

Progressive and Efficient Multi-Resolution Representations for Brain Tractograms

C. Mercier^{1,2}, P. Gori¹, D. Rohmer², M-P. Cani², T. Boubekeur¹, J-M. Thiery¹ and I. Bloch¹

¹LTCl, Télécom ParisTech, Université Paris-Saclay, Paris, France

²LIX, École Polytechnique, Palaiseau, France

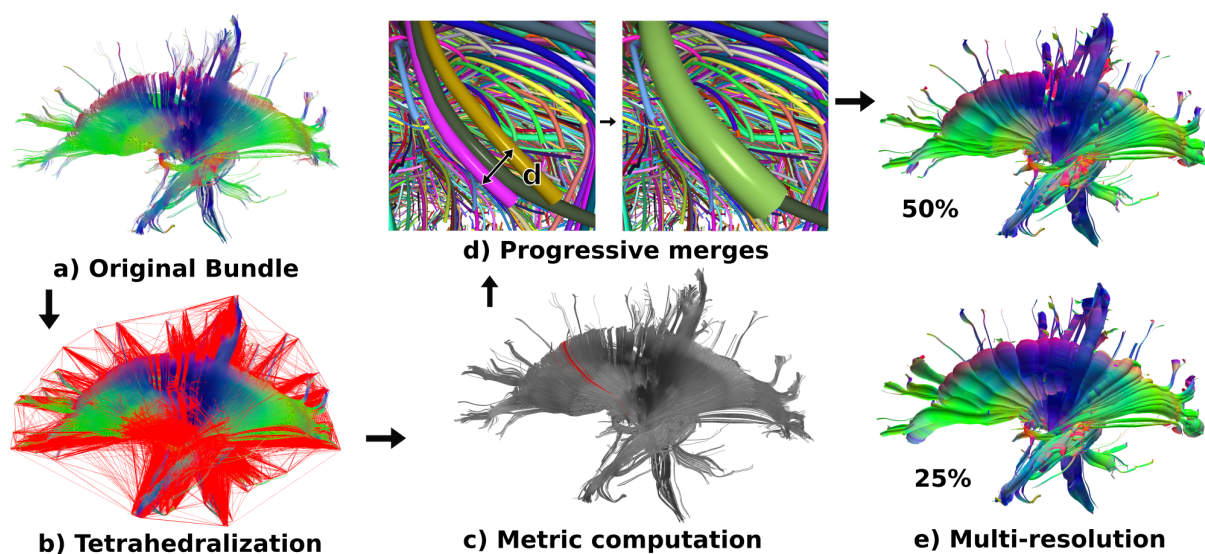


Figure 1: Multi-resolution pipeline: a) original bundle tractogram, here the thalamocortical one, b) connections created by the Delaunay tetrahedralization based on the extremities of the fibers, c) the similarity is computed for each fiber, taking into account only the neighbors found in step b), d) couples of closest fibers are progressively merged into generalized cylinders, e) the final multi-resolution representation makes it possible to navigate through the different levels of detail in real-time. The percentage refers to the fraction of employed generalized cylinders compared to the original number of fibers. Color code depends on the orientation of the fiber: red for left-right, blue for inferior-superior and green for anteroposterior.

Abstract

Current tractography methods generate tractograms composed of millions of 3D polylines, called fibers, making visualization and interpretation extremely challenging, thus complexifying the use of this technique in a clinical environment. We propose to progressively simplify tractograms by grouping similar fibers into generalized cylinders. This produces a fine-grained multi-resolution model that provides a progressive and real-time navigation through different levels of detail. This model preserves the overall structure of the tractogram and can be adapted to different measures of similarity. We also provide an efficient implementation of the method based on a Delaunay tetrahedralization. We illustrate our method using the Human Connectome Project dataset.

1. Introduction

Tractography from diffusion MRI is currently the only technique able to non-invasively explore the white matter architecture of the brain. It results in a tractogram which is a bundle of 3D polylines,

usually called fibers, which are estimates of the trajectories of large groups of neural tracts. Tractography has been proven to be an invaluable tool for clinicians and researchers. It is nowadays used on a daily basis by neurosurgeons for pre-operating planning and

during surgical operations [JDML17]. It also offers important information for studying pathological processes in neurological diseases [CCJB*08].

Recent tractography methods produce up to one million of fibers [TML11]. This can complicate the rendering, visualization and interpretation of tractograms, thus limiting the aforementioned clinical applications. Furthermore, the considerable number of fibers can make computationally intractable processes such as non-linear registration or atlas construction [GCMK*16], which are important for research purposes. Many fibers might have a similar trajectory and connectivity, making the tractogram redundant. For this reason, several authors have proposed new geometric representations and visualization techniques to simplify tractograms. One of the most popular approaches consists in grouping similar fibers into clusters [GCMK*16, GBC*12, GPR*11, MW*07, ZL02] which are then approximated with one representative fiber usually called *prototype* [GBC*12, GPR*11]. Other authors have also proposed to represent the spatial extent of the clusters using an encompassing geometry [MW*07]. These methods are usually controlled by one parameter, e.g. a threshold [ZL02], thus presenting only one level of resolution at a time. Furthermore, information such as the number of fibers or the spatial extent (i.e. the volume) of the cluster might be lost in the process. To improve the visualization quality and efficiency, the geometric models used to represent the fibers are often computed directly on the GPU [PFK07, RBE*06, ESM*05]. Other methods, such as in [EBB*15], ease visualization by bringing similar fibers closer to each other. In this paper, we do not focus on single resolution approximations (i.e. clustering), but propose a multi-resolution representation based on progressive mergings, that preserves the overall structure of the tractogram and whose continuous levels of resolution can be traversed in real-time.

Progressive simplification methods, which focus on decimating complex 3D objects while preserving both important geometric features and topological relationships, have so far not been explored to approximate large scale tractograms. In [ZL02], Zhang and Laidlaw propose to use a hierarchical clustering algorithm to progressively group together similar fibers. They only applied it on bundles composed of a small number of fibers. Their goal was to divide tractograms into clusters and not to propose a new geometric representation or visualization technique (as it is the case in this paper). Taking inspiration from error-driven surface mesh simplification [GH97], we propose a progressive merging strategy for grouping fibers into generalized cylinders. The proposed method reduces the redundancy of the tractogram, producing a multi-resolution structure, which is organized into a nested hierarchy of levels of detail. Every fusion of fibers (or cylinders) represents a new level of resolution. Once the entire multi-resolution representation is computed, it is possible for the user to navigate through different levels of detail in a continuous fashion and in real-time, while maintaining the overall structure of the original tractogram. Furthermore, we also propose an efficient implementation based on a Delaunay tetrahedralization which makes it possible to use our method on large tractograms containing millions of fibers.

2. Method

Fiber Decimation We propose a tractogram simplification method (see Fig.1) based on the progressive mesh methodology [GH97,

HDD*93]. Given a tractogram with N fibers, we first look for the two most similar fibers based on a similarity measure. Once detected, the couple is collapsed into a single generalized cylinder – all input fibers can be seen as generalized cylinders with a null radius. The process is then iterated until obtaining a single large cylinder (i.e. $N - 1$ iterations) or using a stopping criterion to prevent over-simplification.

The proposed “progressive brain tractograms” algorithm is general and can be used with any similarity measure. It is inherently multi-resolution, where every level of detail corresponds to the fusion of two cylinders. Once the sequence of fusions is computed, we can visualize the tractogram at any resolution and switch among levels in real-time.

In the following, we will use three dissimilarity measures: the Mean of Closest distances (MC) [GPR*11], the minimum average direct-flip (MDF) [GBC*12] and one based on the computational model of Weighted Currents (WC) [GCMK*16]. Let $X = \{x_i, i = 1 \dots N\}$ and $Y = \{y_j, j = 1 \dots M\}$ be two fibers composed of N and M points respectively the MC distance is defined as:

$$MC(X, Y) = \text{mean}(d_m(X, Y), d_m(Y, X)) \quad (1)$$

where $d_m(X, Y) = \frac{1}{N} \sum_{i=1}^N \min_{y_j \in Y} \|x_i - y_j\|_2$. The MDF is defined in a similar way, by assuming point-wise correspondence between fibers. The similarity measure of WC is instead defined as:

$$WC(X, Y) = |K_a(\|f^a - t^a\|) K_b(\|f^b - t^b\|) \sum_{i=1}^{N-1} \sum_{j=1}^{M-1} \alpha_i^T K_g(\|c_i - d_j\|) \beta_j| \quad (2)$$

where c_i and α_i (respectively d_j and β_j) are the centers and tangent vectors of X (respectively Y), f^a , f^b and t^a , t^b are the corresponding endpoints of X and Y respectively, and K_a , K_b and K_g are three Gaussian kernels parametrized by σ_a , σ_b and σ_g respectively.

Delaunay Tetrahedralization Comparing all fibers to each other leads to quadratic complexity and intractable computations. Furthermore, most of the computations would be useless since similar close-by fibers should be merged in priority, thus indicating that only neighboring fibers need to be compared and considered as candidates for merging. We exploit this idea by computing once and for all an adjacency relationship among fibers, relying on the Delaunay tetrahedralization of their extremities. This tetrahedralization uses the Euclidean distance between all the endpoints of the fibers and is computed using the fast implementation of TetGen [Si15]. The adjacency links are given by the edges of the tetrahedral mesh. By exploiting this tetrahedralization, we drastically reduce the amount of computations since we now compare every fiber with only the adjacent (linked) ones. When merging two fibers together, the adjacency links are updated accordingly: a neighborhood relationship is built between the merged fiber and all the neighbors of its two original fibers, while previous neighborhood relationships involving these two original fibers are suppressed. Relying on the Delaunay tetrahedralization provides a geometrically well distributed set of adjacency links, which is beneficial to our simplification. Furthermore, results obtained with this technique are not biased, namely equivalent to the ones obtained without using it.




Figure 2: Scheme of the merging of two fibers/cylinders C_1 and C_2 into C_3 .

Geometric Representation In this paper, we consider that every fiber is described by its geometry and connectivity. The employed geometric representation should preserve these properties. To this end, we propose to use generalized cylinders with an elliptical basis as geometric representations for the merged fibers. We define one ellipse per vertex of the center-curve such that the resulting cylinder incorporates the trajectory and endpoints of the fibers (or cylinders) of the previous level of resolution.

Given two fibers/cylinders C_1 and C_2 , Fig.2 schematically presents how they are merged into cylinder C_3 . We use as reference the center-curve (or fiber) with less points (C_1 in Fig.2). For each point of C_1 , we look for the closest point in C_2 . The corresponding point in C_3 is computed as their weighted mean, where the weights W_1 and W_2 correspond to the number of original fibers that C_1 and C_2 represent (and we set $W_3 = W_1 + W_2$ accordingly). To avoid flickering, the thickness of rendered cylinders is clamped to a minimum. Moreover, it is important to notice that we use the Euclidean distance between the endpoints of C_1 and C_2 to compute the extremities of C_3 . In order to make sure that the extremities of C_3 would actually lie at the border with the gray matter surfaces, one should use a (computationally expensive) geodesic distance. This is left as future work.

Implementation We implemented our method in C++ using Qt for the graphical user interface and OpenGL for the rendering. To improve efficiency and memory usage, the geometry is computed on the GPU, using the hardware tessellation unit to synthesize on-the-fly our visual approximations. We will make the code publicly available at <https://perso.telecom-paristech.fr/comercier/>. We provide a video of the proposed interactive system as a complementary material (same url).

3. Results and Discussion

Dataset We conducted our experiments on the *HCP* dataset (<https://db.humanconnectome.org/>). Tractograms are obtained using the SDSTREAM deterministic tracking algorithm of *MRTrix3* [TCC12]. All tracts employed in the experiments (*Uncinate Fasciculus*, *Ifof* and *Thalamocortical* bundle) are extracted using either *WMQL* [WMR*16] or manual segmentation. Please note that our method is general and could be used with any streamline tractography algorithm.

Comparison with QuickBundles In Fig.3, we compare our method with QuickBundles (QB) [GBC*12], a well-known approximation algorithm for brain white matter tractograms based on prototypes. We first executed QB on a bundle composed of 19,782 fibers. We used different thresholds in the range 5-10mm as suggested in [GBC*12]. We chose the middle one (7) for comparison. It resulted in 275 prototypes. We then executed our algo-

rithm (once) with the MDF metric, used by QB, and then interactively change the resolution to obtain Fig.3(c,d and e). Our result, in Fig.3e, has as many cylinders as prototypes in Fig.3b. The proposed method, after a single and fast pre-computation (see Tab.1), creates an encompassing representation that well approximates the original bundle at each level of resolution. On the contrary, QB produces prototypes that, depending on a user-defined threshold, might not preserve the overall structure and volume of the original bundle. Moreover, one might need to execute QB several times before finding an optimal threshold (for a given clinical/research application).

Multi-resolution Figure 4 shows results on three different bundles: the *Uncinate Fasciculus* (a, b), the *Ifof* (c), and a whole brain (d). The measures used for the experiments are MC (see (1)) for (a), WC (see (2)) for (b and d), with $\sigma_a = 6mm$, $\sigma_b = 6mm$ and $\sigma_g = 8mm$, and MDF for (c). Computation times for these experiments are given in Tab.1 and were obtained on an Intel Xeon E5-1650V4. These results suggest that the overall structure of the bundles is preserved across the multi-resolution representation, even at low resolutions, e.g. 15%. The graphs on the right of Fig.4, illustrate the mean and max valences across the resolutions. The valence of a fiber is defined as the number of its neighbors. The evolution of the valence depends on the order in which the fibers are merged, and therefore on the employed distance/similarity measure. A measure that favors geometrically well distributed cylinders – composed of a similar number of original fibers – preserves a bounded valence. In this case, the size of the priority list remains linear in the number of original fibers n , and our strategy exhibits a total time complexity of $O(n \log(n))$. These complexities were observed in our experiments (see Tab.1). In this table and in Fig. 4, we notice that the mean valence remains relatively stable using all distance measures, even with one million of fibers (Fig.4d). This suggests that our method, with the proposed dissimilarities, should preserve the overall structure of the brain since fibers are aggregated all around the tractogram and not only in a specific area.

The proposed method results in a multi-resolution representation which approximates the original tractogram with a decreasing precision, and is not intended to produce anatomically reproducible clusters among subjects. Moreover, it is initialized with a Delaunay tetrahedralization of the extremities, which means that we inherently assume that two fibers are similar if their extremities are close to each other. The employed metric, as the ones proposed here, should therefore be consistent with this assumption.

For similarity measures defined on inner product spaces, we also propose an automatic stopping criterion to prevent oversimplification (e.g. a single cylinder). Two cylinders are not merged together if they are orthogonal (or almost) to each other. Furthermore, all fibers not merged at the end of the algorithm (and cylinders representing very few fibers) can be considered as outliers and thus discarded.

4. Conclusion

We introduced a new multi-resolution representation for brain tractograms that reduces the redundancy, easing the visualization and interpretation. It supports any similarity measure between fibers.

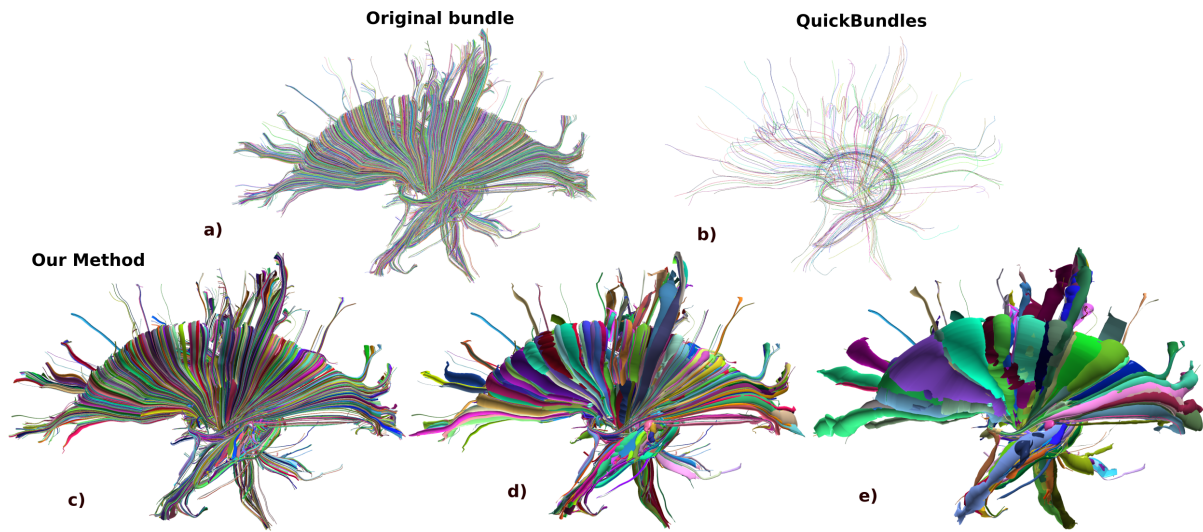


Figure 3: a) Thalamocortical bundle with 19,782 fibers, b) reduced to 275 prototypes with QuickBundles (Threshold=7mm), c), d) and e) reduced to 6925, 1422 and 275 cylinders with our method.

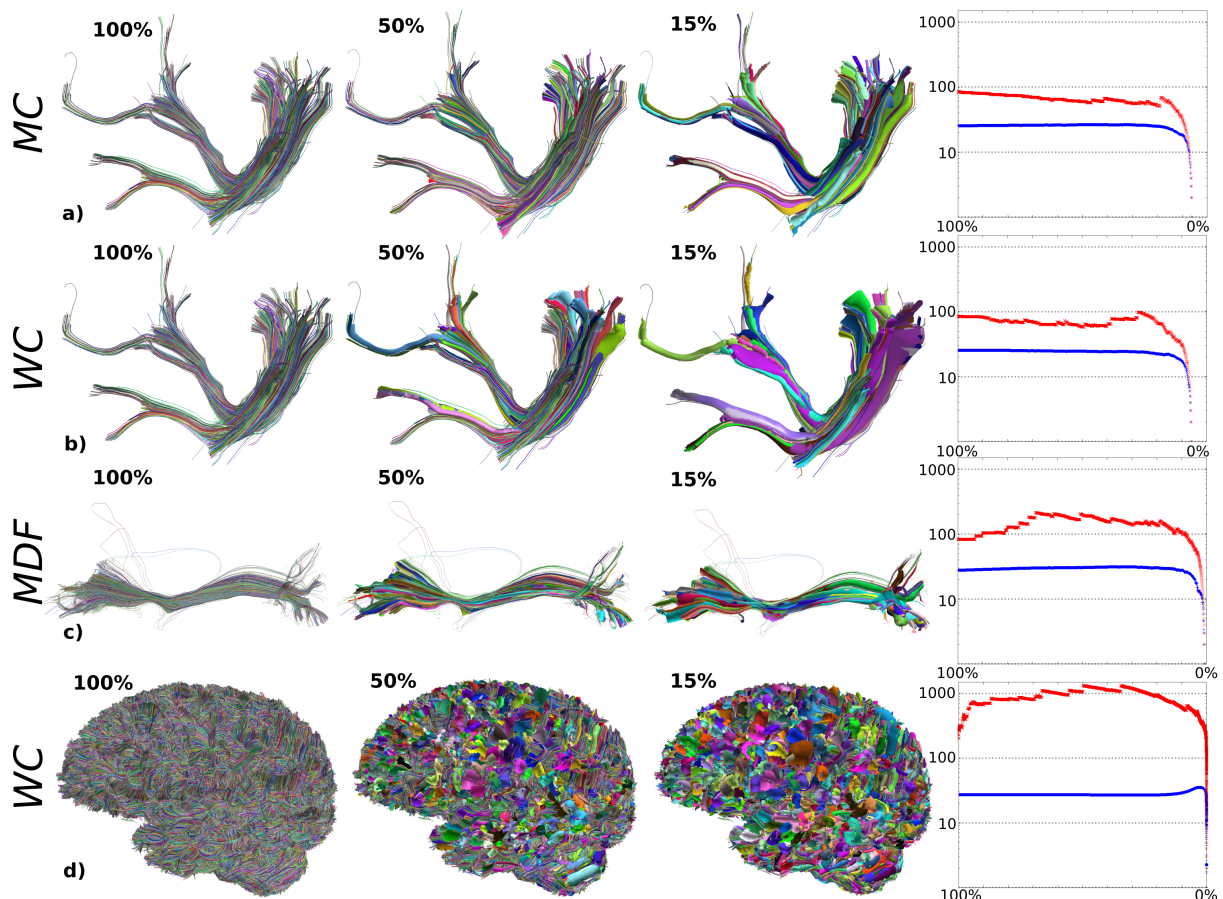


Figure 4: Bundles at different resolutions : a) Uncinate Fasciculus, computed with the mean of closest distances ; b) Uncinate Fasciculus, computed with the dissimilarity of the weighted currents; c) Ifof bundle, computed with the dissimilarity of the weighted currents; d) whole brain tractogram, with 1 million fibers, computed with the dissimilarity of the weighted currents. The graphs represent the valence of the fibers in function of the resolution. The maximum valence is in red, and the average valence is in blue.

Table 1: Maximum and average valence with computational times for some bundles. The computational time for QuickBundles is presented on the last line as comparison. The threshold used for obtaining a single simplification was 7mm.

Bundle	Metric	Number of fibers	Maximum valence	Average valence	Tetrahedralization time (s)	Total time (s)
Uncinate Fasciculus	MC	940	85	25.0 ± 3.2	0.01	1.3
	WC	940	98	23.8 ± 3.2	0.01	1.9
Ifof	MC	1,983	217	28.3 ± 4.3	0.02	4.3
	WC	1,983	197	27.7 ± 3.4	0.02	6.9
	MDF	1,983	216	29.0 ± 3.8	0.02	1.7
Thalamocortical bundle	MC	19,782	760	30.3 ± 3.3	0.28	34.0
	WC	19,782	717	29.3 ± 2.5	0.28	55.6
Whole Brain Tractogram	MC	1,000,000	9,392	29.6 ± 2.5	15.42	1,893.7
	WC	1,000,000	1,331	27.6 ± 1.8	15.42	3,094.2
	MDF	1,000,000	373	29.9 ± 2.7	15.42	616.3
	QB (MDF)	1,000,000	-	-	-	4,825.2

We also proposed a method to determine adequate candidates for efficient merging of groups of fibers. Our progressive method makes it possible to display the input tractogram at any resolution in a continuous and real-time way. From a technical point of view, our two main contributions are a multi-resolution representation for tractograms based on a progressive decimation algorithm, and a combinatorial strategy based on a Delaunay tetrahedralization to make it computationally tractable. Visualizing groups of similar fibers as single generalized cylinders and being able to easily change the level of resolution may be very useful for clinicians. For instance, it can help neurosurgeons identify relevant anatomical tracts which should not be severed during the operation, thus reducing post-operative complications and improving the clinical outcome. It is to note that we received very positive feedbacks from our neurosurgeons colleagues of the *Ste Anne* hospital in Paris.

Future Work To improve visualization, in particular the 3D perception of the bundles, ambient occlusion could be efficiently implemented, exploiting the hierarchical structure. Furthermore, we also plan to expand the proposed representation by adding functional information defined, for instance, as a scalar field (e.g. Fractional Anisotropy (FA)) which could be visualized using color or texture.

Acknowledgements We would like to thank Dr. Alexandre Roux for his feedback on our results. This research was supported by Labex DigiCosme (project ANR11LABEX0045DIGICOSME) operated by ANR as part of the program Investissement d’Avenir IDEX ParisSaclay (ANR11IDEX000302).

References

- [CCJB*08] CICCARELLI O., CATANI M., JOHANSEN-BERG H., CLARK C., THOMPSON A.: Diffusion-based tractography in neurological disorders: concepts, applications, and future developments. *The Lancet Neurology* 7, 8 (2008), 715 – 727. 2
- [EBB*15] EVERTS M., BEGUE E., BEKKER H., ROERDINK J., ISENBURG T.: Exploration of the Brain’s White Matter Structure through Visual Abstraction and Multi-Scale Local Fiber Tract Contraction. *IEEE TVCG* 21, 7 (2015). 2
- [ESM*05] ENDERS F., SAUBER N., MERHOF D., HASTREITER P., NIMSKY C., STAMMINGER M.: Visualization of white matter tracts with wrapped streamlines. In *IEEE Visualization, 2005. VIS 05* (2005), pp. 51–58. 2
- [GBC*12] GARYFALLIDIS E., BRETT M., CORREIA M. M., WILLIAMS G. B., NIMMO-SMITH I.: QuickBundles, a Method for Tractography Simplification. *Frontiers in Neuroscience* 6 (Dec. 2012). 2, 3
- [GCMK*16] GORI P., COLLIOT O., MARRAKCHI-KACEM L., WORBE Y., FALLANI F. D. V., CHAVEZ M., POUPON C., HARTMANN A., AYACHE N., DURRLEMAN S.: Parsimonious Approximation of Streamline Trajectories in White Matter Fiber Bundles. *IEEE TMI* 35, 12 (2016), 2609–2619. 2
- [GH97] GARLAND M., HECKBERT P. S.: Surface Simplification Using Quadric Error Metrics. In *ACM SIGGRAPH* (1997), pp. 209–216. 2
- [GPR*11] GUEVARA P., POUPON C., RIVIÁLRE D., COINTEPAS Y., DESCOTEAUX M., THIRION B., MANGIN J. F.: Robust clustering of massive tractography datasets. *NeuroImage* 54, 3 (2011), 1975–1993. 2
- [HDD*93] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Mesh optimization. In *ACM SIGGRAPH* (1993), pp. 19–26. 2
- [JDML17] JEURISSEN B., DESCOTEAUX M., MORI S., LEEMANS A.: Diffusion MRI fiber tractography of the brain. *NMR in Biomedicine* (2017). 2
- [MWW*07] MADDAH M., WELLS W. M., WARFIELD S. K., WESTIN C.-F., GRIMSON W. E. L.: Probabilistic Clustering and Quantitative Analysis of White Matter Fiber Tracts. In: *IPMI* (2007), 372–383. 2
- [PFK07] PETROVIC V., FALLON J., KUESTER F.: Visualizing Whole-Brain DTI Tractography with GPU-based Tuboids and LoD Management. *IEEE TVCG* 13, 6 (2007), 1488–1495. 2
- [RBE*06] REINA G., BIDMON K., ENDERS F., HASTREITER P., ERTL T.: GPU-based hyperstreamlines for diffusion tensor imaging. In *EuroVis* (2006), vol. 6, pp. 35–42. 2
- [Si15] SI H.: TetGen, a Delaunay-Based Quality Tetrahedral Mesh Generator. *ACM Trans. Math. Softw.* 41, 2 (Feb. 2015), 11:1–11:36. 2
- [TCC12] TOURNIER J.-D., CALAMANTE F., CONNELLY A.: MRtrix: Diffusion tractography in crossing fiber regions. *International Journal of Imaging Systems and Technology* 22, 1 (2012), 53–66. 3
- [TML11] TOURNIER J.-D., MORI S., LEEMANS A.: Diffusion tensor imaging and beyond. *Magnetic Resonance in Medicine* 65, 6 (2011), 1532–1556. 2
- [WMR*16] WASSERMANN D., MAKRIS N., RATHI Y., SHENTON M., KIKINIS R., KUBICKI M., WESTIN C.-F.: The white matter query language: a novel approach for describing human white matter anatomy. *Brain Structure and Function* 221, 9 (2016), 4705–4721. 3
- [ZL02] ZHANG S., LAIDLAW D. H.: Hierarchical Clustering of Streamtubes. *Technical report, Brown University, CS Department* (2002). 2