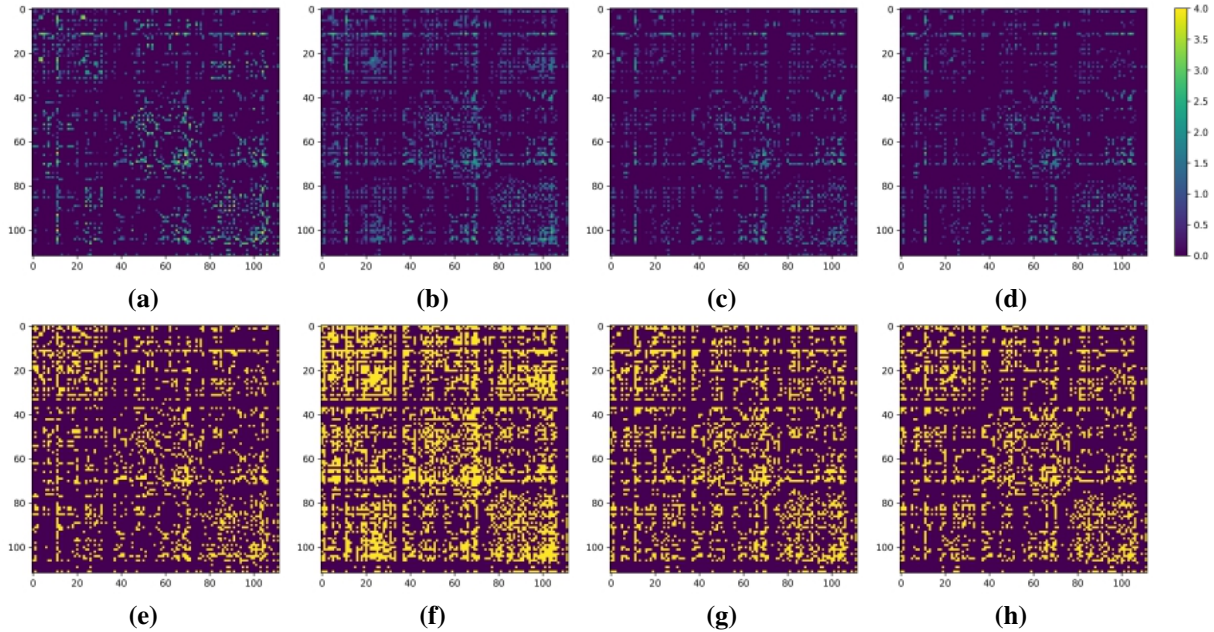


Figure 19: Streamline count connectivity matrices for the ISMRM 2015 Tractography Challenge dataset (local probabilistic tracking). Upper row: log-scale; lower row: binarized. (a) and (e): ground truth; (b) and (f): unfiltered tractogram; (c) and (g): *Tractometer_adv*-scored tractogram; (d) and (h): FINTA-filtered tractogram.

Table 12: Distance measures on the binarized connectivity matrices for the ISMRM 2015 Tractography Challenge dataset.

	Chi-square	Total variation	Correlation
Ground truth vs. Unfiltered	1353.5	0.22	0.49
Ground truth vs. <i>Tractometer_adv</i>	744.5	0.12	0.62
Ground truth vs. FINTA	726.0	0.12	0.61
<i>Tractometer_adv</i> vs. FINTA	242.5	0.04	0.88