



FIESTA: Autoencoders for accurate fiber segmentation in tractography

Félix Dumais^{a,b,*}, Jon Hartz Legarreta^{c,1}, Carl Lemaire^d, Philippe Poulin^{a,b}, François Rheault^e, Laurent Petit^f, Muhamed Barakovic^g, Stefano Magon^g, Maxime Descoteaux^{a,h,2}, Pierre-Marc Jodoin^{b,h,2}, for the Alzheimer's Disease Neuroimaging Initiative³

^a Sherbrooke Connectivity Imaging Lab (SCIL), Department of Computer Science, Université de Sherbrooke, Canada

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

^c Department of Radiology, Brigham and Women's Hospital, Mass General Brigham/Harvard Medical School, USA

^d Centre de Calcul Scientifique, Université de Sherbrooke, Canada

^e Medical Imaging and Neuroinformatics (MINI) Lab, Department of Computer Science, Université de Sherbrooke, Canada

^f Groupe d'Imagerie Neurofonctionnelle (GIN), CNRS, CEA, IMN, GIN, UMR 5293, F-33000 Bordeaux, Université de Bordeaux, France

^g Pharma Research and Early Development, Neuroscience and Rare Diseases, Roche Innovation Center Basel, F. Hoffmann-La Roche Ltd., Basel, Switzerland

^h Imeka Solutions inc, Sherbrooke, Canada

ARTICLE INFO

Dataset link: <https://doi.org/10.5281/zenodo.7562790>, <https://openneuro.org/datasets/ds003900>, <http://www.humanconnectomeproject.org/>, <https://doi.org/10.5281/zenodo.5165374>, <https://adni.loni.usc.edu/>, <https://www.ppmi-info.org/>, <https://github.com/scil-vital/fiesta>, <https://github.com/scil-vital/tractolearn>

Keywords:

Fiber tractography
Bundle segmentation
Autoencoder
Representation learning
Generative sampling
dMRI

ABSTRACT

White matter bundle segmentation is a cornerstone of modern tractography to study the brain's structural connectivity in domains such as neurological disorders, neurosurgery, and aging. In this study, we present FIESTA (*FibEr Segmentation in Tractography using Autoencoders*), a reliable and robust, fully automated, and easily semi-automatically calibrated pipeline based on deep autoencoders that can dissect and fully populate white matter bundles. This pipeline is built upon previous works that demonstrated how autoencoders can be used successfully for streamline filtering, bundle segmentation, and streamline generation in tractography. Our proposed method improves bundle segmentation coverage by recovering hard-to-track bundles with generative sampling through the latent space seeding of the subject bundle and the atlas bundle. A latent space of streamlines is learned using autoencoder-based modeling combined with contrastive learning. Using an atlas of bundles in standard space (MNI), our proposed method segments new tractograms using the autoencoder latent distance between each tractogram streamline and its closest neighbor bundle in the atlas of bundles. Intra-subject bundle reliability is improved by recovering hard-to-track streamlines, using the autoencoder to generate new streamlines that increase the spatial coverage of each bundle while remaining anatomically correct. Results show that our method is more reliable than state-of-the-art automated virtual dissection methods such as *RecoBundles*, *RecoBundlesX*, *TractSeg*, *White Matter Analysis* and *XTRACT*. Our framework allows for the transition from one anatomical bundle definition to another with marginal calibration efforts. Overall, these results show that our framework improves the practicality and usability of current state-of-the-art bundle segmentation framework

1. Introduction

White matter (WM) fiber tractography is a well-established method for brain connectivity analysis. It is currently the only non-invasive method able to investigate brain WM pathways *in vivo*. By using the local water diffusion information from diffusion-weighted Magnetic

Resonance Imaging (dMRI) images, one can infer the local orientation of the underlying WM streamlines (Descoteaux et al., 2009, 2007; Tournier et al., 2007) and use it to numerically reconstruct WM pathways. Over the years, many challenges have been tackled to improve this modeling technique such as *Global Tractography* (Kreher et al., 2008; Mangin et al., 2013; Christiaens et al., 2015), *Probabilistic Tractography* (Descoteaux et al., 2009; Tournier et al.,

* Correspondence to: Faculty of Sciences, Université de Sherbrooke, 2500, boulevard de l'Université, Sherbrooke, Québec, Canada, J1K 2R1.

E-mail address: felix.dumais@usherbrooke.ca (F. Dumais).

¹ Work done while at the SCIL and VITAL labs, Department of Computer Science, Université de Sherbrooke, Canada.

² Co-senior author. These authors contributed equally.

³ Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data, but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf.

2011, 2012), *Particle Filtering Tractography* (PFT) (Girard et al., 2014), *Bundle-Specific Tractography* (Wasserthal et al., 2018a; Rheault et al., 2019), or *Surface-Enhanced Tractography* (SET) (St-Onge et al., 2018).

The usability of WM tractography often comes from one's ability to filter and group streamlines into WM bundles. In this work, we refer to **filtering** as any method able to remove implausible streamlines from whole-brain tractograms (filtering methods yield implausible-free whole-brain tractograms). Filtering will later be discussed as it is used inside the bundle segmentation module to filter and segment tractograms (c.f. Fig. 1b and Section 2.3). Furthermore, there exists automated grouping methods that are 100% data-driven, fully unsupervised, whilst other methods follow pre-defined anatomical definitions. Thus, unsupervised grouping methods will be referred to as **clustering** (clustering methods yield streamline clusters), whilst grouping methods following pre-defined anatomical definitions will be referred to as **segmenting** (segmenting methods yield streamline bundles). The most faithful way to extract WM bundles from a tractogram is by asking a neuroanatomist to manually dissect bundles of interest (BOI). Manual dissection of bundles is long and tedious, and is prone to large inter- and intra-expert variability (Rheault et al., 2020a, 2022). Automated methods, such as *QuickBundles* (QB) (Garyfallidis et al., 2012), *QuickBundlesX* (QBx) (Garyfallidis et al., 2015a), *Deep Fiber Clustering* (DFC) (Chen et al., 2023), *RecoBundles* (RB) (Garyfallidis et al., 2018), *RecoBundlesX* (RBx) (Rheault, 2020), *TractSeg* (Wasserthal et al., 2018b, 2019, 2020), *XTRACT* (Warrington et al., 2020), and *White Matter Analysis* (WMA) (O'Donnell and Westin, 2007; O'Donnell et al., 2012; Zhang et al., 2018), amongst others, have been proposed to accelerate and increase the reproducibility of this process.

Unfortunately, these grouping methods are not void of limitations. First, clustering methods are fully unsupervised, thus excluding prior knowledge of each WM bundle's class membership (*QuickBundles*, *QuickBundlesX*, DFC). Second, even if there is a lack of consensus in the community over each bundle's anatomical definition (Rheault et al., 2020b), some state-of-the-art methods do not allow an easy modification of anatomical definitions (*TractSeg*, *XTRACT*). In fact, if a change in bundle definitions is needed with *TractSeg* (Wasserthal et al., 2018b), a full reannotation of the training dataset is required alongside with a complete retraining of its three neural networks. Contrastingly, *XTRACT* (Warrington et al., 2020) heavily relies on regions of interest (ROI) drawn by expert neuroanatomists to work properly. Those ROIs are usually hard to get and prone to inter-expert variability. Third, methods such as WMA are impractical in a large scale context because they rely on an affinity matrix computed using the pairwise distance of each streamline in a whole-brain tractogram. Such methods have prohibitive memory usage when the tractogram has more than 500,000 streamlines (Wasserthal et al., 2018b). Fourth, methods such as *RecoBundles* (Garyfallidis et al., 2018) and *RecoBundlesX* (Rheault, 2020) require the non-trivial calibration of several parameters and are more or less reliable in a test-retest analysis (see *RecoBundles* and *RecoBundlesX* results in Section 3). Finally, knowing all possible limitations of dMRI tractography (Rheault et al., 2020b), segmentation and clustering methods are tied to the tracking algorithms' ability to recover hard-to-track WM bundles in the first place.

In this paper, we present FIESTA (*FibEr Segmentation in Tractography using Autoencoders*), a reliable and robust, fully automated, and easily semi-automatically calibrated pipeline based on deep autoencoders that can dissect and fully populate WM bundles. Thus, FIESTA is a **filtering** and **bundle segmentation** (not **clustering**) pipeline whose output is subsequently improved by an autoencoder-based generative streamline sampling method. FIESTA allows an easy change in its bundle definitions, depending on the need, with a marginal calibrating time. This pipeline is built upon FINTA, CINTA, and GESTA methods (Legarreta et al., 2021, 2022, 2023) that demonstrated how autoencoders can be used successfully for filtering, bundle segmentation, and streamline generation in tractography.

1.1. Related work

Over the years, many methods have been developed to ease the interpretability of whole-brain tractograms, such as filtering methods (Legarreta et al., 2021; Maier-Hein et al., 2017; Petit et al., 2021; Jörgens et al., 2021; Sotiropoulos and Zalesky, 2019) and clustering methods (Garyfallidis et al., 2012, 2015a; Chen et al., 2021, 2022; O'Donnell and Westin, 2007; Wang et al., 2011; Visser et al., 2011). Unfortunately, filtering methods do not allow to easily target anatomical regions as tractograms are not grouped in streamline bundles. On the other hand, streamline clustering algorithms, a class of methods designed to group streamlines with similar properties, are typically built upon unsupervised machine learning approaches such as mixture models, spectral clustering and hierarchical clustering. Unfortunately, clustering methods suffer from a major drawback: they offer no control to which cluster belongs to which WM anatomical region.

Therefore, bundle segmentation comes as a solution for this problem. It aims to give streamlines an anatomical label. Many methods have been proposed in the literature to bundle whole-brain tractograms, such as *TRACULA* (Yendiki, 2011), *TractQuerier* (Wassermann et al., 2016), or *Classifyber* (Bertò et al., 2021). In this work, we will confine our reference methods to *RecoBundles* (Garyfallidis et al., 2018) and *RecoBundlesX* (Rheault, 2020), WMA (O'Donnell and Westin, 2007; O'Donnell et al., 2012; Zhang et al., 2018), *TractSeg* (Wasserthal et al., 2018a,b, 2019, 2020) and *XTRACT* (Warrington et al., 2020). Reference methods for this work are based on different criteria. The main goal is to capture the representativeness of the different existing state-of-the-art methods. *RecoBundles* and *RecoBundlesX* were selected for their ease of implementation, and because author F.R. developed the latter. *XTRACT* was selected to represent ROI-based bundle segmentation methods. WMA was selected to represent methods based on a non-physical embedding space. Finally, *TractSeg* was selected to capture the power of supervised deep learning methods.

RecoBundles is a well-established method based upon atlas bundles pre-segmented by medical experts. When a new tractogram needs to be bundled, a streamline-based linear registration (Garyfallidis et al., 2015b) is done between the target and the atlas tractograms, bringing the tractogram and the atlas in the same space. After this co-registration, a sequence of operations is performed for each bundle in the atlas. First, the target tractogram is pruned to remove irrelevant streamlines to the current bundle of interest. Then, a local streamline-based linear registration is performed to better match the target anatomy. Finally, a second pruning operation is performed using a stricter distance threshold to isolate the streamlines that have a high shape similarity to the current bundle of interest in the atlas. All the aforementioned steps are performed using centroids rather than individual streamlines to speed up processing.

RecoBundlesX is an improved version of *RecoBundles* using many iterations of *RecoBundles* over multiple atlases with different sets of parameters. Varying parameters are (i) whole-brain tractograms *QuickBundles* clustering thresholds, (ii) atlas bundle *QuickBundles* clustering thresholds, and (iii) pruning thresholds. While being much more reliable than *RecoBundles*, this method is time-consuming since *RecoBundles* is launched many times.

White Matter Analysis (WMA), proposed by O'Donnell and Westin (2007), is an atlas-based bundle recognition method where streamline classification is done in an embedding space built with spectral clustering (O'Donnell et al., 2006; O'Donnell and Westin, 2005). The atlas streamlines are represented as a point in the embedding space, where each point is assigned to a bundle class. New tractogram streamlines are bundled using the closest atlas centroid as the most probable class. One drawback of the method is that it requires computing a $N \times N$ pairwise streamline distance affinity matrix, which is computationally prohibitive for a reasonably sized tractogram. Therefore, the Nystrom approximation (Fowlkes et al., 2004) is done as a trade-off to approximate the affinity matrix.

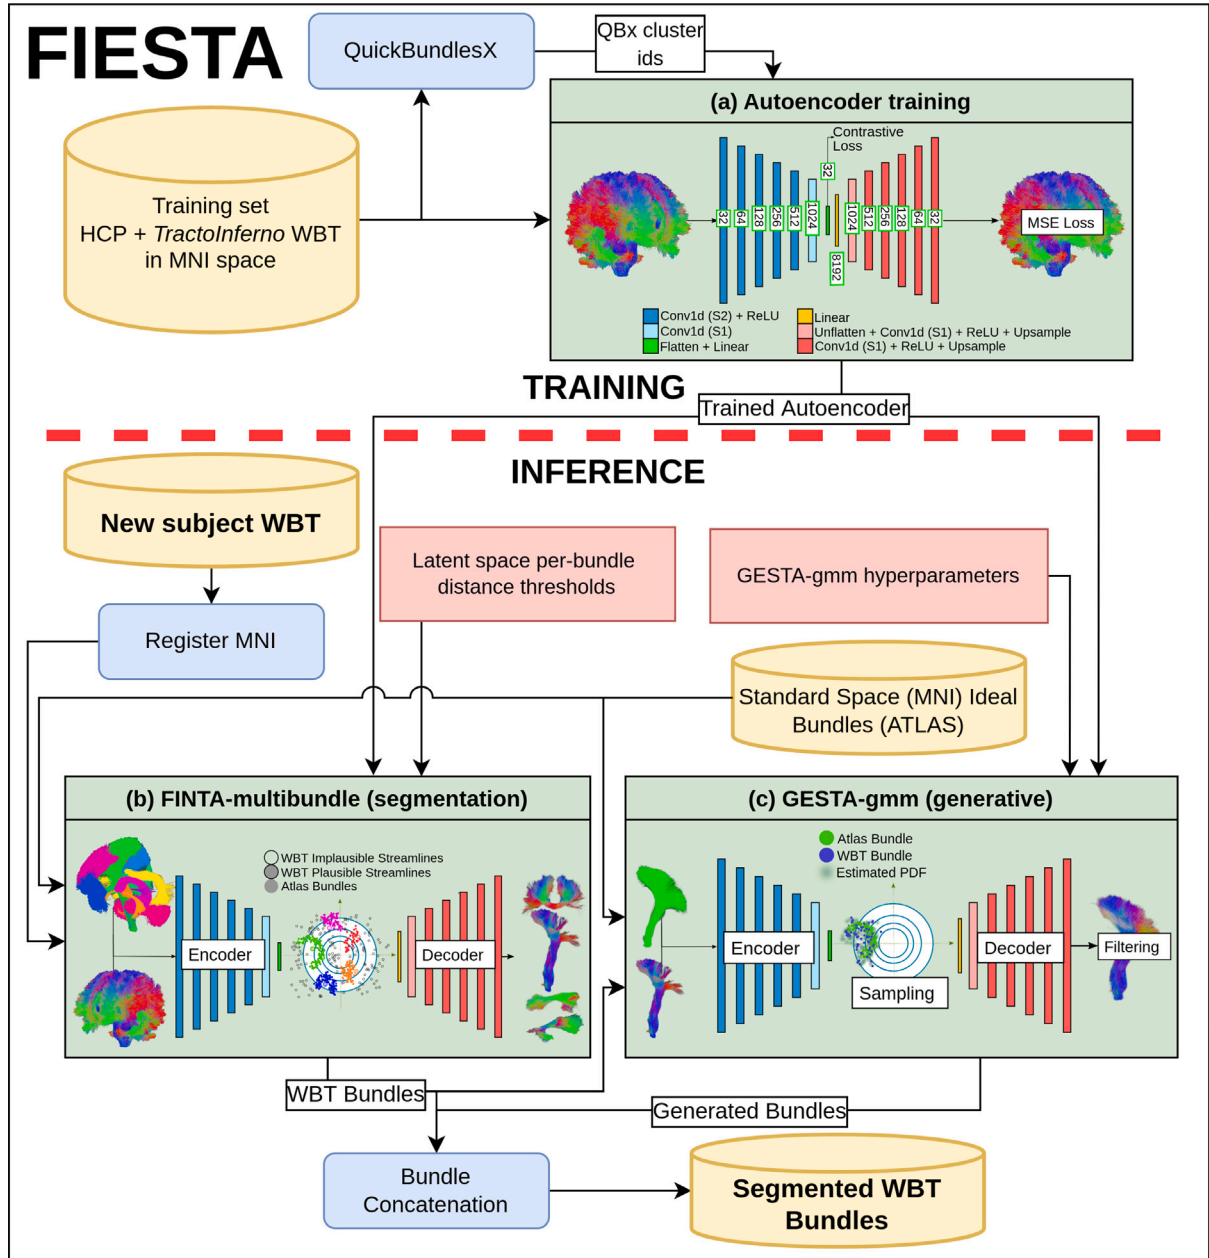


Fig. 1. FIESTA pipeline. (a) Training of a convolutional autoencoder on raw dMRI whole-brain tractograms (WBT) with a combination of a mean squared error (MSE) loss and a contrastive loss. A streamline is a 1D signal with 3 channels. The contrastive loss' positive and negative pairs are based on *QuickBundlesX* streamline cluster memberships. (b) The trained autoencoder is used to filter and bundle streamlines based on a k-nearest neighbors algorithm and a Euclidean distance threshold. The key idea behind the autoencoder's latent space is that plausible streamlines fall closer to the cleaned, filtered, and averaged bundles in the standard atlas than implausible streamlines. (c) A probability distribution function (PDF) is empirically estimated with a Parzen estimator (Bishop, 2006) based on the embedding of each bundle and its atlas counterpart. Rejection sampling is used to generate new streamline samples from the estimated PDF. Sampled streamlines are then filtered by checking their fit to the underlying diffusion signal, a length range, a maximum winding angle, and a WM coverage rate, finally yielding a bundle with better spatial coverage.

ROI-based methods, such as (Zhang et al., 2010), including the White Matter Query Language (WMQL) (Wassermann et al., 2016), use a pre-defined atlas of ROIs to accomplish the bundle segmentation task. ROIs are defined in a standard space and contain starting areas, inclusion and exclusion zones, and termination areas for each bundle. The regions are non-linearly registered to the diffusion space and used to bundle raw tractograms. More recently, Warrington et al. (2020) proposed the *XTRACT* protocol, which leverages ROI-based methods in order to achieve an improved generalizability.

TractSeg is a state-of-the-art deep learning bundle segmentation method. Presented over many articles (Wasserthal et al., 2018a,b, 2019, 2020), the latest version works in three steps. First, a constrained spherical deconvolution (CSD) is used over the diffusion orientation

distribution function (dODF) to extract the three principal peaks from the obtained fiber orientation distribution function (fODF). Next, peak maps are given to three U-Net neural networks (Ronneberger et al., 2015) yielding $B = 72$ bundle segmentation maps, $2 \times B$ start and end bundle segmentation maps, and B bundle-specific *Tract Orientation Maps* (TOM). Finally, probabilistic tractography is done over each TOM within their respective bundle segmentation maps and making sure start and end regions are respected for each streamline.

1.2. Contributions

In this paper, we present a semi-supervised bundle segmentation method called FIESTA. This work is the natural extension of FINTA

(Legarreta et al., 2021), CINTA (Legarreta et al., 2022), and GESTA (Legarreta et al., 2023), which use a deep convolutional autoencoder to project streamlines into a lower, more structured, smoother and locally linear dimensional space to either filter, segment, or generate WM streamlines (c.f. Supplemental section A.6 for more details). To better structure the latent representation, we trained the autoencoder with a contrastive loss (Hadsell et al., 2006) informed by *QuickBundlesX* clustering. FIESTA uses an autoencoder to filter and bundle whole-brain tractograms based on a given bundle atlas. It also takes advantage of its latent space sampling strategy to synthesize new streamlines and improve the coverage of its bundles. To avoid confusion and to be sure to distinguish the slight method variations between previous works and the current implemented methods, we renamed the implemented 2-in-1 filtering and bundle segmentation method as FINTA-multibundle, and the implemented generative method as GESTA-gmm (Gaussian mixture model) (c.f. Section 2). This work makes the following five contributions:

1. FIESTA leverages the power of previous fiber autoencoder methods (Legarreta et al., 2021, 2022, 2023) and overcomes their individual limitations, yielding a pipeline able to work in an end-to-end *in vivo* dMRI data analysis scenario;
2. FIESTA is more reliable than current state-of-the-art automatic bundle segmentation methods;
3. We show that the usage of generative sampling improves the bundle-wise volumetric reliability;
4. We show that we can train and use contrastive learning based on *QuickBundlesX* clusters to build a useful latent representation of a whole-brain tractogram;
5. Bundle definitions are easily editable in FIESTA without the need to re-train a neural network.

2. Methods

FIESTA's three main underlying key concepts are summarized in Fig. 1. First, an autoencoder is used to learn a streamline representation latent space (c.f. Fig. 1a). Next, given a new tractogram and a reference set of bundles (typically provided by a trained specialist), FINTA-multibundle extracts false-positive-free bundles from a given whole-brain tractogram (c.f. Fig. 1b). The tractogram is dissected by applying a filtering procedure to every bundle of interest individually (i.e., against all other groups, including implausible, and plausible streamlines belonging to other bundles). Finally, GESTA-gmm populates each bundle by generating plausible streamlines supported by a filtering process to increase the spatial coverage (Fig. 1c). FIESTA is a software suite built upon previous works that yields WM bundles from the concatenation of the output bundles of FINTA-multibundle and GESTA-gmm.

2.1. Data

Five datasets were used for the development and the testing of FIESTA — namely *TractoInferno* (Poulin et al., 2022), and the Human Connectome Project (HCP) (Glasser et al., 2013, 2016) as mentioned in Fig. 1, and *MyeloInferno* (Edde et al., 2023), the Alzheimer's Disease Neuroimaging Initiative (ADNI) (<https://adni.loni.usc.edu>), and the Parkinson's Progression Markers Initiative (PPMI) (www.ppmi-info.org/access-data-specimens/download-data) for the evaluation. The choice of datasets was based on several criteria, such as ease of access, multiple time-points per subject (*MyeloInferno*, ADNI, and PPMI), high-quality data (HCP), and the availability of multiple WM tracking types (*TractoInferno*). *TractoInferno* and HCP data were used for the training of the convolutional autoencoder and early evaluation stage, while *MyeloInferno*, ADNI, and PPMI were employed to evaluate the reliability of the pipeline. Also, *TractoInferno* 'silver standard' (Theaud et al.,

2022) bundles were used for the threshold calibration steps. HCP subjects were also used for adjusting GESTA-gmm's parameters.

TractoInferno (Poulin et al., 2021) is a publicly available, widely heterogeneous database designed for machine learning purposes in dMRI tractography. It is composed of 354 subjects from 6 different datasets acquired with 5 different MRI scanners with various image resolutions, acquisition parameters, and subject ages. MRI acquisitions underwent a manual quality control (QC) process before using an ensemble of four tracking types — namely local deterministic, local probabilistic, PFT (Girard et al., 2014), and SET (St-Onge et al., 2018) — before being processed into *RecoBundlesX*, yielding its silver standard. *TractoInferno* whole-brain tractograms were used for the autoencoder training (c.f. Section 2.2.1), whilst *TractoInferno* silver standard bundles were only used for the latent space distance threshold calibration (c.f. Section 2.3.1).

MyeloInferno (Edde et al., 2023) is a non-public dMRI dataset. We have divided the *MyeloInferno* dataset into 2 single time-point subsets and 2 five-time-point subsets. One single time-point subset, *MyeloInferno-HC*, was composed of 24 young and healthy control (HC) subjects (mean age 36 years \pm 4.7 [standard deviation]; 17 women) and the other, *MyeloInferno-MS*, was composed of 21 young subjects with multiple sclerosis (MS) (mean age 38 years \pm 6.8 [standard deviation]; 16 women). Written informed consent was obtained from participants and were recruited following the ethics protocol of the Centre de Recherche du Centre Hospitalier Universitaire de Sherbrooke (Sherbrooke, Canada). Most HC and MS subjects in *MyeloInferno* had 5 time-points or more. Thus, for the creation of our 2 test-retest subsets, we only kept data with 5 time-points or more and excluded time-points above 5 for our analysis. We named those 2 new subsets *MyeloInferno-HC-TR* and *MyeloInferno-MS-TR*. *MyeloInferno-HC-TR* contains $N = 18 \times 5 = 90$ acquisitions, whilst *MyeloInferno-MS-TR* contains $N = 19 \times 5 = 95$ acquisitions. Thus, the data are ideal to evaluate the reproducibility of our pipeline. All images were acquired on a 3T Ingenia MR scanner with a 32-channel head coil (Philips Healthcare, Best, Netherlands) with T1w images acquired with a resolution of 1 mm isotropic voxels and dMRI images acquired with a resolution of 2 mm isotropic voxels, 100 gradient directions uniformly distributed over three shells ($b = \{300(8), 2000(32), 3000(60)\} \text{ s/mm}^2$) (the number of directions per shell is in parentheses), and 7 unweighted images. Each tractogram was an ensemble of tractograms generated using the *TractoFlow* pipeline (Theaud et al., 2020), which provided a PFT probabilistic tracking and a local probabilistic tracking that were concatenated. PFT tractograms were generated using interface seeding with 30 seeds per voxel, while local tractograms were generated using white matter seeding with 10 seeds per voxel. Both tracking methods were constrained with a streamline length range between 20 and 200 mm. More details can be found in Edde et al. (2023) and on the corresponding website at <https://high-frequency-mri-database-supplementary.readthedocs.io/en/latest/index.html>.

The Human Connectome Project (HCP) Young Adult dataset (Glasser et al., 2013, 2016) is a dataset of approximately 1200 subjects (age range 22–35 y/o) composed of different MRI modalities acquired on Siemens scanners. For the purposes of this project, only the T1w and dMRI acquisitions were used. All images were acquired on a 3T MR scanner with T1w images acquired with a resolution of 0.7 mm isotropic voxels and dMRI images acquired with a resolution of 1.25 mm isotropic voxels, 270 gradient directions equally distributed over three shells ($b = \{1000, 2000, 3000\} \text{ s/mm}^2$) and 6 unweighted images. Again, the tracking was done using the *TractoFlow* pipeline (Theaud et al., 2020) based on fODF estimated using $b = \{0, 1000, 2000, 3000\} \text{ s/mm}^2$ shells, which provided a PFT probabilistic tracking and a local probabilistic tracking that were concatenated. PFT tractograms were generated using interface seeding with 60 seeds per voxel, while local tractograms were generated using white matter seeding with 30 seeds per voxel. Both tracking methods were constrained with a streamline length range between 20 and 200 mm.

The Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset (<https://adni.loni.usc.edu>) was used in the preparation of this article. The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD). For this study, T1w and dMRI acquisitions were used. All images were acquired on a 3T MR scanner from different vendors with T1w images acquired with a resolution of 1 mm isotropic voxels and dMRI images acquired with a resolution of 2 mm isotropic voxels, 114 gradient directions distributed over three shells ($b = \{500(6), 1000(48), 2000(60)\}$ s/mm²) (the number of directions per shell is in parentheses) and 12 unweighted images. Tracking was done using the *TractoFlow* pipeline (Theaud et al., 2020) based on fODF estimated using $b = \{0, 1000, 2000\}$ s/mm² shells, which provided a PFT probabilistic tracking and a local probabilistic tracking that were concatenated. PFT tractograms were generated using interface seeding with 4 seeds per voxel, while local tractograms were generated using white matter seeding with 3 seeds per voxel. Both tracking methods were constrained with a streamline length range between 20 and 200 mm. For our analysis we created a test-retest subset (ADNI-TR) comprised of 23 subjects with 5 time-points each (baseline, 3 months, 6 months, 12 months, 24 months) yielding a total of 115 images. Also, we created a single time-point subset (ADNI-HC) with 21 old healthy subjects.

The Parkinson's Progression Markers Initiative (PPMI) database (www.ppmi-info.org/access-data-specimens/download-data) was also used for this article. For this study, T1w and dMRI acquisitions were used. All images were acquired on a 3T MR scanner with T1w images acquired with a resolution of 1 mm isotropic voxels and dMRI images acquired with a resolution of 2 mm isotropic voxels, 64 gradient directions distributed over 1 shells ($b = 1000$ s/mm²) and 1 unweighted images. Again, the tracking was done using the *TractoFlow* pipeline (Theaud et al., 2020) based on fODF estimated using $b = \{0, 1000\}$ s/mm² shells, which provided a PFT probabilistic tracking and a local probabilistic tracking that were concatenated. PFT tractograms were generated using interface seeding with 4 seeds per voxel, while local tractograms were generated using white matter seeding with 3 seeds per voxel. Both tracking methods were constrained with a streamline length range between 20 and 200 mm. 34 subjects with 3 time-points each were used (baseline, 12 months, 24 months) yielding a total of 102 images used from this dataset.

TractoFlow (Theaud et al., 2020) was used to process HCP, *MyeloInferno*, ADNI, and PPMI subjects, yielding PFT and local probabilistic whole-brain tractograms, while *TractoInferno* authors were reached to get access to original raw tractograms (PFT, SET (St-Onge et al., 2018), local deterministic tracking and local probabilistic tracking).

2.1.1. Bundle atlas

As shown in Fig. 1, standard space ideal bundles (i.e., an atlas) are needed as a reference for FIESTA to work properly. For the development and evaluation of the current pipeline, we used an in-house but yet public Population Average of WM (PAWM) atlas to evaluate our framework (Rheault, 2021). The PAWM atlas was built in the context of *RecoBundlesX* works by Garyfallidis et al. (2018), Rheault (2020) based on ExTractor (Petit et al., 2022) and well-curated by a neuroanatomist to finally have bundles that fit their normative shapes (Rheault, 2021). Fig. 2 presents a representation of all the bundles in the PAWM atlas. Full names, abbreviations, and labels of bundles used in this work are indicated in Supplemental section A.1, whilst more context on the method used to generate these bundles is given in Supplemental section A.5.

2.2. Tractogram embedding

Fig. 1a shows a summary of the autoencoder's embedding and training scheme of the whole human brain streamlines. Autoencoders (Hinton and Salakhutdinov, 2006) are a type of neural network designed to compress and reconstruct input data as faithfully as possible. They are made of an encoder (in blue) and a decoder (in red), two 1D convolutional neural networks in our case. The encoder projects a sample from the input space to a latent space, and the decoder reverses that operation. An undercomplete autoencoder (Goodfellow et al., 2016) is one whose latent representation dimensionality is smaller than that of the input data. This forces the autoencoder to learn a latent space which embeds the most salient features of the input data. A well-constructed autoencoder encodes streamlines with similar properties (i.e., shapes, anatomical location, etc.) close to each other in the latent space.

Let $S = \{S_1, \dots, S_N\}$ be a set of N streamlines non-linearly registered in a common space (i.e., MNI (Fonov et al., 2009, 2011) using T1w image-based registration for the current pipeline) with $S_i = [s_i^1, \dots, s_i^k]$ where k is the number of vertices per streamline and $s_i^j \in \mathbb{R}^3$. By making sure that $k = D$ for all streamlines, with D being a constant, we train a 1D convolutional autoencoder on raw whole-brain tractograms to learn a useful WM streamline latent representation. The training loss contains two terms: an MSE reconstruction loss and a contrastive loss to help the bundle segmentation problem.

The contrastive loss runs over pairs of samples S_i, S_j to enforce streamlines of the same bundle to be as close as possible in the latent space. Following Hadsell et al. (2006), the contrastive loss is built upon a binary label y where $y = 0$ if S_i and S_j have similar properties and $y = 1$, if S_i and S_j are dissimilar. Also, let the parameterized Euclidean distance between two streamlines be defined as

$$D_\theta(S_i, S_j) = \|p_\theta(S_i) - p_\theta(S_j)\|_2, \quad (1)$$

where p_θ is the encoder function. The contrastive loss function is

$$C(\theta) = \sum_{i=1}^N C(\theta, y, S_i, S_j), \quad (2)$$

with

$$C(\theta, y, S_i, S_j) = (1 - y)C_S(D_\theta) + yC_D(D_\theta). \quad (3)$$

C_S and C_D are designed such that minimizing C w.r.t θ would result in low values of D_θ when S_i, S_j are within the same bundle and high values for D_θ otherwise. Therefore,

$$C(\theta, y, S_i, S_j) = (1 - y)\frac{1}{2}(D_\theta)^2 + y\frac{1}{2}\{\max(0, m - D_\theta)\}^2, \quad (4)$$

where the contrastive loss margin was set to $m = 1.25$, as in Hadsell et al. (2006).

Since whole-brain tractograms come with no annotation, similar and dissimilar pairs of streamlines are determined based on *QuickBundlesX* clusters (Garyfallidis et al., 2015a) obtained with $\{40, 30, 20, 10\}$ mm as input parameters. Thus, similar (or positive) streamline pairs are taken from the same *QuickBundlesX* cluster, whilst dissimilar (or negative) streamline pairs are taken from two different clusters.

Finally, the MSE reconstruction loss is defined as

$$\mathcal{M}(\theta, \phi) = \sum_{i=1}^N \|\hat{S}_i - S_i\|_2^2, \quad (5)$$

where $\hat{S}_i = q_\phi(p_\theta(S_i))$ is a reconstructed streamline and $q_\phi(\cdot)$ is the decoder. Therefore, the overall training loss is

$$\mathcal{L}(\theta, \phi) = \mathcal{M}(\theta, \phi) + \lambda C(\theta), \quad (6)$$

with λ being a hyperparameter. The hyperparameter λ was empirically determined to approximately balance the MSE loss with the contrastive loss on a randomly initialized autoencoder. In our case, the optimal value was $\lambda = 400$.

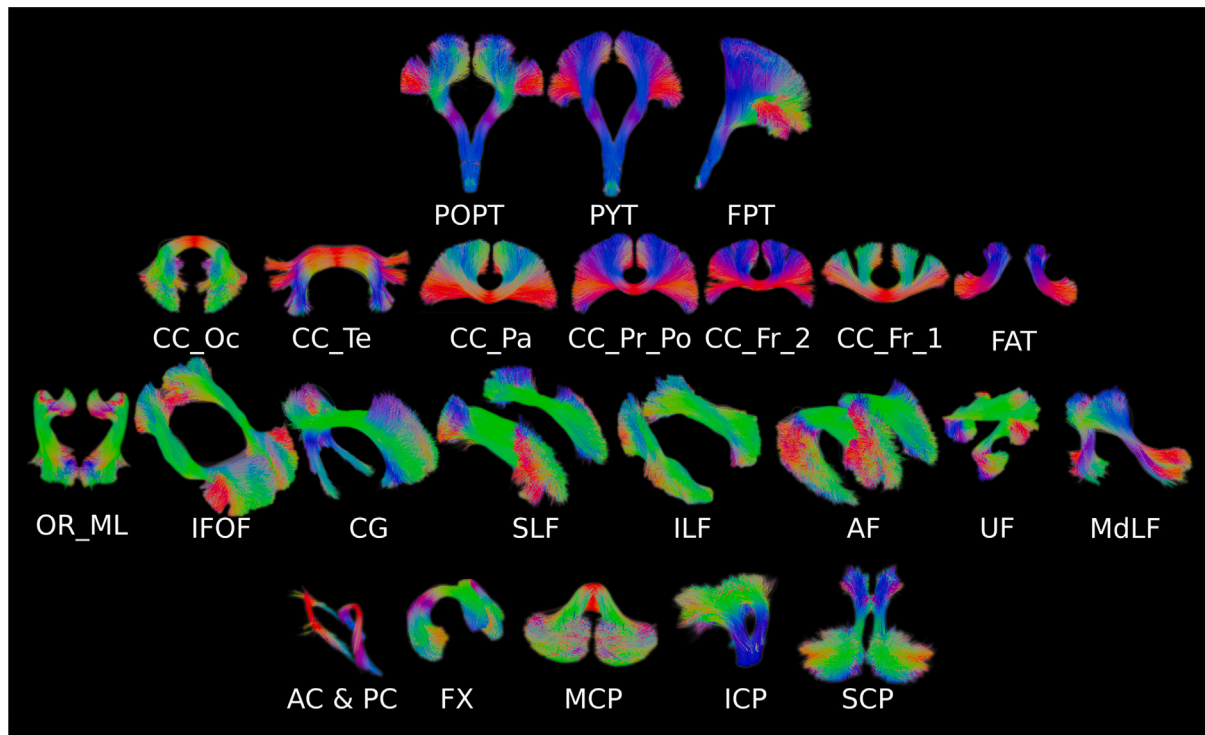


Fig. 2. PAWM atlas bundles used for FIESTA development and evaluation. All symmetric bundles are joined together (left and right). Abbreviations are described in full detail in section A.1.

2.2.1. The autoencoder training scheme

We trained the autoencoder in a train-validation-test scheme. Our whole dataset is composed of 120 subjects randomly shuffled with an 80/20 split between *TractoInferno* and HCP datasets (c.f. Section 2.1). The train/validation/test split was therefore 100/10/10. 200,000 streamlines per subject were used, with 100,000 streamlines per HCP whole-brain tractogram ($N = 2$ — namely PFT and local probabilistic tracking) and 50,000 streamlines per *TractoInferno* whole-brain tractogram were randomly sampled from each raw tractograms ($N = 4$ — namely PFT, SET, local deterministic tracking and local probabilistic tracking). In total, our dataset contained 24,000,000 streamlines. Also, each streamline is resampled to a fixed dimensionality of $D = 256$ 3D vertices and the latent space dimensionality is fixed to $d = 32$.

2.3. Tractogram segmentation

Once the autoencoder is trained, one can use it to filter and bundle a new tractogram. As displayed in Fig. 1b, every streamline that ought to be filtered and segmented is non-linearly registered into the MNI reference space using T1w image-based registration and projected into the latent space using the encoder neural network. This is done alongside the bundle streamlines of the atlas. A k -nearest neighbors (k -NN) algorithm is used to assign each streamline (in the latent space) to the majority class label among the nearest atlas streamline (also in the latent space). The absolute Euclidean distance between the WM streamline, and its atlas counterpart is then compared to a pre-determined threshold. To determine the threshold that best balances true positives and false positives, a ROC curve was used, as suggested by Legarreta et al. (2021), alongside manual adjustment (c.f. Section 2.3.1). Manual adjustments is mandatory to qualitatively match atlas bundles. Thus, if the distance is smaller than the threshold, the streamline is kept for downstream tasks. Otherwise, the streamline is considered implausible and, therefore, discarded.

2.3.1. Latent space distance threshold calibration steps

Bundle segmentation methods require a calibration to account for each bundle's unique shape; for example, with *RecoBundles* and *RecoBundlesX*, there are multiple parameters to set for each bundle. Fortunately, FIESTA has only one threshold per bundle (c.f. Fig. 4 to see threshold effect on bundles). For calibration, a *near-optimal* per-bundle threshold is determined automatically followed by manual adjustments to obtain the desired bundle shape (c.f. Fig. 3). *Near-optimal* thresholds were found by equally distributing the per-bundle true positive and false positive streamlines, using the ROC curve analysis method presented by Legarreta et al. (2021). We used *TractoInferno* silver standard (Theaud et al., 2022) bundles in the validation set as true positive streamlines. We consider as implausible the streamlines from the raw tractogram that are absent from the silver standard bundles. To avoid any confusion for the latent space calibration step, the following list explains each data used.

1. **IST**: Implausible streamlines from the *TractoInferno* validation set;
2. **SST**: Silver standard streamlines from the *TractoInferno* validation set;
3. **WBTT**: Whole-brain tractograms from the *TractoInferno* validation set, where $IST + SST = WBTT$;
4. **PAWM**: Ideal atlas of bundles comprised of plausible streamlines (c.f. Section 2.1.1).

The bundle-wise thresholds are thus determined as follows:

1. Encode atlas bundles (PAWM) and their flipped versions (to make the model independent of its direction) into latent vectors;
2. Encode plausible silver standard bundles (SST) and implausible streamlines (IST) and their flipped version from the validation set;
3. For each latent vector from the validation set (both plausible (SST) and implausible (IST)), find the closest streamline in

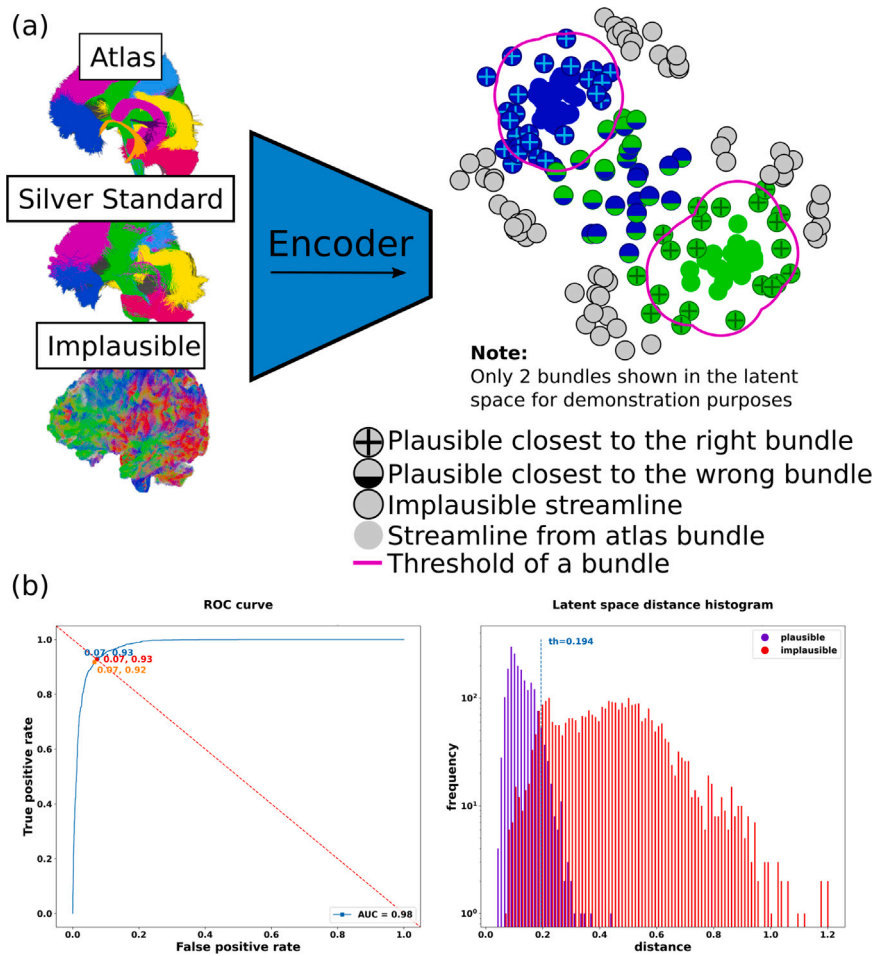


Fig. 3. Threshold calibration steps where (a) atlas, silver standard, and implausible streamlines in MNI space are projected in the latent space. *Near-optimal* thresholds are found by assigning all silver standard and implausible streamlines to the class of the closest atlas streamline and by balancing true positives and false positives for each bundle. (b) For each bundle in the atlas, the inverse diagonal of the ROC curve, built from the right silver standard streamlines assigned to the current atlas bundle (interpreted as positive streamlines) and the wrong silver standard combined with the implausible streamlines assigned to the current atlas bundle (interpreted as negative streamlines), is used to find the *near-optimal* threshold. **Left:** *near-optimal* threshold for 1 bundle found at the intersection of the ROC curve with the inverse diagonal and **Right:** histogram of one bundle with the *x*-axis giving the distance to the closest streamline in the atlas to each streamline in the silver standard (plausible) and implausible streamlines. The *near-optimal* distance threshold found with the ROC curve analysis is indicated.

the atlas using a k-NN algorithm. This step forms groups of candidate plausible and implausible streamlines with a class label;

- For each group of streamlines, optimize the Euclidean distance threshold that maximizes the number of true positive streamlines and minimize the number of false positive streamlines. Note that during this stage, candidate plausible streamlines assigned to the wrong class are considered implausible. This step is equivalent to taking the threshold value at the intersection between each ROC (bundle-wise) curve and the inverse diagonal line;
- Qualitatively manually adjust the *near-optimal* thresholds, if needed, to improve atlas bundle similarity with the validation set bundles.

To understand the effect of the choice of threshold, we visually inspected the shape and the quality of the resulting bundles. We set up an experiment with a subset of the bundles from the atlas presented in Section 2.1.1. We studied the effect of 8 uniformly spaced threshold values centered on the *near-optimal* threshold of the AF_L, the CC_Pr_Po, the IFOF_L, the OR_ML_L, the PYT_L and the UF_L bundles. Fig. 4 illustrates the results of that experiment. We set the values using 20% increments based on the ROC curves analysis thresholds. Therefore, it is possible to observe that low thresholds seems to increase bundle specificity at the expense of the sensitivity. On the other hand, high

threshold values, classifying more streamlines as positives, increase the sensitivity at the expense of the bundle specificity. Red circles in Fig. 4 indicate streamlines that should not be part of final bundles based on the atlas ideal bundles. Interestingly, such streamlines are not always present when the threshold is higher than the *near-optimal* threshold. In fact, we see that spurious streamlines are still present in the OR_ML even with a 40% threshold reduction. Thus, this analysis shows that threshold manual adjustment is mandatory, and threshold values based only on ROC curves analysis should not all be used as is. In our case, for some bundles, we needed to decrease the latent space distance threshold, thus increasing the specificity, to more closely match, qualitatively, bundles from the PAWM atlas (c.f. Fig. 2).

2.4. Bundle coverage improvement using generative sampling (GESTA-gmm)

As shown in Fig. 1c, GESTA-gmm is used, after FINTA-multibundle, to better populate each bundle. This module compensates the imperfections of the latent space-based threshold segmentation method and the missing streamlines in the original whole-brain tractogram, especially in hard-to-track bundles (Rheault et al., 2019). Thus, for each bundle in the atlas, we consider the streamlines labeled by FINTA-multibundle as belonging to the current bundle (see Fig. 1b). An input ratio 1:1 between the number of streamlines to use from each bundle (atlas

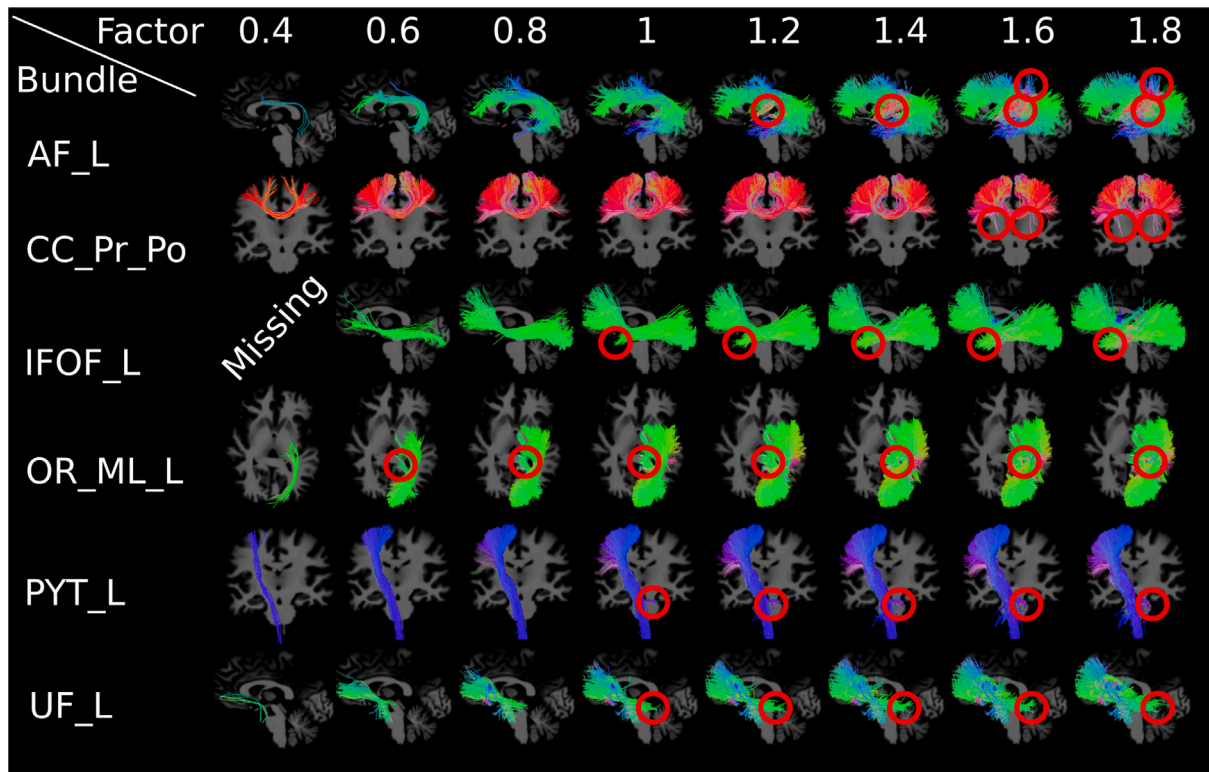


Fig. 4. Effect of various factors based on automated thresholds found with ROC curves analysis where a factor of 1 represents the original thresholds. Factors below 1 mean that original thresholds were reduced, while factors above 1 mean that original thresholds were increased. Red circles indicate false positive streamlines according to atlas ideal bundles presented in Fig. 2. Missing means that the method did not extract any streamline for a particular bundle.

and FINTA-multibundle output, c.f. Supplemental section A.3 for more details) is specified alongside the maximum total number of streamlines to use. The decision to estimate the latent space PDF with streamlines from the subject bundles and the atlas bundles was motivated by the fact that ideal bundles are built to fully cover the target anatomical region. Therefore, missing streamlines from the original bundles can be recovered if a portion of the streamline seeds (Legarreta et al., 2023) for the latent space PDF estimation is based on the atlas bundles. Then, all streamlines used to estimate the PDF are embedded with the encoder. The PDF is empirically estimated using a Parzen estimator (Bishop, 2006) over all embedded points with a Gaussian kernel. The kernel bandwidth is automatically estimated using Silverman's rule of thumb (Silverman, 1986).

As described in Legarreta et al. (2023), Painchaud et al. (2020), since streamline latent representations live in a locally linear manifold, rejection sampling (RS) (Bishop, 2006) can be used to sample new data from an empirically estimated PDF. Thus, giving a set of latent vectors \mathcal{Z} with $\mathbf{z}_i \in \mathbb{R}^d, \mathbf{z}_i \in \mathcal{Z}$ and $\mathbf{z}_i = p_\theta(S_i)$, our goal is to estimate a new set of latent vectors $\mathcal{Z}' \not\subset \mathcal{Z}$ where the new PDF of estimated vectors $P(\mathbf{z}')$ is close to $P(\mathbf{z})$. As described previously, the PDF of $P(\mathbf{z})$ is unknown, and we estimate it with a Parzen estimator. Because $P(\mathbf{z})$ is difficult to sample, we sample an easier PDF $Q(\mathbf{z})$. In our case, $Q(\mathbf{z})$ is a mixture of Gaussians estimated with the expectation-maximization algorithm (Bishop, 2006). The RS procedure works as follows: one first generate a random sample from $Q(\mathbf{z})$ as well as a random number u_0 i.i.d. of a uniform distribution between $[0, KQ(\mathbf{z})]$ where K is a constant. If $u_0 > P(\mathbf{z})$, the sample is rejected, otherwise it is accepted.

Finally, the accepted vectors are decoded to generate new streamlines. As described in GESTA (Legarreta et al., 2023), those streamlines need to fit certain anatomical constraints. Therefore, we adopt the four proposed constraints — namely a length range, a WM coverage, a maximum curving angle, and the local streamline orientation to fODF peak angle (c.f. Supplemental section A.2). Furthermore, we trim off the vertices at each end of generated streamlines that overshoots the gray

matter with a WM mask. The final bundles used for evaluation are the simple concatenation (without further filtering) of FINTA-multibundle and GESTA-gmm bundles.

2.4.1. Number of generated streamlines

To assess the capacity of GESTA-gmm to generate diverse streamlines that “fill” the anatomy well, we produced saturation curves (Gauvin, 2016; Rheault, 2020), which plot the bundle volume with respect to the streamline count. We expect that if adding new streamlines does not affect the bundle volume significantly, the bundle is saturated, i.e., it is “well filled”. We analyzed saturation curves (Gauvin, 2016; Rheault, 2020) of bundle volumes from the GESTA-gmm process compared to the streamline count over 5 bundles based on a whole-brain tractogram of 1 million streamlines. To generate such saturation curves, we produced bundles with 25,000 streamlines **post GESTA-gmm filtering**. After, we logarithmically randomly sampled with replacement 50 subsampled bundles from each bundle. The volume was then computed for each sub-bundle. Fig. 5 presents the saturation curves for those 5 bundles — namely the AF_L, the UF_L, the CC_Pr_Po, the IFOF_L, and the PYT_L. We see that, after 5000 generated streamlines, each bundle volume starts to saturate and after 15,000, no real volume gain is obtained from the generative process. Therefore, we fixed the number of latent-sampled vectors to 25,000 prior to decoding in our experiments, where about half the generated bundles had a final count of more than 5000 streamlines. Finally, bundles with less final streamlines were generally caused, not by a wrongful generative process, but by imperfect WM masks or bundle atlas alignments with the underlying standard space anatomy.

2.5. Reliability evaluation metrics

FIESTA's test-retest reliability was assessed using the *MyeloInferno* dataset according to five overlap metrics — namely (i) the voxel-wise Dice coefficient, (ii) the weighted voxel-wise Dice coefficient, (iii)

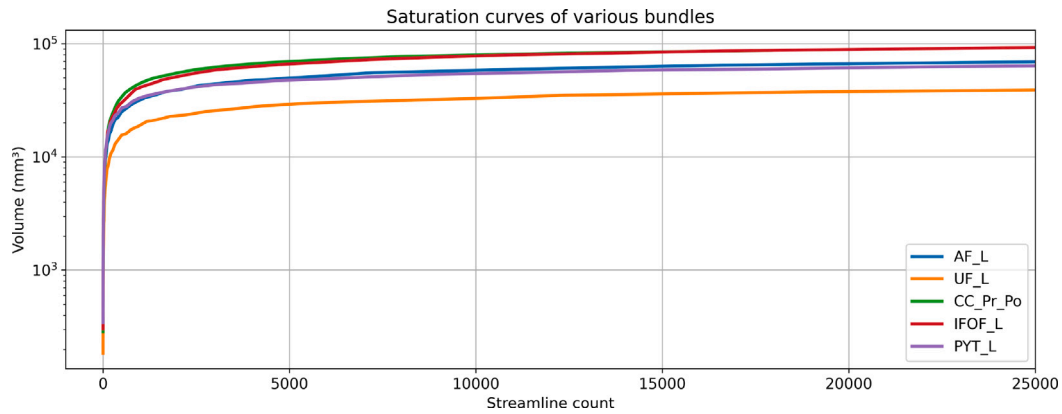


Fig. 5. Saturation curves of various bundles generated by GESTA-gmm showing that between 5000 and 15,000 is a reasonable amount of generated streamlines to cover the entirety of the bundle. After 15,000 streamlines, the volume of each bundle does not increase significantly.

the voxel-wise bundle adjacency, (iv) the streamline-wise bundle adjacency, and (v) the streamline density correlation (Rheault et al., 2020a; Cousineau et al., 2017; Dice, 1945). The Dice and the weighted Dice coefficients are overlap measures between 2 bundles varying from 0 to 1 where 1 means a perfect overlap. The weighted Dice score is weighted by the number of streamlines in each voxel. Bundle adjacency measures are distance metrics, similar to the Hausdorff distance (Rockafellar and Wets, 2009), which yield an average distance between 2 bundles. Metrics are reported in mm and, while there is no cap for those metrics, a value of 0 mm means that there is no distance between 2 bundles. Finally, the streamline density correlation yields a Pearson correlation between 2 bundle density maps going from 0 to 1 with 1 meaning that those 2 bundles are perfectly correlated. Intra-subject tractography variability was assessed pairwise, i.e., for a total of J acquisitions per subject, metrics were computed between images $[1 \rightarrow 2, 1 \rightarrow 3, 1 \rightarrow 4, \dots, J-1 \rightarrow J]$ equaling to a total of $J(J-1)/2$ comparisons per bundle. Also, to gauge FIESTA's reproducibility we computed the intraclass correlation coefficient (ICC) over volume, and average bundle lengths. Based on Koo and Li (2016) (Koo and Li, 2016), ICC estimates and their 95% confidence intervals were computed based on a two-way mixed-effects, consistency, multiple raters/measurements model (ICC (3,k)). Implementation was done using the Pingouin statistical Python package (v0.5.1) (Vallat, 2018).

2.6. Hyperparameters

The final fixed hyperparameters used for FIESTA were a generative sampling of 25,000 new streamlines based on 10,000 seed points from an atlas/subject bundle ratio of 1:1 followed by the GESTA-gmm filtering process. It should also be noted that, for bundles where FINTA-multibundle was not able to give 50% of the 10,000 required seeds, atlas bundle streamlines were used to compensate.

2.7. Methods' comparison

To have a broad range of comparisons, we benchmarked our method against 5 state-of-the-art bundle segmentation methods — namely *RecoBundles* (RB) (Garyfallidis et al., 2018), *RecoBundlesX* (RBx) (Rheault, 2020), *TractSeg* (Wasserthal et al., 2018a,b, 2019, 2020), *XTRACT* (Warrington et al., 2020), and *White Matter Analysis* (WMA) (O'Donnell and Westin, 2007; O'Donnell et al., 2012; Zhang et al., 2018). We also compared FINTA-multibundle against FIESTA (FINTA-multibundle + GESTA-gmm) to see the effect of generative sampling on bundles' reliability. While *RecoBundles*, *RecoBundlesX*, WMA, FINTA-multibundle, and FIESTA used the whole-brain tractogram from *TractoFlow*, we used the default tracking algorithm for *TractSeg* and *XTRACT*. Also, as *XTRACT* software yields bundle density maps, the evaluation was

done on thresholded maps with a value of 0.5% of the highest bundle value. WMA whole-brain tractograms had to be downsampled to 500,000 streamlines due to algorithmic computational limitations. In comparison, *RecoBundles*, *RecoBundlesX*, FINTA-multibundle, and FIESTA whole-brain tractograms were not downsampled and used the ~4M streamlines yielded by *TractoFlow*. A list of the main parameters for each method is presented in Supplemental section A.8 alongside their approximate computation time for a standard desktop computer.

Since the 5 bundle segmentation methods against which FIESTA is compared do not have the same bundle names and/or definitions, as a matter of fairness, only the similar bundles from all methods were compared together. Therefore, 27 similar bundles are kept for the analysis since some bundles in *TractSeg*, WMA, and *XTRACT* had completely different definitions than that of FIESTA. All compared bundles from each method can be found in Supplemental section A.4. It should be noted that WMA misses the SCP bundle, while *XTRACT* only shared 17 similar bundles compared to the other methods. Thus, the global average results were computed on those 17 bundles, whilst individual results were computed on the 27 bundles shared amongst the method.

2.8. FIESTA qualitative evaluation

To gain a better intuition into the different tested methods, we conducted a qualitative analysis on a random subject from the *MyeloInferno-HC-TR* subset. We qualitatively compared various bundles from the different benchmarked methods in order to understand the pitfalls and the quality of each one of them. As one subject is not sufficient to conclude on the bundles' anatomical accuracy, we conducted a more extensive qualitative assessment using the dMRIQCpy toolbox (Theaud, 2022) on the 95 individual images of the *MyeloInferno-MS-TR* dataset. The qualitative evaluation was conducted using a 3 score classification following the user guide of dMRIQCpy on the 27 similar bundles shared amongst methods. The 3 possible scores were either **pass**, **warning**, or **fail**. For each bundle, or a pair of bundles if applicable, a **pass** score was given when the bundle clearly respected the anatomical definition of the atlas, overlapped the correct brain region, and was located in the expected hemispheric and anatomical region. A **warning** score was given if the bundles respected every previous condition but were sparsely populated. Finally, a **fail** score was given if any of the previous three conditions were not respected. We decided to do this test on a cohort of subjects with a neurological disease as a good test to understand the behavior of FIESTA on disease subjects.

2.9. Generated fibers faithfulness

To investigate the impact of generative sampling on tractography and the faithfulness of generated streamlines, we conducted two experiments. We used the 24 acquisitions from the *MyeloInferno-HC* subset

and the 21 from the ADNI-HC subset. All selected subjects were healthy and had no diagnosed neurological disease. In the first experiment, we aimed to evaluate the effectiveness of the 4 anatomical constraints, especially the constraint on the diffusion signal. Since we chose an atlas of bundles made of streamlines from young adults, we assumed that the number of streamlines rejected based on the 4 anatomical constraints should be greater on older adults than on younger ones. Two reasons pushes us to this hypothesis. First, because the MNI template brain was constructed based on young adults (aged 18.5–43.5 years) (Fonov et al., 2011), the registration needed to bring streamlines from the native space to the template space for FIESTA to work properly should be more accurate on younger brains as brain structures are more similar. Second, as the streamlines from the atlas used for our tests were constructed from young adults, their shape should be closer to streamlines from young adult brains. Therefore, when such data is used to seed part of the autoencoder latent space for the GESTA-gmm module, the estimated probability distribution should be closer to that of younger subjects. Thus, while sampling into this distribution, a greater percentage of generated streamlines should be accepted after passing the 4 constraints on young subjects.

The goal of the second experiment is to assess the impact of GESTA-gmm on the streamline-wise bundle adjacency metric between the 2 healthy populations, as this metric can be interpreted as how far 2 populations of streamlines are from each other. The goal is to avoid having a substantial closer streamline-wise bundle adjacency between young and old bundles. Thus, for every subject in the *MyeloInferno-HC* population, we computed the streamline-wise bundle adjacency with every bundle of every subject in the ADNI-HC 21 subjects population. We repeated the experiment once on the bundles segmented by the FINTA-multibundle module, and once on the final output of FIESTA comprising the bundles from the FINTA-multibundle module and the GESTA-gmm module. As the number of streamlines will increase using the FIESTA output compared to the FINTA-multibundle output, we should expect a slight improvement in the streamline-wise bundle-adjacency, but it must not be substantial.

2.10. Performance on disease cohorts

In order to grasp the performance of FIESTA on disease subjects, we extended the test–retest analysis to subjects with neurodegenerative diseases. Therefore, the ADNI-TR, PPMI, and *MyeloInferno-MS-TR* datasets were used with their multiple time-points per subject. ADNI-TR subjects used for this analysis included a mixture of normal aging, MCI, and AD subjects. PPMI subjects were only composed of subjects with Parkinson's disease (PD). *MyeloInferno-MS-TR* contained subjects with multiple sclerosis disease. All three datasets contain multiple time-points, with 5 for ADNI-TR and *MyeloInferno-MS-TR*, and 3 for PPMI. Thus, we repeated the reliability evaluation on those 3 datasets following the same procedure presented in Section 2.5.

3. Results

3.1. Qualitative analysis

To gain insight into the different methods, we conducted a qualitative analysis on a random subject from the *MyeloInferno-HC-TR* dataset. Fig. 6 lays out the same bundles across each bundle segmentation method for a given subject. First, it is qualitatively seen that FIESTA's results (which contains generated streamlines), seems to have the best hemispheric homogeneity. This is particularly true for the PYT and the CC_Pr_Po bundles. *TractSeg* misses some WM streamlines on the left PYT and the CC_Pr_Po bundles. Otherwise, *TractSeg* and WMA seem to output cleaned bundles with few spurious streamlines. The major drawback of WMA comes from its inability to process the original tractograms due to memory issues (~4M streamlines in the *MyeloInferno* dataset for each whole-brain tractogram). *RecoBundles* and

RecoBundlesX seem to be the least effective methods. For *RecoBundles*, the left UF is almost empty, while the left ICP is missing. The same goes for *RecoBundlesX* with a completely missing ICP (left and right) and a missing left UF. Also, comparing their SCP to that from the atlas (c.f. Fig. 2), it is seen that several of those streamlines belong to the medulla, which were not included in the atlas bundles. *XTRACT* retrieves almost all streamlines for each bundle. Unfortunately, *XTRACT* has a high sensitivity but a low specificity for some bundles such as the AF, the UF, and the OR. Finally, FINTA-multibundle clusters correctly all bundles, but with an inhomogeneous fanning for the PYT. The only qualitative issue found with FIESTA and FINTA-multibundle is seen on the IFOF. Indeed, some outlier streamlines seem to appear on the inferior part of the bundle.

Concerning the qualitative analysis made on 95 acquisitions of the *MyeloInferno-MS-TR* dataset using the dMRIQCpy (Theaud, 2022) toolbox, all the 27 evaluated bundles over the whole dataset were given a score of **pass**. Such results confirm that, qualitatively, FIESTA performs well on a disease cohort where all bundles are fully populated, located in the right hemisphere and in the correct anatomical region. During this exercise, the shape of all bundles were compared to the shape of the bundle in the PAWM atlas. All bundles presented a qualitative similar shape to the atlas of bundles used for the development of FIESTA.

3.2. Reliability analysis

Table 1 presents the results on the *MyeloInferno* dataset of the average similarity metrics for each method benchmarked over 17 similar bundles, all shared among the different methods. Quantitative average results show that FIESTA yielded the best Dice, voxel-wise Dice, voxel-wise bundle adjacency, and streamline density correlation with 0.74 ± 0.08 , 0.96 ± 0.04 , 0.43 ± 0.21 mm, and 0.85 ± 0.13 respectively, whereas FINTA-multibundle obtained the best streamline-wise bundle adjacency with a value of 5.02 ± 0.54 mm. The voxel-wise bundle adjacency shows that, on average, two binary masks of the same bundle from the same subject are ~0.5 mm apart using FIESTA. For each metric where FIESTA did not come up first, it was in the margin of error of the first one. *TractSeg* yielded the best ICC length and ICC volume with values of with 0.79 ± 0.05 and 0.82 ± 0.06 , respectively, whereas FIESTA yielded 0.79 ± 0.06 and 0.81 ± 0.07 , respectively. FINTA-multibundle yielded the best streamline-wise bundle adjacency (5.02 ± 0.54 mm), whilst FIESTA yielded 5.24 ± 0.44 mm. Finally, according to Koo and Li (2016) (Koo and Li, 2016), ICCs yielded by FIESTA show that the method has a good reliability.

Fig. 7 presents per-bundle similarity metric results of the 27 benchmarked bundles. It should be noted that *XTRACT* only presents 17 bundles and WMA 25 (the SCP is not present). Except for the streamline-wise bundle adjacency, where FINTA-multibundle is clearly superior, FIESTA is globally more reliable with better overall scores. The improvement is particularly important over small and hard-to-track bundles such as the ICP, the SCP, the OR_ML or the UF. More specifically, the streamline density correlation, the Dice score and the weighted Dice score (Fig. 7a, 7b, and 7c), show that FIESTA yields higher within-bundle similarity. Fig. 7a shows that the streamline density between different acquisitions is systematically more correlated when using FIESTA than with the other methods. The voxel-wise bundle adjacency of FIESTA shows that the average distance between 2 bundle masks of the same subject is ~0.5 mm. Knowing that the voxel size of the *MyeloInferno* dataset is 2 mm isotropic, this underlines how reproducible FIESTA is. Finally, Fig. 7e shows that, depending on the studied bundle, the *TractSeg* streamline-wise bundle adjacency is close to FIESTA and FINTA-multibundle, with many bundles surpassing the proposed method, whilst never surpassing FINTA-multibundle.

According to Fig. 7e, FINTA-multibundle shows the closest scores to FIESTA with *RecoBundles*, *RecoBundlesX*, and WMA's being less reproducible. In comparison, *RecoBundles*, *RecoBundlesX*, FINTA-multibundle, and FIESTA had about 4,000,000 input streamlines. This

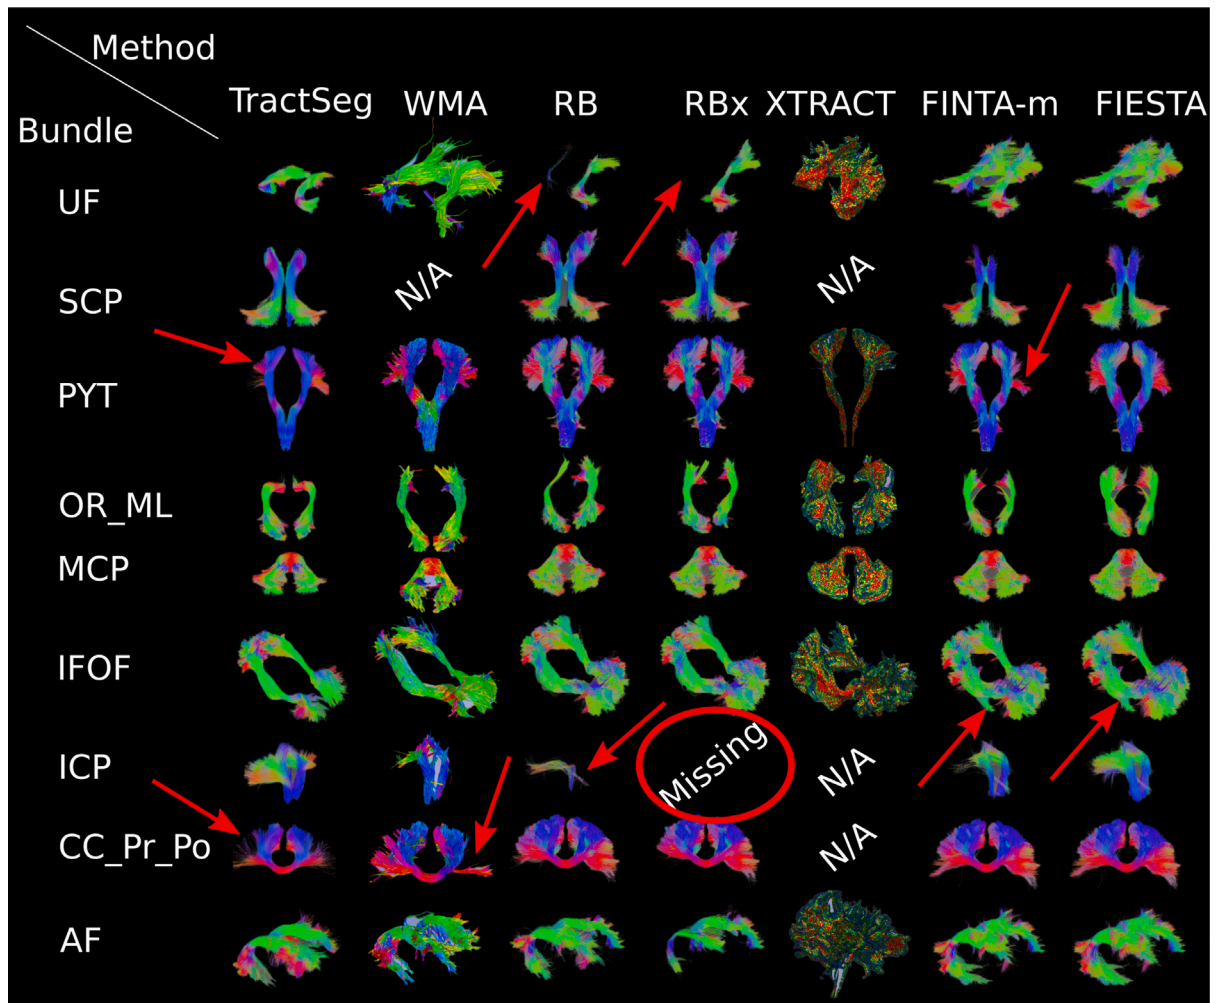


Fig. 6. Qualitative results for a subset of bundles from the different benchmarked methods from one *MyeloInferno* subject. *N/A* means that the bundle was not included in the method, while *Missing* means that the method failed to extract that bundle. Lateral bundles (in both hemispheres) are shown jointly: UF, SCP, PYT, OR_ML, IFOF, ICP, and AF. *XTRACT* bundles show density maps (thresholded at 0.5%). *FINTA-m* stands for *FINTA-multibundle*. Red arrows point to suspected missing streamlines or wrongly bundled.

Table 1

Average (\pm SD) similarity metric values and ICC for the 17 studied bundles, shared across different state-of-the-art methods for the *MyeloInferno* dataset. Up and down arrows mean that higher and lower are better respectively for the studied metric. V. and S. stand for Voxels and Streamlines. L. stands for Length. BA stands for Bundle Adjacency. The bundle adjacency streamlines and the ICC length were not computed for *XTRACT* because it outputs streamline density maps instead of WM streamlines for each bundle. *FINTA-m* stands for *FINTA-multibundle*.

17 Bundles	TractSeg	RB	RBx	WMA	XTRACT	FINTA-m	FIESTA
Dice V. (\uparrow)	0.65 \pm 0.10	0.56 \pm 0.22	0.58 \pm 0.23	0.56 \pm 0.11	0.64 \pm 0.15	0.71 \pm 0.09	0.74 \pm 0.08
W-Dice V. (\uparrow)	0.85 \pm 0.09	0.78 \pm 0.29	0.79 \pm 0.29	0.78 \pm 0.14	0.84 \pm 0.14	0.94 \pm 0.06	0.96 \pm 0.04
B. A. V. (mm) (\downarrow)	0.70 \pm 0.45	1.65 \pm 2.03	1.46 \pm 2.01	0.88 \pm 0.39	0.89 \pm 0.63	0.51 \pm 0.27	0.43 \pm 0.21
B. A. S. (mm) (\downarrow)	5.37 \pm 0.92	7.82 \pm 3.01	6.94 \pm 2.76	6.07 \pm 0.65	N/A	5.02 \pm 0.54	5.24 \pm 0.44
Density Corr. (\uparrow)	0.64 \pm 0.18	0.60 \pm 0.30	0.64 \pm 0.30	0.58 \pm 0.22	0.63 \pm 0.21	0.81 \pm 0.15	0.85 \pm 0.13
ICC L. (\uparrow)	0.79 \pm 0.05	0.55 \pm 0.28	0.67 \pm 0.19	0.78 \pm 0.08	N/A	0.78 \pm 0.06	0.79 \pm 0.06
ICC V. (\uparrow)	0.82 \pm 0.06	0.24 \pm 1.54	0.70 \pm 0.13	0.73 \pm 0.15	0.67 \pm 0.17	0.80 \pm 0.09	0.81 \pm 0.07

might explain why WMA seems to perform poorly, as we had to down sample the input whole-brain tractogram due to the computational cost of the method.

According to the charts in Fig. 7b and Fig. 7c, FIESTA is less effective for the ICP bundle. To better understand this behavior, we further investigated the case. Fig. 8 shows the left ICP produced by FIESTA on the subject with the poorest similarity metrics. We see that the variability for this particular bundle is explained by the positioning of the MRI field of view, which seems to cut at various points in the brainstem. Therefore, the starting zones of the ICP are variable. Despite this bundle variability, Fig. 8 shows that FIESTA is robust to such starting zone variations.

To complement the previous results, Fig. 9 shows the reliability scores with respect to bundle-wise metrics — namely volume and average length. Following Koo and Li (2016) (Koo and Li, 2016) guidelines, an ICC between 0.75 and 0.9 and above 0.9 are respectively indicative of good and excellent reliability. Fig. 9a presents the ICC of the bundle average length, which provides insight into the reliability of each method to obtain the average bundle length values. FIESTA outperforms *RecoBundles* and *RecoBundlesX* in most bundles. WMA and *TractSeg* are more competitive than anticipated. Careful analysis reveals that although average overlap metrics for those methods are much worse than FIESTA, they yield streamlines with consistent lengths across various time points. It can be seen in Fig. 7e, that *TractSeg* seems

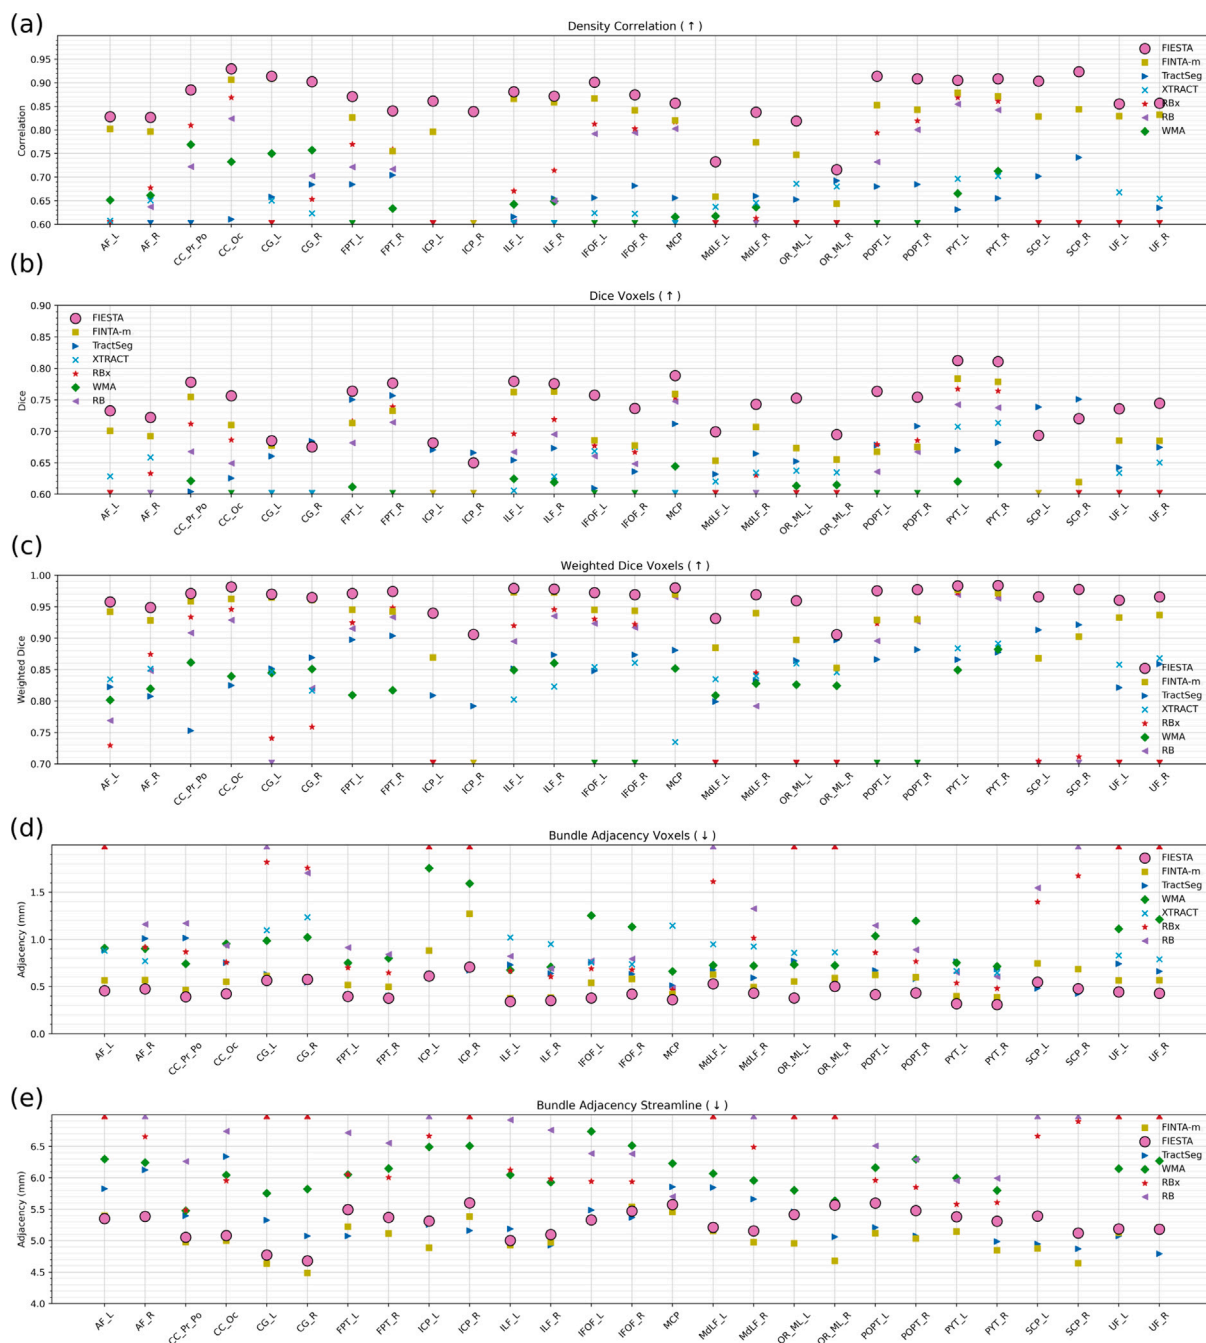


Fig. 7. Average per-bundle overlap metrics for the *MyeloInferno* dataset. Each metric was computed over 5 acquisitions from 18 subjects registered to the MNI152 template with per subject pairwise comparison yielding 10 comparisons per subject. For visualization purposes, only the 27 most similar bundles across methods are shown. It should be noted that the SCP was not present in the bundles from WMA. *XTRACT* also misses all CC bundles, the FPT, the ICP, the POPT, and the SCP. (a) Streamline Density Correlation, (b) Voxel-Wise Dice Score Coefficient, (c) Weighted Voxel-Wise Dice Score Coefficient, (d) Voxel-Wise Bundle Adjacency, and (e) Streamline-Wise Bundle Adjacency. The streamline-wise bundle adjacency was not computed for *XTRACT* because the software outputs streamline density volume maps instead of WM streamlines for each bundle. The legend of each chart presents, from top to bottom, the method with the best average metric value according to Table 1. Outlier values are clipped to the bottom or top of each chart and are displayed as small triangles. FINTA-m stands for FINTA-multibundle.

to produce a high streamline-wise bundle adjacency. The per-bundle volume measurement results in Fig. 9b shows a similar trend with *TractSeg*, *FINTA-multibundle*, and *FIESTA* having comparable results, and *TractSeg* and *FINTA-multibundle* outperforming *FIESTA* for some bundles. This might be explained by the fact that *TractSeg* uses probabilistic bundle specific tracking instead of whole-brain tracking. But although the *TractSeg* bundle-wise volume is stable through different time-points, its test-retest similarity metrics are much worse than *FIESTA*. Finally, even if some methods seem to yield higher ICC values

than *FIESTA*, most *FIESTA*'s ICC were higher than 0.75, thus producing good reliability measurement (c.f. Table 1).

Table 2 presents the results of *FIESTA* on three disease cohorts — namely ADNI-TR, PPMI, and *MyeloInferno-MS-TR*. Results show that performances on those cohorts are similar to what were obtained on healthy subjects, presented in Table 1. Thus, it shows the potential of using *FIESTA* on a wide range of population. Also, results presented in Table 2 shows the usefulness of using *GESTA-gmm* to improve most reliability metrics. Finally, a subset of bundles (CC_Pr_Po and IFOF) alongside the segmented lesions are presented in Fig. 10 from one




Fig. 8. Visualization of FIESTA's ICP across 5 acquisitions of the subject with the largest variability in scores. Each color stands for a different acquisition. The variability is mostly explained by the MRI field of view, and not by FIESTA's natural variability. In that case, the ICP starts at various points of the medulla (or the pons for the second image) because the original image is cut at various points in the brainstem. Red arrows point to the medulla on the left most image.




Fig. 9. Bundle ICCs from the 27 benchmarked bundles, where (a) is the ICC based on bundle average length and (b) is the ICC based on the per-bundle voxel volume. It should be noted that the ICC length was not computed for *XTRACT* because the software outputs streamline density volume maps instead of WM streamlines for each bundle. Green dotted lines indicate good reliability (0.75) and green hard lines indicate excellent reliability (0.9). FINTA-m stands for FINTA-multibundle.

Table 2

Results of the performance of FIESTA on three disease cohorts — namely ADNI-TR (AD), PPMI (PD), and MyeloInferno-MS-TR (MS) on the 27 benchmarked bundles; **Left:** Reliability results of the FINTA-multibundle module; **Right:** Reliability results of the FIESTA segmented bundles. FINTA-m stands for FINTA-multibundle.

27 bundles	AD FINTA-m	AD FIESTA	PD FINTA-m	PD FIESTA	MS FINTA-m	MS FIESTA
Dice V. (↑)	0.69 ± 0.10	0.73 ± 0.08	0.68 ± 0.09	0.74 ± 0.07	0.70 ± 0.10	0.76 ± 0.05
W-Dice V. (↑)	0.91 ± 0.12	0.95 ± 0.09	0.91 ± 0.10	0.96 ± 0.06	0.93 ± 0.09	0.97 ± 0.04
B. A. V. (mm) (↓)	0.48 ± 0.23	0.43 ± 0.21	0.53 ± 0.29	0.40 ± 0.23	0.49 ± 0.30	0.37 ± 0.14
B. A. S. (mm) (↓)	4.63 ± 0.51	5.40 ± 0.43	5.35 ± 0.48	5.35 ± 0.49	5.23 ± 0.48	5.27 ± 0.38
Density Corr. (↑)	0.82 ± 0.19	0.90 ± 0.14	0.83 ± 0.16	0.91 ± 0.11	0.87 ± 0.12	0.93 ± 0.08
ICC L. (↑)	0.95 ± 0.03	0.94 ± 0.03	0.92 ± 0.03	0.92 ± 0.03	0.92 ± 0.08	0.94 ± 0.05
ICC V. (↑)	0.95 ± 0.04	0.96 ± 0.02	0.90 ± 0.06	0.91 ± 0.05	0.92 ± 0.10	0.94 ± 0.08

subject of the *MyeloInferno-MS-TR* dataset showing how the bundle are able to cross MS lesions as long as the diffusion signal is not altered in those regions.

3.3. GESTA-gmm generated streamlines faithfulness analysis

Fig. 11 presents the results showing the ratio between the streamline count for each bundle of the generated streamlines ($N = 25,000$) before the four anatomical constraints filtering and after. Thus, a ratio of 0.5 indicates that, for a particular bundle, for each remaining streamline at the end of the filtering process, 2 must be generated by the rejection

sampling process. **Fig. 11** shows that our initial hypothesis, stipulating that a greater percentage of generated streamlines should be accepted after passing the 4 constraints on young subjects, explains such results. Indeed, for all bundles, more streamlines passed the filtering process for the young population in comparison to the old population. Such results suggest that the remaining streamlines are faithful to the underlying anatomy and respect the diffusion signal for both the young and old subjects.

Fig. 12 shows the results of the streamline-wise bundle adjacency between the young subjects and the old subjects for each bundle. We see that, as anticipated, the metric improves with the usage of




Fig. 10. Lesions from one subject in the *MyeloInferno-MS-TR* dataset; **Top**: Lesions in the CC_Pr_Po bundle; **Bottom**: Lesions in the left and right IFOF bundles.




Fig. 11. Ratio between the streamline count for each bundle before and after the four anatomical constraint filtering. * indicates a p value < 0.05 , n.s. means non statistically significant difference.




Fig. 12. Streamline-wise, bundle adjacency for the FIESTA and the FINTA-multibundle bundles between the young and the old healthy subjects from the *MyeloInferno-HC* and ADNI-HC datasets. FINTA-m stands for FINTA-multibundle.

GESTA-gmm. On average, the streamline-wise bundle adjacency for FINTA-multibundle was 6.08 ± 0.65 mm, whilst it was 5.98 ± 0.48 mm when using FIESTA. However, such a small augmentation of 0.1 mm, in comparison to the smallest ADNI voxel size (1.36 mm), indicates that the generated streamlines that pass the 4 anatomical constraints, even if the atlas bundles and anatomy were based on a younger population, stay faithful to the current subject anatomy.

4. Discussion

Automatic bundle segmentation is an essential processing step in tractography to isolate specific white matter pathways. In this study, we presented a novel semi-supervised autoencoder-based automatic bundle segmentation method called FIESTA (*FibEr Segmentation in Tractography using Autoencoders*). We compared FIESTA's reliability to state-of-the-art

automatic bundle segmentation methods — namely *RecoBundles* (RB) (Garyfallidis et al., 2018), *RecoBundlesX* (RBx) (Rheault, 2020), *TractSeg* (Wasserthal et al., 2018a,b, 2019, 2020), *XTRACT* (Warrington et al., 2020), and *White Matter Analysis* (WMA) (O'Donnell and Westin, 2007; O'Donnell et al., 2012; Zhang et al., 2018). We found that methods such as *RecoBundles* and *RecoBundlesX* are prone to large intra and inter-subject variability. Bundle definitions from *TractSeg* are fixed, requiring a complete retraining of the system for a different set of bundles. From a practical point of view, it would be hardly doable to change a specific bundle definition overnight, as one would need to create a completely new training set to train all three U-Nets. WMA is too computationally expensive. We had to downsample input tractograms to 500,000 streamlines. *XTRACT* ROIs seem to yield bundles with a low specificity and are hardly editable because they are defined by 5 anatomical experts, which would be time-consuming to change. Also, we found that FIESTA is robust to these problems present in current state-of-the-art bundle segmentation methods. Moreover, we showed that FIESTA is robust and reliable on various kinds of brain populations, such as disease cohorts, as long as the diffusion signal is not altered (or slightly altered) by such diseases.

As mentioned, the main goal of this work is to assess the reliability of FIESTA alongside state-of-the-art automatic bundle segmentation methods. As per many segmentation algorithms in biomedical image analysis, we claim that reliability analysis is as important as standard comparisons to some ground truth (or silver standard). It gives a better insight on the robustness of the method, even if the comparison to the ground truth is poor. Therefore, both validation methods are as important, but the reliability is often rarely reported. Results presented in Table 1 show that FIESTA yields reliable bundles suggested by different similarity metrics (Dice, weighted-Dice, voxel-wise bundle adjacency, streamline-wise bundle adjacency and streamline density correlation) in a test-retest manner over different dMRI time points. Quantitative average results showed that FIESTA yielded the best Dice, voxel-wise bundle adjacency, and streamline density correlation or is in the margin of error of the best values (streamline-wise bundle adjacency, ICC length, and ICC volume). The voxel-wise bundle adjacency shows that, on average, two binary masks of the same bundle from the same subject are ~0.5 mm apart using FIESTA. This value is minimal considering that the acquisition resolution of the dataset used is 2 mm isotropic. Also, variability induced by registration, necessary for similarity metric computations, did not overly affect the voxel-wise bundle adjacency average distance. Higher density correlations obtained by FIESTA might be explained, in part, by the use of generative sampling. Interestingly, a closer analysis of the source of variability in FIESTA reveals that it is mainly due to the MRI field of view positioning (c.f. Fig. 8).

ICC results are not as consistent as it is the case for similarity metrics. It is seen that bundles from *TractSeg* are more reliable for volume and average length. This is probably due to the fact that *TractSeg* uses probabilistic bundle-specific tracking and not whole-brain tracking for bundle segmentation later on. Thus, it probably yields more consistent metrics at various time points depending on the studied bundle. Following those results, we claim that the interpretation of any ICCs, especially in biomedical image analysis, needs to be interpreted cautiously. Indeed, a bad algorithm could, for example, produce, by chance, bundles with the same volume without being in the same anatomical region. Thus, it would yield high ICCs, with low overlap values. Therefore, the interpretability of an ICC must always be paired, within the context of a tractography analysis, with similarity metrics and a good qualitative analysis.

Analysis of the generative sampling scheme shows that generated streamlines remaining after the filtering process, respect the subject anatomy and do not generalize to a generic ensemble of streamlines (c.f. Figs. 11 and 12). Such behavior could be assumed to happen due to the fact that part of the latent space is seeded using generic streamlines based on the PAWM atlas. However, the constraint on the diffusion signal removes all streamlines that are not faithful to one's subject

anatomy. Such constraint stipulates that any streamline generated using the generative process that has a local orientation angle to fODF peak greater than a certain angle value (i.e., 30°) ought to be discarded. Therefore, we would argue that streamlines that are kept after this step should be closely similar to a streamline generated from a traditional tractography process. Another way to see this filtering step is, instead of tracking the streamline on the diffusion signal, we proposed streamlines that are likely to be close to the subject anatomy, but we only keep the ones that respect the diffusion signal.

FIESTA is not a method without flaws. First, we see that the use of generative sampling does worsen the streamline-wise bundle adjacency. We hypothesized that such behavior might be due to the usage of atlas bundles as seeds for the latent space combined with the lossy process of autoencoders. Thus, combining generative sampling with generic seed points and another kind of tractography process might produce variations that increase the streamline-wise bundle adjacency. Such behavior is currently under investigation, and generative sampling filtering parameters might need further adjustment. Second, the threshold calibration, even though it can be done quickly compared to other state-of-the-art methods, is probably the most time-consuming process. A solution would be to get rid of the threshold method by including a new *bundle* in the atlas with all whole-brain implausible WM streamlines. Therefore, streamlines closer to the implausible *bundle* would be classified as implausible. New bundle definitions could easily be implemented, as only an ideal bundle in a standard space would be needed for the change to be effective, without the need to find/calibrate any threshold. As WMA uses this method for implausible streamlines, it might be pertinent to test. Unfortunately, this method also has its limitations, as defining implausible streamlines is an open and hard problem to solve. Also, even if we presented a solution to find the *near-optimal* threshold in Section 2.3.1, the most easy way to find the desired threshold, whilst being the least sophisticated method, is just to test a few thresholds manually, starting completely randomly and slowly converging around the desired final values. Moreover, the generative sampling method takes a long time for certain bundles. It would be interesting to include per-bundle generative sampling parameters to optimize the pipeline (Legarreta et al., 2023). Furthermore, FIESTA was not developed to work properly on heavily impaired and deformed brain anatomy (such as tumor). Such, behavior was not characterized, and we are currently investigating the best approach for such specific brain deformed anatomy. However, we showed that FIESTA is robust to small white matter lesions as long as the diffusion signal is not (or slightly) altered. Also, we showed that the brain deformation on normal aging population does not affect the performances of FIESTA. Finally, we can see in Table A.5 that GESTA-gmm is still a slow process, even though it was improved for this pipeline. Thus, more optimization work is needed for the generative process. Even if a method like *TractSeg* is still faster than FIESTA, its rigid framework prevents rapid iterations when new bundles are needed. Here, we demonstrated that it is possible to have a bundle segmentation pipeline that works on raw tractograms and is more reliable in a test-retest manner than state-of-the-art bundle segmentation methods.

This work mainly focuses on the bundle segmentation of long-range white matter streamlines, that connect distant cortical areas of the brain. We did not assess the behavior of FIESTA on short association fibers (U-Fibers) (Shastin et al., 2022) that connect adjacent cortical regions. Future works will include the evaluation of FIESTA on such fibers.

5. Conclusion

We presented a new tractography bundle segmentation pipeline named FIESTA (*FibEr Segmentation in Tractography using Autoencoders*) that leverages the power of unsupervised and self-supervised learning and that is intended as a solution to the limitations of current state-of-the-art bundle segmentation pipelines. We showed that FIESTA lets

users easily and rapidly edit its bundle definitions, and is highly reliable in test–retest. FIESTA is more reliable than other state-of-the-art methods such as *TractSeg*, *RecoBundles*, *RecoBundlesX*, *XTRACT*, and *WMA*. We believe FIESTA reaches an optimal compromise between computational burden, ease-of-use, and reconstruction quality and reliability. Thus, FIESTA might be a promising avenue as an easy and reliable way for bundle segmentation in tractography, especially so for hard-to-track bundles.

Ethics statement

Written informed consent was obtained from participants and were recruited following the ethics protocol of the Centre de Recherche du Centre Hospitalier Universitaire de Sherbrooke (Sherbrooke, Canada) (for the MyeloInferno dataset).

CRediT authorship contribution statement

Félix Dumais: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing – original draft, Visualization, Project administration. **Jon Haitz Legarreta:** Conceptualization, Methodology, Software, Writing – review & editing. **Carl Lemaire:** Software, Validation, Writing – review & editing. **Philippe Poulin:** Conceptualization, Writing – review & editing. **François Rheault:** Data curation, Writing – review & editing. **Laurent Petit:** Writing – review & editing. **Muhammed Barakovic:** Writing – review & editing. **Stefano Magon:** Writing – review & editing. **Maxime Descoteaux:** Conceptualization, Funding acquisition, Resources, Supervision, Writing – review & editing. **Pierre-Marc Jodoin:** Conceptualization, Funding acquisition, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

1. Maxime Descoteaux and Pierre-Marc Jodoin report membership and employment with Imeka Solutions inc.
2. Patent #17/337,413 is pending to Imeka Solutions inc. with inventors Jon Haitz Legarreta, Maxime Descoteaux and Pierre-Marc Jodoin.
3. Muhammed Barakovic is an employee of Hays plc and a consultant for F. Hoffmann-La Roche Ltd.
4. Stefano Magon is an employee and shareholder of F. Hoffmann-La Roche Ltd.

Data/code availability

All data used to develop and test FIESTA were acquired from different public (TractoInferno, HCP, PAWM atlas, ADNI, PPMI) and private datasets (MyeloInferno). Data from the MyeloInferno dataset cannot be made public to respect ethics protocol of the institution. FIESTA config files, training data, and the deep learning model are made publicly available at <https://doi.org/10.5281/zenodo.7562790> (Dumais and Legarreta, 2023). URL links for public datasets are:

- TractoInferno: <https://openneuro.org/datasets/ds003900>
- HCP: <http://www.humanconnectomeproject.org/>
- PAWM atlas: <https://doi.org/10.5281/zenodo.5165374>
- ADNI: <https://adni.loni.usc.edu/>
- PPMI: <https://www.ppmi-info.org/>

The source code of the FIESTA pipeline can be found under the GitHub repository <https://github.com/scil-vital/fiesta>. Original scripts and library implemented in the pipeline can be found under the GitHub repository <https://github.com/scil-vital/tractolearn>.

Acknowledgments

This research was conducted in the context of the Acuity-QC group funded by the Fond d'Accélération des Collaborations en Santé (FACS) and Consortium Québécois sur la Découverte du Médicament (CQDM). It was enabled, in part, by support provided by Calcul Québec (www.calculquebec.ca) and Digital Research Alliance of Canada (www.alliancecacan.ca).

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.

Data were provided, in part, by the OpenNeuro database (<https://openneuro.org/datasets/ds003900>). Its accession number is ds003900.

Data collection and sharing for this project was funded, in part, by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

Data used in the preparation of this article were obtained from the Parkinson's Progression Markers Initiative (PPMI) database (www.ppmi-info.org/access-data-specimens/download-data). For up-to-date information on the study, visit www.ppmi-info.org. PPMI – a public–private partnership – is funded by the Michael J. Fox Foundation for Parkinson's Research and funding partners, including [list the full names of all of the PPMI funding partners found on the <https://www.ppmi-info.org/about-ppmi/who-we-are/study-sponsors>].

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.neuroimage.2023.120288>.

References

- Bertò, G., Bullock, D., Astolfi, P., Hayashi, S., Zigiottò, L., Annicchiarico, L., Corsini, F., De Benedictis, A., Sarubbo, S., Pestilli, F., Avesani, P., Olivetti, E., 2021. Classifier, a robust streamline-based linear classifier for white matter bundle segmentation. *NeuroImage* 224, 117402. <http://dx.doi.org/10.1016/j.neuroimage.2020.117402>, URL <https://www.sciencedirect.com/science/article/pii/S1053811920308879>.
- Bishop, C.M., 2006. Pattern Recognition and Machine Learning. In: Information science and statistics, Springer, New York.
- Chen, Y., Zhang, C., Song, Y., Makris, N., Rath, Y., Cai, W., Zhang, F., O'Donnell, L.J., 2021. Deep Fiber Clustering: Anatomically Informed Unsupervised Deep Learning for Fast and Effective White Matter Parcellation. *arXiv:2107.04938* [cs] arXiv: 2107.04938.

- Chen, Y., Zhang, C., Xue, T., Song, Y., Makris, N., Rath, Y., Cai, W., Zhang, F., O'Donnell, L.J., 2022. DFC: Anatomically Informed Fiber Clustering with Self-supervised Deep Learning for Fast and Effective Tractography Parcellation. *arXiv: 2205.00627* [cs].
- Chen, Y., Zhang, C., Xue, T., Song, Y., Makris, N., Rath, Y., Cai, W., Zhang, F., O'Donnell, L.J., 2023. Deep fiber clustering: Anatomically informed fiber clustering with self-supervised deep learning for fast and effective tractography parcellation. *NeuroImage* 273, 120086. <http://dx.doi.org/10.1016/j.neuroimage.2023.120086>, URL <https://www.sciencedirect.com/science/article/pii/S105381192300232X>.
- Christiaens, D., Reiser, M., Dhollander, T., Sunaert, S., Suetens, P., Maes, F., 2015. Global tractography of multi-shell diffusion-weighted imaging data using a multi-tissue model. *NeuroImage* 123, 89–101. <http://dx.doi.org/10.1016/j.neuroimage.2015.08.008>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811915007168>.
- Cousineau, M., Jodoin, P.-M., Garyfallidis, E., Côté, M.-A., Morency, F.C., Rozanski, V., Grand'Maison, M., Bedell, B.J., Descoteaux, M., 2017. A test-retest study on Parkinson's PPMI dataset yields statistically significant white matter fascicles. *NeuroImage: Clinical* 16, 222–233. <http://dx.doi.org/10.1016/j.nicl.2017.07.020>, URL <https://www.sciencedirect.com/science/article/pii/S2213158217301869>.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2007. Regularized, fast, and robust analytical Q-ball imaging. *Magn. Reson. Med.* 58 (3), 497–510. <http://dx.doi.org/10.1002/mrm.21277>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21277>.
- Descoteaux, M., Deriche, R., Knosche, T.R., Anwander, A., 2009. Deterministic and Probabilistic Tractography Based on Complex Fibre Orientation Distributions. *IEEE Trans. Med. Imaging* 28 (2), 269–286. <http://dx.doi.org/10.1109/TMI.2008.2004424>, Conference Name: IEEE Transactions on Medical Imaging.
- Dice, L.R., 1945. Measures of the Amount of Ecologic Association Between Species. *Ecology* 26 (3), 297–302. <http://dx.doi.org/10.2307/1932409>, URL <http://www.jstor.org/stable/1932409>, Publisher: Ecological Society of America.
- Dumais, F., Legarreta, J.H., 2023. FINTA/CINTA/GESTA/FIESTA Datasets. <http://dx.doi.org/10.5281/zenodo.7562790>.
- Edde, M., Theaud, G., Dumont, M., Théberge, A., Valcourt-Caron, A., Gilbert, G., Houde, J.-C., Maltais, L., Rheault, F., Spagnolo, F., Barakovic, M., Magon, S., Descoteaux, M., 2023. High-frequency longitudinal white matter diffusion- and myelin-based MRI database: Reliability and variability. *Hum. Brain Mapping* n/a (n/a), <http://dx.doi.org/10.1002/hbm.26310>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.26310>.
- Fonov, V., Evans, A.C., Botteron, K., Almli, C.R., McKinstry, R.C., Collins, D.L., 2011. Unbiased average age-appropriate atlases for pediatric studies. *NeuroImage* 54 (1), 313–327, URL <http://www.proquest.com/docview/1549926845/abstract/4E0A8BE73F6F4AE9PQ/1>, Num Pages: 313-327 Place: Amsterdam, United States Publisher: Elsevier Limited.
- Fonov, V., Evans, A., McKinstry, R., Almli, C., Collins, D., 2009. Unbiased nonlinear average age-appropriate brain templates from birth to adulthood. *NeuroImage* 47, S102. [http://dx.doi.org/10.1016/S1053-8119\(09\)70884-5](http://dx.doi.org/10.1016/S1053-8119(09)70884-5), URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811909708845>.
- Fowlkes, C., Belongie, S., Chung, F., Malik, J., 2004. Spectral grouping using the nystrom method. *IEEE Trans. Pattern Anal. Mach. Intell.* 26 (2), 214–225. <http://dx.doi.org/10.1109/TPAMI.2004.1262185>, URL <http://ieeexplore.ieee.org/document/1262185/>.
- Garyfallidis, E., Brett, M., Correia, M.M., Williams, G.B., Nimmo-Smith, I., 2012. QuickBundles, a Method for Tractography Simplification. *Front. Neurosci.* 6, 175. <http://dx.doi.org/10.3389/fnins.2012.00175>, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3518823/>.
- Garyfallidis, E., Côté, M.-A., Rheault, F., Descoteaux, M., 2015a. QuickBundlesX: Sequential clustering of millions of streamlines in multiple levels of detail at record execution time. In: ISMRM 2016.
- Garyfallidis, E., Côté, M.-A., Rheault, F., Sidhu, J., Hau, J., Petit, L., Fortin, D., Cunanne, S., Descoteaux, M., 2018. Recognition of white matter bundles using local and global streamline-based registration and clustering. *NeuroImage* 170, 283–295. <http://dx.doi.org/10.1016/j.neuroimage.2017.07.015>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811917305839>.
- Garyfallidis, E., Ocegueda, O., Wassermann, D., Descoteaux, M., 2015b. Robust and efficient linear registration of white-matter fascicles in the space of streamlines. *NeuroImage* 117, 124–140. <http://dx.doi.org/10.1016/j.neuroimage.2015.05.016>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811915003961>.
- Gauvin, A., 2016. Assurance qualité en dissection virtuelle des faisceaux de la matière blanche par tractographie (Master's thesis). Université de Sherbrooke, p. 36.
- Girard, G., Whittingstall, K., Deriche, R., Descoteaux, M., 2014. Towards quantitative connectivity analysis: reducing tractography biases. *NeuroImage* 98, 266–278. <http://dx.doi.org/10.1016/j.neuroimage.2014.04.074>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811914003541>.
- Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., Smith, S.M., Van Essen, D.C., 2016. A multi-modal parcellation of human cerebral cortex. *Nature* 536 (7615), 171–178. <http://dx.doi.org/10.1038/nature18933>, URL <http://www.nature.com/articles/nature18933>.
- Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., Van Essen, D.C., Jenkinson, M., 2013. The minimal preprocessing pipelines for the Human Connectome Project. Mapping the Connectome, *NeuroImage Mapping the Connectome*, 80, 105–124. <http://dx.doi.org/10.1016/j.neuroimage.2013.04.127>, URL <https://www.sciencedirect.com/science/article/pii/S1053811913005053>.
- Goodfellow, I., Bengio, Y., Courville, A., 2016. Deep Learning. The MIT Press.
- Hadsell, R., Chopra, S., LeCun, Y., 2006. Dimensionality Reduction by Learning an Invariant Mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 2 (CVPR'06), vol. 2. IEEE, New York, NY, USA, pp. 1735–1742. <http://dx.doi.org/10.1109/CVPR.2006.100>, URL <http://ieeexplore.ieee.org/document/1640964/>.
- Hinton, G.E., Salakhutdinov, R.R., 2006. Reducing the Dimensionality of Data with Neural Networks. *Science* 313 (5786), 504–507. <http://dx.doi.org/10.1126/science.1127647>, URL <https://www.science.org/doi/10.1126/science.1127647>.
- Jörgens, D., Descoteaux, M., Moreno, R., 2021. Challenges for Tractogram Filtering. In: Özarslan, E., Schultz, T., Zhang, E., Fuster, A. (Eds.), *Anisotropy Across Fields and Scales*. In: Mathematics and Visualization, Springer International Publishing, Cham, pp. 149–168. http://dx.doi.org/10.1007/978-3-030-56215-1_7.
- Koo, T.K., Li, M.Y., 2016. A Guideline of Selecting and Reporting Intraclass Correlation Coefficients for Reliability Research. *J. Chiropractic Med.* 15 (2), 155–163. <http://dx.doi.org/10.1016/j.jcm.2016.02.012>, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4913118/>.
- Kreher, B., Mader, I., Kiselev, V., 2008. Gibbs tracking: A novel approach for the reconstruction of neuronal pathways. *Magn. Reson. Med.* 60 (4), 953–963. <http://dx.doi.org/10.1002/mrm.21749>, URL <http://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21749>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.21749>.
- Legarreta, J.H., Petit, L., Jodoin, P.-M., Descoteaux, M., 2022. Clustering in tractography using autoencoders (CINTA). In: *Computational Diffusion MRI. CDMRI 2022. In: Lecture Notes in Computer Science*, vol. 13722, Springer, Cham.
- Legarreta, J.H., Petit, L., Jodoin, P.-M., Descoteaux, M., 2023. Generative Sampling in Bundle Tractography using Autoencoders (GESTA). *Med. Image Anal.* 85, 102761, URL <https://www.sciencedirect.com/science/article/pii/S1361841523000221>.
- Legarreta, J.H., Petit, L., Rheault, F., Theaud, G., Lemaire, C., Descoteaux, M., Jodoin, P.-M., 2021. Filtering in tractography using autoencoders (FINTA). *Med. Image Anal.* 72, 102126. <http://dx.doi.org/10.1016/j.media.2021.102126>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1361841521001729>.
- Maier-Hein, K.H., Neher, P.F., Houde, J.-C., Côté, M.-A., Garyfallidis, E., Zhong, J., Chamberland, M., Yeh, F.-C., Lin, Y.-C., Ji, Q., Reddick, W.E., Glass, J.O., Chen, D.Q., Feng, Y., Gao, C., Wu, Y., Ma, J., He, R., Li, Q., Westin, C.-F., Deslauriers-Gauthier, S., González, J.O., Paquette, M., St-Jean, S., Girard, G., Rheault, F., Sidhu, J., Tax, C.M.W., Guo, F., Mesri, H.Y., Dávid, S., Froeling, M., Heemskerk, A.M., Leemans, A., Boré, A., Pinsard, B., Bedetti, C., Desrosiers, M., Brambati, S., Doyon, J., Sarica, A., Vasta, R., Cerasa, A., Quattrone, A., Yeatman, J., Khan, A.R., Hodges, W., Alexander, S., Romascano, D., Barakovic, M., Auria, A., Esteban, O., Lemkaddem, A., Thiran, J.-P., Cetingul, H.E., Odry, B.L., Mailhe, B., Nadar, M.S., Pizzagalli, F., Prasad, G., Villalon-Reina, J.E., Galvis, J., Thompson, P.M., Requejo, F.D.S., Laguna, P.L., Lacerda, L.M., Barrett, R., Dell'Acqua, F., Catani, M., Petit, L., Caruyer, E., Daducci, A., Dyrby, T.B., Holland-Letz, T., Hilgetag, C.C., Stieltjes, B., Descoteaux, M., 2017. The challenge of mapping the human connectome based on diffusion tractography. *Nature Commun.* 8 (1), 1349. <http://dx.doi.org/10.1038/s41467-017-01285-x>, URL <http://www.nature.com/articles/s41467-017-01285-x>.
- Mangin, J.-F., Fillard, P., Cointepas, Y., Le Bihan, D., Frouin, V., Poupon, C., 2013. Toward global tractography. *NeuroImage* 80, 290–296. <http://dx.doi.org/10.1016/j.neuroimage.2013.04.009>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811913003364>.
- O'Donnell, L., Kubicki, M., Shenton, M., Dreusicke, M., Grimson, W., Westin, C., 2006. A Method for Clustering White Matter Fiber Tracts. *AJNR Am. J. Neuroradiol.* 27 (5), 1032–1036, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2768142/>.
- O'Donnell, L.J., Wells, W.M., Golby, A.J., Westin, C.-F., 2012. Unbiased Groupwise Registration of White Matter Tractography. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012*. In: *Lecture Notes in Computer Science*, Springer, Berlin Heidelberg, pp. 123–130. http://dx.doi.org/10.1007/978-3-642-33454-2_16.
- O'Donnell, L., Westin, C.-F., 2005. White Matter Tract Clustering and Correspondence in Populations. In: Hutchison, D., Kanade, T., Kittler, J., Kleinberg, J.M., Mattern, F., Mitchell, J.C., Naor, M., Nierstrasz, O., Pandu Rangan, C., Steffen, B., Sudan, M., Terzopoulos, D., Tygar, D., Vardi, M.Y., Weikum, G., Duncan, J.S., Gerig, G. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*, vol. 3749. In: *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 140–147. http://dx.doi.org/10.1007/11566465_18, URL http://link.springer.com/10.1007/11566465_18 Series Title.
- O'Donnell, L., Westin, C.-F., 2007. Automatic Tractography Segmentation Using a High-Dimensional White Matter Atlas. *IEEE Trans. Med. Imaging* 26, 1562–1575. <http://dx.doi.org/10.1109/TMI.2007.906785>.
- Painchaud, N., Skandarani, Y., Judge, T., Bernard, O., Lalonde, A., Jodoin, P.-M., 2020. Cardiac Segmentation With Strong Anatomical Guarantees. *IEEE Trans. Med. Imaging* 39 (11), 3703–3713. <http://dx.doi.org/10.1109/TMI.2020.3003240>, Conference Name: IEEE Transactions on Medical Imaging.

- Petit, L., Ali, K.M., Rheault, F., Boré, A., Cremona, S., Corsini, F., De Benedictis, A., Descoteaux, M., Sarubbo, S., 2022. The structural connectivity of the human angular gyrus as revealed by microdissection and diffusion tractography. *Brain Struct. Funct.* <http://dx.doi.org/10.1007/s00429-022-02551-5>, URL <https://link.springer.com/10.1007/s00429-022-02551-5>.
- Petit, L., Rheault, F., Descoteaux, M., Tzourio-Mazoyer, N., 2021. Half of the streamlines built in a whole human brain tractogram is anatomically uninterpretable. *F1000Research* 1, [version 1; not peer reviewed].
- Poulin, P., Theaud, G., Rheault, F., St-Onge, E., Bore, A., Renaud, E., de Beaumont, L., Guay, S., Jodoin, P.-M., Descoteaux, M., 2021. *TractoInferno* : A large-scale, open-source, multi-site database for machine learning dMRI tractography. pp. 1–32, *bioRxiv* 2021.11.29.470422v1.
- Poulin, P., Theaud, G., Rheault, F., St-Onge, E., Bore, A., Renaud, E., de Beaumont, L., Guay, S., Jodoin, P.-M., Descoteaux, M., 2022. *TractoInferno* - A large-scale, open-source, multi-site database for machine learning dMRI tractography. *Sci. Data* 9 (1), 725. <http://dx.doi.org/10.1038/s41597-022-01833-1>, URL <https://www.nature.com/articles/s41597-022-01833-1>, Number: 1 Publisher: Nature Publishing Group.
- Rheault, F., 2020. Analyse et reconstruction de faisceaux de la matière blanche (Ph.D. thesis). Université de Sherbrooke, Sherbrooke.
- Rheault, F., 2021. Population average atlas for RecobundlesX (1.1). <http://dx.doi.org/10.5281/zenodo.5165374>.
- Rheault, F., De Benedictis, A., Daducci, A., Maffei, C., Tax, C.M.W., Romascano, D., Caverzasi, E., Morency, F.C., Corrivetti, F., Pestilli, F., Girard, G., Theaud, G., Zemmoura, I., Hau, J., Glavin, K., Jordan, K.M., Pomiecko, K., Chamberland, M., Barakovic, M., Goyette, M., Poulin, P., Chenot, Q., Panesar, S.S., Sarubbo, S., Petit, L., Descoteaux, M., 2020a. Tractostorm: The what, why, and how of tractography dissection reproducibility. *Hum. Brain Map.* 41 (7), 1859–1874. <http://dx.doi.org/10.1002/hbm.24917>, URL <http://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.24917>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.24917>.
- Rheault, F., Poulin, P., Caron, A.V., St-Onge, E., Descoteaux, M., 2020b. Common misconceptions, hidden biases and modern challenges of dMRI tractography. *J. Neural Eng.* 17 (1), 011001. <http://dx.doi.org/10.1088/1741-2552/ab6aad>, Publisher: IOP Publishing.
- Rheault, F., Schilling, K.G., Valcourt-Caron, A., Théberge, A., Poirier, C., Grenier, G., Guberman, G.I., Begnoche, J., Legarreta, J.H., Y. Cai, L., Roy, M., Edde, M., Caceres, M.P., Ocampo-Pineda, M., Al-Sharif, N., Karan, P., Bon-tempi, P., Obaid, S., Bosticardo, S., Schiavi, S., Sairanen, V., Daducci, A., Cutting, L.E., Petit, L., Descoteaux, M., Landman, B.A., 2022. Tractostorm 2: Optimizing tractography dissection reproducibility with segmentation protocol dissemination. *Hum. Brain Map.* 43 (7), 2134–2147. <http://dx.doi.org/10.1002/hbm.25777>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.25777>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/hbm.25777>.
- Rheault, F., St-Onge, E., Sidhu, J., Maier-Hein, K., Tzourio-Mazoyer, N., Petit, L., Descoteaux, M., 2019. Bundle-specific tractography with incorporated anatomical and orientational priors. *NeuroImage* 186, 382–398. <http://dx.doi.org/10.1016/j.neuroimage.2018.11.018>.
- Rockafellar, R.T., Wets, R.J.-B., 2009. *Variational Analysis*. Springer Science & Business Media.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (Eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. In: *Lecture Notes in Computer Science*, Springer International Publishing, Cham, pp. 234–241. http://dx.doi.org/10.1007/978-3-319-24574-4_28.
- Shastin, D., Genc, S., Parker, G.D., Koller, K., Tax, C.M., Evans, J., Hamandi, K., Gray, W.P., Jones, D.K., Chamberland, M., 2022. Surface-based tracking for short association fibre tractography. *NeuroImage* 260, 119423. <http://dx.doi.org/10.1016/j.neuroimage.2022.119423>, URL <https://linkinghub.elsevier.com/retrieve/pii/S1053811922005407>.
- Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. In: *Monographs on Statistics and Applied Probability*, vol. 26, Chapman and Hall.
- Sotiropoulos, S.N., Zalesky, A., 2019. Building connectomes using diffusion MRI: why, how and but. *NMR Biomed.* 32 (4), e3752. <http://dx.doi.org/10.1002/nbm.3752>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nbm.3752>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nbm.3752>.
- St-Onge, E., Daducci, A., Girard, G., Descoteaux, M., 2018. Surface-enhanced tractography (SET). *NeuroImage* 169, 524–539. <http://dx.doi.org/10.1016/j.neuroimage.2017.12.036>, URL <https://www.sciencedirect.com/science/article/pii/S1053811917310583>.
- Theaud, G., 2022. dMRIQCpy: a python based toolbox for diffusion MRI quality control and beyond. In: (ISMRM 2022). URL <https://archive.ismrm.org/2022/3906.html>.
- Theaud, G., Edde, M., Dumont, M., Zotti, C., Zucchelli, M., Deslauriers-Gauthier, S., Deriche, R., Jodoin, P.-M., Descoteaux, M., 2022. DORIS: A diffusion MRI-based 10 tissue class deep learning segmentation algorithm tailored to improve anatomically-constrained tractography. *Front. Neuroimaging* 1, URL <https://www.frontiersin.org/articles/10.3389/fnimg.2022.917806>.
- Theaud, G., Houde, J.-C., Boré, A., Rheault, F., Morency, F., Descoteaux, M., 2020. TractoFlow: A robust, efficient and reproducible diffusion MRI pipeline leveraging Nextflow & Singularity. *NeuroImage* 218, 116889. <http://dx.doi.org/10.1016/j.neuroimage.2020.116889>, URL <https://www.sciencedirect.com/science/article/pii/S105381192030375X>.
- Tournier, J.-D., Calamante, F., Connelly, A., 2007. Robust determination of the fibre orientation distribution in diffusion MRI: Non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* 35 (4), 1459–1472. <http://dx.doi.org/10.1016/j.neuroimage.2007.02.016>, URL <https://www.sciencedirect.com/science/article/pii/S1053811907001243>.
- Tournier, J.-D., Calamante, F., Connelly, A., 2012. MRtrix: Diffusion tractography in crossing fiber regions. *Int. J. Imaging Syst. Technol.* 22 (1), 53–66. <http://dx.doi.org/10.1002/ima.22005>, URL <http://onlinelibrary.wiley.com/doi/abs/10.1002/ima.22005>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ima.22005>.
- Tournier, J.-D., Mori, S., Leemans, A., 2011. Diffusion tensor imaging and beyond. *Magn. Reson. Med.* 65 (6), 1532–1556. <http://dx.doi.org/10.1002/mrm.22924>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.22924>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/mrm.22924>.
- Vallat, R., 2018. Pingouin: statistics in Python. *J. Open Source Softw.* 3 (31), 1026. <http://dx.doi.org/10.21105/joss.01026>, URL <https://joss.theoj.org/papers/10.21105/joss.01026>.
- Visser, E., Nijhuis, E., Buitelaar, J., Zwiers, M., 2011. Partition-based mass clustering of tractography streamlines. *NeuroImage* 54 (1), 303–312.
- Wang, X., W., W.G., Westin, C., 2011. Tractography segmentation using a hierarchical dirichlet processes mixture model. *NeuroImage* 54 (1), 290–302.
- Warrington, S., Bryant, K.L., Khrapitchev, A.A., Sallet, J., Charquero-Ballester, M., Douaud, G., Jbabdi, S., Mars, R.B., Sotiropoulos, S.N., 2020. XTRACT - Standardised protocols for automated tractography in the human and macaque brain. *NeuroImage* 217, 116923. <http://dx.doi.org/10.1016/j.neuroimage.2020.116923>, URL <https://www.sciencedirect.com/science/article/pii/S1053811920304092>.
- Wassermann, D., Makris, N., Rathi, Y., Shenton, M., Kikinis, R., Kubicki, M., Westin, C.-F., 2016. The white matter query language: a novel approach for describing human white matter anatomy. *Brain Struct. Funct.* 221 (9), 4705–4721. <http://dx.doi.org/10.1007/s00429-015-1179-4>.
- Wasserthal, J., Maier-Hein, K.H., Neher, P.F., Northoff, G., Kubera, K.M., Fritze, S., Harneit, A., Geiger, L.S., Tost, H., Wolf, R.C., Hirjak, D., 2020. Multiparametric mapping of white matter microstructure in catatonia. *Neuropsychopharmacology* 45 (10), 1750–1757. <http://dx.doi.org/10.1038/s41386-020-0691-2>, URL <https://www.nature.com/articles/s41386-020-0691-2>, Number: 10 Publisher: Nature Publishing Group.
- Wasserthal, J., Neher, P.F., Hirjak, D., Maier-Hein, K.H., 2019. Combined tract segmentation and orientation mapping for bundle-specific tractography. *Med. Image Anal.* 58, 101559. <http://dx.doi.org/10.1016/j.media.2019.101559>, URL <https://www.sciencedirect.com/science/article/pii/S136184151930101X>.
- Wasserthal, J., Neher, P.F., Maier-Hein, K.H., 2018a. Tract orientation mapping for bundle-specific tractography. Technical Report, (arXiv:1806.05580), arXiv, URL <http://arxiv.org/abs/1806.05580>, arXiv:1806.05580 [cs] type: article.
- Wasserthal, J., Neher, P., Maier-Hein, K.H., 2018b. TractSeg - Fast and accurate white matter tract segmentation. *NeuroImage* 183, 239–253. <http://dx.doi.org/10.1016/j.neuroimage.2018.07.070>, URL <https://www.sciencedirect.com/science/article/pii/S1053811918306864>.
- Yendiki, A., 2011. Automated probabilistic reconstruction of white-matter pathways in health and disease using an atlas of the underlying anatomy. *Front. Neuroinform.* 5, <http://dx.doi.org/10.3389/fninf.2011.00023>, URL <http://journal.frontiersin.org/article/10.3389/fninf.2011.00023/abstract>.
- Zhang, F., Wu, Y., Norton, I., Rigolo, L., Rathi, Y., Makris, N., O'Donnell, L.J., 2018. An anatomically curated fiber clustering white matter atlas for consistent white matter tract parcellation across the lifespan. *NeuroImage* 179, 429–447. <http://dx.doi.org/10.1016/j.neuroimage.2018.06.027>, URL <https://www.sciencedirect.com/science/article/pii/S1053811918305342>.
- Zhang, Y., Zhang, J., Oishi, K., Faria, A.V., Jiang, H., Li, X., Akhter, K., Rosa-Neto, P., Pike, G.B., Evans, A.C., Toga, A.W., Woods, R., Mazziotta, J.C., Miller, M.I., van Zijl, P.C.M., Mori, S., 2010. Atlas-guided tract reconstruction for automated and comprehensive examination of the white matter anatomy. *NeuroImage* 52 (4), 1289–1301. <http://dx.doi.org/10.1016/j.neuroimage.2010.05.049>.