

Tubular Fiber Bundle Segmentation for Diffusion Weighted Imaging

M. Niethammer^{1,2}, C. Zach¹, J. Melonakos³, and A. Tannenbaum^{3,4}

¹ Department of Computer Science

² Biomedical Research Imaging Center, School of Medicine
University of North Carolina, Chapel Hill, NC, USA

³ School of Electrical and Computer Engineering

⁴ School of Biomedical Engineering
Georgia Institute of Technology, Atlanta, GA, USA

Abstract. This paper proposes a methodology to segment tubular fiber bundles from diffusion weighted magnetic resonance images (DW-MRI). Segmentation is simplified by locally reorienting diffusion information based on large-scale fiber bundle geometry. Segmentation is achieved through simple global statistical modeling of diffusion orientation. Utilizing a modification of a recent segmentation approach by Bresson et al. [19] allows for a convex optimization formulation of the segmentation problem, combining orientation statistics and spatial regularization. The approach compares favorably with segmentation by full-brain streamline tractography.

1 Introduction

Diffusion weighted (DW) magnetic resonance imaging (MRI) allows for in-vivo measurements of water diffusion in tissues such as the human brain. While brain white matter appears uniform in structural MRI, DW-MRI measurements can provide estimates of macroscopic fiber bundle direction as well as indicate changes in tissue properties. However, the relation between DW-MRI signal and white matter ultra-structure is only known partially. For example, how axonal organization and geometry relates to a measured diffusion profile in general remains an open question. Fiber bundle direction correlates with the major diffusion direction in fiber bundle areas comprised of large numbers of approximately unidirectional axons [1]. This allows for the estimation of distinct fiber bundles from DW-MRI measurements.

A variety of approaches to extract white matter bundles from diffusion weighted images exist. They may be classified into streamline-based approaches and voxel-based approaches. The streamline-based approaches utilize streamline tractography to come up with bundle segmentations. This can for example be direct voxelization of the streamlines, voxelization preceded by streamline clustering [21], or stochastic tractography [16, 11]. Voxel-based approaches aim at extracting white matter bundles directly from the voxel data without using streamline tractography. Approaches include voxel-based clustering [18], surface-evolution using global statistics [6, 23] or local similarity terms [10], optimal

connectivity methods [17], region-growing [9], hidden Markov measure fields [15] and fuzzy segmentation [14].

This paper proposes a segmentation approach based on reorienting the diffusion measurements. Reorientation information is derived from large-scale fiber bundle geometry; it facilitates region-based fiber bundle segmentation with global statistics. The approach is computationally efficient, is simple, allows for reliable optimization, and is robust to local noise effects.

Briefly summarizing the remainder of this paper, in Sec. 2, we give an overview of the system. Sec. 3 introduces the *local* coordinate system used for the reorientation of diffusion information. Sec. 4 describes how to extend the local coordinate system to the complete image volume. The reorientation of diffusion data is described in Sec. 5. Sec. 6 and 7 describe the statistical modeling of fiber bundle direction and its use for bundle segmentation respectively. Results are given in Sec. 8. Sec. 9 concludes the paper with a discussion of the approach, and an outlook on possible future work.

2 System Overview

This section summarizes the key steps in the proposed segmentation approach. The overall goal of the method is to be able to segment tubular fiber bundles from diffusion weighted images. Segmentation requires a suitable similarity measure for voxel grouping into object foreground and object background. While a multitude of segmentation methods for diffusion weighted images exists (see Sec. 1) arguably the methods used in practice are based on streamlining: direct voxelization of streamlining results, clustering of streamlines, or stochastic tractography. This is surprising, because (i) streamlining approaches are sensitive to noise and (ii) volumetric segmentation algorithms developed outside the area of diffusion weighted imaging have either not been applied to DW-MRI or only with limited success. A major impediment to adopting existing volumetric segmentation approaches for DWI segmentation is the nature of DWI data. DWI data is (i) vector-valued (tensor-valued if diffusion tensors are computed), is (ii) axial (identifying antipodal directions), typically has (iii) low signal to noise ratio and is of relatively low resolution, and is (iv) spatially non-stationary (i.e., large scale orientation changes are expected to occur within individual fiber bundles).

Fig. 1 illustrates diffusion tensors changing direction along a fiber bundle and the same set of diffusion tensors when realigned relative to a representative fiber tract. This realignment process is at the core of the approach proposed in this paper. Realignment simplifies the original problem by making it spatially stationary. Segmentation methods for vector-valued images can then be employed for fiber bundle segmentation. Note that standard streamline tractography usually incorporates a weak, implicit form of spatial realignment by disallowing orientation changes considered too drastic.

The proposed approach is:

- 1) For every candidate point in the image volume, find the *closest* point on the representative fiber tract.




Fig. 1. Tensor reorientation concept. The spatially varying tensor orientation can largely be removed by reorientation with respect to a representative fiber tract (blue).

- 2) Regard the candidate point as part of the fiber bundle if its diffusion information is *similar* to the diffusion information at the closest point.
- 3) Create a *spatially consistent* segmentation based on the similarities of 2).

The key questions are, what is meant by “closest,” “similar,” and “spatially consistent.” The direct approach to measure closeness is to look at Euclidean distance. Euclidean distance typically does not yield unique point to point correspondences. Sec. 4 thus proposes a method based on frame diffusion. Since the focus of this paper is the segmentation of tubular fiber bundles, the overall fiber bundle geometry can be approximately described by the space curve given by a representative fiber tract. The (regularized) Frenet frame of the space curve can then be used as a local coordinate frame and as the basis for frame diffusion; see Sec. 3.

Many probabilistic and deterministic similarity measures have been proposed for diffusion weighted imaging (in particular, for diffusion tensor imaging; see for example [6, 10]). One of the simplest measures of diffusion similarity is to measure angular deviations of the major directions of diffusion. This is in line with streamline tractography which typically uses only the principal diffusion direction for streamline propagation⁵ and will be used in a probabilistic formulation in this paper as discussed in Sec. 6. To achieve spatial consistency, which cannot be achieved by local segmentation decisions based on directional statistics and reorientation of diffusion measurements alone, regularization is necessary. Sec. 7 describes the proposed segmentation approach based on a slight modification of the convex optimization formulation by Bresson et al. [19].

3 The Regularized Axial Frenet Frame

To parameterize tubular fiber bundles, a suitable coordinate system is necessary. For space-curves, the Frenet frame can be used. Given a parameterized curve $\mathcal{C}(p) : [0, 1] \mapsto \mathbb{R}^3$, such that $\mathcal{C}_{ss} \neq 0$, $\mathcal{C}_s \neq 0$ (i.e., without singular points of order 0 and 1 [4]) the Frenet frame is given by

$$\mathcal{T}_s = \kappa \mathcal{N}, \quad \mathcal{N}_s = -\kappa \mathcal{T} - \tau \mathcal{B}, \quad \mathcal{B}_s = \tau \mathcal{N}, \quad \frac{\partial}{\partial s} = \frac{1}{\|\mathcal{C}_p\|} \frac{\partial}{\partial p}.$$

⁵ Tensor derived measures other than principal diffusion direction are typically only used as tract termination criteria.

$\mathcal{T} = \frac{\mathcal{C}_p}{\|\mathcal{C}_p\|}$ is the unit tangent vector, \mathcal{N} and \mathcal{B} are the normal and the binormal, κ and τ denote curvature and torsion respectively, and s denotes arc-length. See Fig. 2 for a depiction of the Frenet frame. Computing \mathcal{T} from \mathcal{C} is immediate. Computing \mathcal{N} yields $\mathcal{B} = \mathcal{T} \times \mathcal{N}$ and thus the desired local coordinate frame.

In this paper the space curve is given by a representative fiber tract. For the experiments of Sec. 8 streamline tractography was used to compute the representative fiber. For a more robust approach, streamlining should be replaced by an optimal path method [8]. In what follows, a known representative fiber tract is assumed.

Since the Frenet frame is based on differential properties of the space curve it is sensitive to noise. Instead of using the Frenet frame directly, the frame diffusion is instead based on a regularized version of the Frenet frame. Fig. 2 (right) shows a progressively more regularized Frenet frame. Note that for the reorientation of diffusion information (see Sec. 5) the Frenet axes can be flipped. All computations in this paper identify antipodal directions; derivatives are computed by prealigning all vector-valued quantities locally before derivative computation.




Fig. 2. The Frenet frame $\{\mathcal{T}, \mathcal{N}, \mathcal{B}\}$ consisting of tangent, normal, and binormal to \mathcal{C} . Regularization helps to obtain smoothly varying frames from noisy data (right).

4 Frame Diffusion

Instead of declaring a point in space to correspond to its closest point (measured by Euclidean distance) on the representative tract, here, correspondence is established implicitly through a diffusion process. This allows for smoother correspondences avoiding orientation jumps which occur at shock points for the Euclidean distance map. Since orientation is the quantity of interest, the orientation information is diffused *away* from the representative tract. Tschumperle and Deriche [22] regularize diffusion tensor fields by evolutions on frame fields. This can be used to define the diffusion of the frame field off the reference tract. Formally,

$$\mathbb{F}_\theta(\mathbf{x}, \theta) = \Delta_{\mathbf{x}} \mathbb{F}, \quad \mathbf{x} \in \Omega \setminus \mathcal{C}, \quad \mathbb{F}(\mathbf{x}, \theta) = \mathbb{F}_b, \quad \mathbf{x} \in \mathcal{C}, \quad (1)$$

where $\mathbb{F} = \{\mathcal{T}^a, \mathcal{N}^a, \mathcal{B}^a\}$ is the set of the axes implied by the regularized Frenet frame, \mathbb{F}_b denotes the boundary condition given by the Frenet-frame-implied axes on the tract, $\mathbf{x} \in \mathbb{R}^3$ denotes spatial position, θ artificial evolution time, and $\Delta_{\mathbf{x}} = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}$ the spatial Laplacian operator. The frame diffusion

problem (1) can be solved [22] by evolving a set of three coupled vector diffusions:

$$\begin{cases} \mathcal{T}_\theta &= \Delta\mathcal{T} - (\Delta\mathcal{T} \cdot \mathcal{T})\mathcal{T} - (\Delta\mathcal{N} \cdot \mathcal{T})\mathcal{N} - (\Delta\mathcal{B} \cdot \mathcal{T})\mathcal{B}, \\ \mathcal{N}_\theta &= \Delta\mathcal{N} - (\Delta\mathcal{T} \cdot \mathcal{N})\mathcal{T} - (\Delta\mathcal{N} \cdot \mathcal{N})\mathcal{N} - (\Delta\mathcal{B} \cdot \mathcal{N})\mathcal{B}, \\ \mathcal{B}_\theta &= \Delta\mathcal{B} - (\Delta\mathcal{T} \cdot \mathcal{B})\mathcal{T} - (\Delta\mathcal{N} \cdot \mathcal{B})\mathcal{N} - (\Delta\mathcal{B} \cdot \mathcal{B})\mathcal{B} \end{cases}$$

which may be rewritten [22] as the rotations

$$\mathcal{T}_\theta = R \times \mathcal{T}, \quad \mathcal{N}_\theta = R \times \mathcal{N}, \quad \mathcal{B}_\theta = R \times \mathcal{B},$$

where $R = \mathcal{T} \times \Delta\mathcal{T} + \mathcal{N} \times \Delta\mathcal{N} + \mathcal{B} \times \Delta\mathcal{B}$ and $\mathbb{F} = \{\mathcal{T}^a, \mathcal{N}^a, \mathcal{B}^a\}$ is given by identifying antipodal directions. While the statistics used for the segmentation in Sec. 7 only use directional information, diffusing the complete frame information specifies a local rotation. This allows for easy extension of the methodology to formulations using for example the full tensor information or orientation distribution functions. Fig. 3 shows two 2D examples of frame diffusion. The resulting diffused frame field is smoother. Interestingly, the partial half-circle example shows that, to a limited extent, frame diffusion can be used to fill in missing information. This is a useful feature in case it is not possible to obtain one connected representative fiber tract.

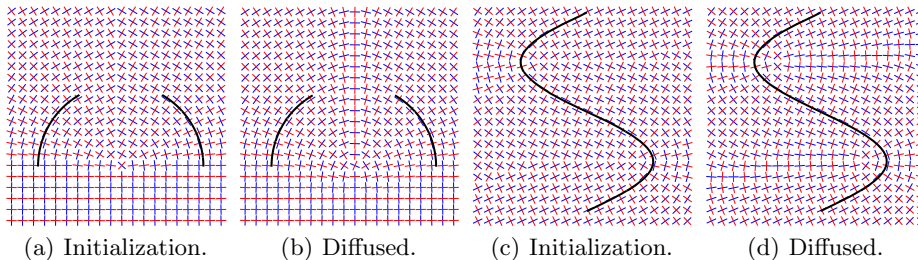


Fig. 3. Frame diffusion yields smooth frame fields and thus smooth reorientations. Initializations using Euclidean distance point correspondences show frame discontinuities.

5 Frame Reorientation

The diffused frames can be used to reorient diffusion measurements locally to a canonical frame M^6 . This reorientation can be applied to any representation of diffusion information, e.g., the diffusion tensor, orientation distribution functions, etc. For clarity, reorientation is explained here for the case of diffusion tensors T . Given the diffused frame $\{\mathcal{T}, \mathcal{N}, \mathcal{B}\}$ and the associated rotation matrix $F = [\mathcal{T}, \mathcal{N}, \mathcal{B}]$ a tensor T is reoriented by applying the relative rotation MF^T , i.e., by

$$T^r = MF^T T F M^T.$$

The tensor reorientation yields tight tensor statistics while allowing a segmentation algorithm to apply spatial regularizations in the original space. It greatly simplifies computations by avoiding an explicit warping to straighten a curved fiber bundle.

⁶ See Sec. 6 for a way to determine the canonical frame automatically

6 Orientation Statistics

We now describe the probabilistic modeling of fiber bundle orientations.

6.1 Watson Distribution

The Watson distribution is one of the simplest distributions for directional random variables [24, 2, 20]. It is radially symmetric around a mean direction μ , with a spread controlled by the concentration parameter k .

The Watson distribution on the unit sphere S^2 has probability density

$$p_w(q|\mu, k) = \frac{1}{4\pi {}_1F_1(\frac{1}{2}; \frac{3}{2}; k)} e^{k(\mu^T q)^2}, \quad p_w(q|\mu, 0) = \frac{1}{4\pi},$$

where μ is the mean direction vector, k the concentration parameter⁷, $q \in S^2$ is a direction represented as a column vector, and ${}_1F_1(\cdot; \cdot)$ denotes the confluent hypergeometric function. The Watson distribution is bipolar for $k > 0$, with maxima at $\pm\mu$ and uniform for $k = 0$. To model the interior of a fiber bundle, μ is set to the tangential direction of the canonical frame M . Reorienting diffusion information results in a tight Watson distribution with large concentration parameter k . The statistics outside the fiber bundle are modeled using the uniform distribution, since no preferred direction can be assumed in general in the fiber exterior.

While it is possible to use more complicated probability distributions (e.g., the Bingham distribution, or distributions on the diffusion tensor directly) to model a fiber tract orientation distribution, the Watson distributions chosen (in conjunction with the reorientation scheme) have the advantage of modeling the interior and the exterior of the fiber bundle with only one free parameter, the concentration k , greatly simplifying the estimation task and allowing for an easy interpretation of the estimated probability distribution.

6.2 Parameter estimation for the Watson distribution

The distribution parameters k and μ are easy to estimate. Given a set of N points $\mathbf{q}_i \in S^2$ (written as column vectors and representing spatial directions), the maximum likelihood estimate of μ is the major eigenvector of the sample covariance [5] $C = \frac{1}{N} \sum_{i=1}^N \mathbf{q}_i \mathbf{q}_i^T$ and $1 - \lambda_1$ (with λ_1 the largest eigenvalue of C) is the maximum likelihood estimate of $\frac{1}{k}$. Estimation of μ is performed only as a means of estimating the canonical frame direction and computed only on the representative tract. It is assumed fixed throughout the segmentation process described in Sec. 7. Only the concentration parameter k is estimated during bundle segmentation. For increased estimation robustness, robust estimators for the concentration parameter k may be used to account for cases where orientation measurements are either incorrect or cannot be reliably determined (as for example for isotropic tensors).

⁷ To avoid ambiguities the concentration is denoted as k ; κ denotes curvature in this paper.

7 Segmentation

We now integrate the diffusion data reorientation method and the statistical modeling described in Sec. 6 within a probabilistic version of the Chan-Vese [3] segmentation framework [7] using the probability distributions of Sec. 6.

7.1 Optimization Problem

The probabilistic Chan-Vese segmentation approach [3] is a piecewise-constant approximation problem, minimizing the energy functional

$$E_{cv}(\Omega_i, p_1, p_2) = \int_{\partial\Omega_i} ds + \lambda \int_{\Omega_i} (-\log p_1(f(x)) d\Omega + \lambda \int_{\Omega \setminus \Omega_i} (-\log p_2(f(x)) d\Omega, \quad (2)$$

where $f(\cdot)$ denotes an image feature (here, direction), p_1 and p_2 are the likelihoods for the interior and the exterior of the segmentation respectively, Ω is the computational domain and Ω_i is the interior domain. Choosing

$$p_1(q) = p_w(q|\mu, k), \quad p_2(q) = p_w(q|\mu, 0) = \frac{1}{4\pi},$$

constitutes the segmentation approach. See Sec. 6.1 for a discussion of this choice.

7.2 Numerical Solution

According to a slight modification of the solution approach in [19], the *probabilistic* Chan-Vese energy minimization problem 2 (on log-likelihood functions instead of image intensities) can be recast as the minimization of

$$E_{cvb}(u, c_1, c_2) = \int_{\Omega} \|\nabla u(x)\| d\Omega + \int_{\Omega} \lambda r_1(x, c_1, c_2)u + \alpha \nu(u) d\Omega \quad (3)$$

where

$$\nu(\zeta) = \max\{0, 2|\zeta - \frac{1}{2}| - 1\}, \quad (\text{the exact penalty function}),$$

$$r_1(x, c_1, c_2) = \log \frac{p_2(f(x))}{p_1(f(x))} = \log \left({}_1F_1\left(\frac{1}{2}; \frac{3}{2}; k\right) - k(\mu^T q)^2 \right).$$

The boundary is recovered as $\Omega_i = \{x : u(x) > \xi\}$, $\xi \in [0, 1]$. Eq. 3 can be solved efficiently through a dual formulation of the total-variation norm [19]:

1) Solve for u keeping v fixed:

$$\min_u \left\{ \int_{\Omega} \|\nabla u\| dx + \frac{1}{2\theta} \|u - v\|_{L^2}^2 \right\} \quad (4)$$

2) Solve for v keeping u fixed:

$$\min_v \left\{ \frac{1}{2\theta} \|u - v\|_{L^2}^2 + \int_{\Omega} \lambda r_1(x, p_1, p_2)v + \alpha \nu(v) dx \right\} \quad (5)$$

VIII

3) Repeat until convergence.

Eq. 5 has the solution $v = \min\{\max\{u(x) - \theta\lambda r_1(x, p_1, p_2), 0\}, 1\}$ and Eq. 4 can be solved using a fixed-point iteration

$$u = v - \theta \operatorname{div} p, \quad p^{n+1} = \frac{p^n + \delta t \nabla(\operatorname{div}(p^n) - \frac{v}{\theta})}{1 + \delta t |\nabla(\operatorname{div}(p^n) - \frac{v}{\theta})|}, \quad p = (p^1, p^2, p^3), \quad \delta t \leq \frac{1}{6}.$$

To enforce segmenting a bundle containing the representative tract set

$$\begin{cases} v = 1 & \text{for all points on the representative tract,} \\ v = 0 & \text{for all points at a distance } d \geq d_{max} \text{ from the representative tract.} \end{cases}$$

The segmented fiber bundle is the set of voxels with $u \geq \frac{1}{2}$ which are contained in the connected component containing the voxels of the representative tract. This is also the volume which is used throughout the evolution to update the estimation of the concentration k of the fiber bundle's Watson distribution.

8 Results

This section gives results for the proposed segmentation approach. Synthetic examples are discussed in Sec. 8.1. Sec. 8.2 presents results for a real DW-MRI of the brain and compares them to segmentation results obtained through streamline tractography.

8.1 Synthetic example

A synthetic tensor example was generated. Tensors are assumed of uniform shape with eigenvalues $(1.5, 0.5, 0.5)e - 3$ oriented along a circular path to model a fiber bundle. Tensors oriented orthogonally to the circular path model the outside. Diffusion weighted images were generated using the Steijskal Tanner equation $S_k = S_0 e^{-b \mathbf{g}_k^T T \mathbf{g}_k}$, where S_k denotes the diffusion weighted image acquired by applying a gradient direction \mathbf{g}_k with b-value b , and T the diffusion tensor. Parameters were $S_0 = 1000$, $b = 1000$ with 46 gradient directions equally spaced on the unit sphere. Rician noise of $\sigma = 70$ was introduced to the baseline image S_0 (non-diffusion weighted) and the diffusion weighted images S_k . Fig. 4 shows the original data and the resulting segmentation on the top row (with the streamline indicating the computed representative tract) and the reoriented data with associated segmentation on the bottom row. For this synthetic example, reorientation results in an almost perfectly uniform tensor distribution on the inside and the outside of the simulated fiber bundle. Consequently, while the proposed approach fails at segmenting the original data, it segments the reoriented data well. Note, that the failure to segment the original data is not merely a result of the segmentation method employed. Any segmentation relying purely on region-based statistics will either have to include some of the background in its bundle segmentation or will severely under-segment the bundle itself, since

background and foreground are not clearly separable based on global statistics. While including edge-based terms may improve the segmentation of the original data, regional terms will be of limited use and will locally counteract the edge influence requiring a delicate balance between region-based and edge-based energies to faithfully segment the simulated fiber bundle.




Fig. 4. Synthetic segmentation example overlaid on a color by orientation representation. Reorientation results in a successful segmentation ($k = 10$, $\theta = 0.01$, $\lambda = 0.7$).

8.2 Real example

The real example was computed for the cingulum bundle using a 3T DW-MRI upsampled to isotropic resolution (0.93 mm^3) with 8 baseline images and 51 uniformly distributed gradient directions ($b=586$). The representative tract was computed using streamline tractography.

Fig. 5 shows color by orientation representations for a sagittal slice through the brain with the cingulum bundle (mainly in green) before and after reorientation. The reoriented image shows a consistently green cingulum bundle, whereas in the original image the cingulum bundle is colored blue when wrapping posteriorly around the corpus callosum, indicating a change of orientation from anterior-posterior to superior-inferior. Example segmentation results of the proposed approach are shown for the reoriented and the original data. Algorithm parameters were set to $\theta = 0.01$, $\lambda = 0.5$. The concentration parameter was set to $k = 100$ and converged to $k = 19.5$ throughout the evolution for the reoriented dataset. The surface models generated from the computed segmentations show that the segmentation for the reoriented data approximates the cingulum bundle more faithfully.

Finally, to demonstrate the strength of the reorientation approach, Fig. 6 gives an example for the cingulum bundle segmentation at a posterior slice of the cingulum bundle where the cingulum bundle wraps around the corpus callosum. While in the reoriented case the segmentation is successful and the direction of the cingulum bundle is uniform (green), the segmentation on the original data fails in this part of the fiber bundle.

To compare the proposed methods to alternative segmentation approaches, the cingulum bundle was segmented using a region of interest based approach (the same regions of interest used to generate the representative fiber tract for reorientation). Two small axial regions of interest were defined for the cingulum bundle (superiorly to the corpus callosum). Streamline tractography with

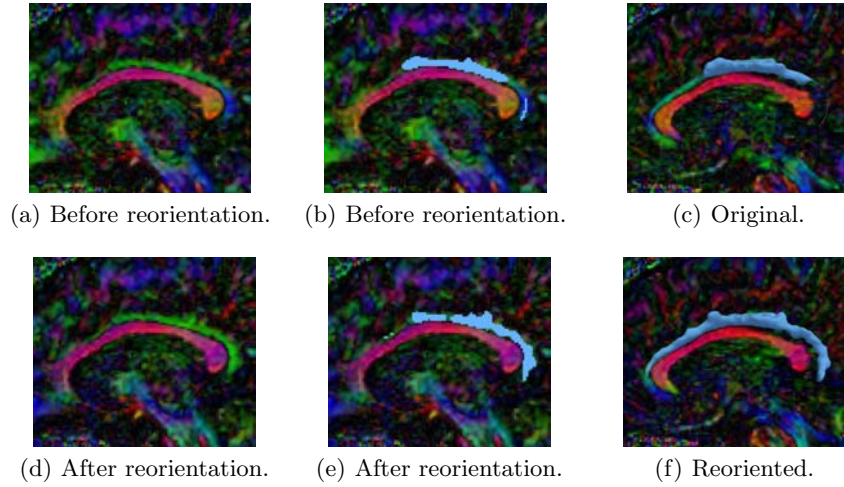


Fig. 5. Sagittal slice of the cingulum bundle, before and after tensor reorientation. The cingulum bundle appears more uniform in direction (green) after reorientation. Reorientation greatly improves the segmentation result of the proposed approach.




Fig. 6. Posterior coronal slice: Color by orientation shows advantage of reorientation. voxelization, full brain streamline tractography with voxelization, as well as segmentation on the original and reoriented data using the proposed approach was performed. Fig. 7 illustrates segmentation results for these methods for coronal slices in the superior part of the cingulum bundle (where the cingulum bundle is *not* strongly curved). As expected streamline tractography and full brain streamline tractography mainly capture the interior of the fiber bundle, with full brain tractography performing qualitatively better than standard region of interest based streamline tractography (streamlines were seeded one per voxel in the regions of interest). The proposed segmentation approach captures the cingulum bundle well for the reoriented and for the original data, showing the utility of segmenting in orientation space. However, the reoriented segmentation results are better where the cingulum bundle curves strongly, as shown in Fig. 6.

9 Conclusion and Discussion

This paper proposed a new segmentation method for tubular fiber bundles. It is based on reorientation of diffusion measurements resulting in more uniform data




Fig. 7. Superior coronal slices: Original data; results for streamline (S) and full brain streamline (FBS) tractography, for the proposed segmentation on original data (O) and on reoriented data (RO). Only the proposed approach segments up to the perceived bundle boundary in orientation space.

distributions inside the fiber bundle of interest. Segmentation is performed by an efficient convex approximation of the probabilistic Chan-Vese energy using region-based directional statistics. The approach compares favorably to streamline approaches for bundle segmentation. However, since no realistic ground-truth for DW-MRI is available to date quantitative analysis is difficult. Extensions to sheet-like structures are conceivable, where the representative tract would be replaced by a representative sheet [12, 13] (using the major diffusion direction combined with the normal to the medial sheet to define a frame for reorientation). Population studies could be performed by either performing segmentation in atlas space, or by using an atlas defined representative tract and subject-specific bundle segmentations.

Acknowledgments This work is part of the National Alliance for Medical Image Computing (NAMIC), funded by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>.

References

1. D. Le Bihan. Looking into the functional architecture of the brain with diffusion MRI. *Nature Neuroscience*, 4(6):469–480, 2003.
2. C. Bingham. An antipodally symmetric distribution on the sphere. *The Annals of Statistics*, 2(6):1201–1225, 1974.
3. T. F. Chan and L.A. Vese. Active contours without edges. *TIP*, 10(2):266–277, 2001.
4. M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice Hall, 1976.
5. A. Schwartzman et al. False discovery rate analysis of brain diffusion direction maps. *Annals of Applied Statistics*, 2(1):153–175, 2008.
6. C. Lenglet et al. DTI segmentation by statistical surface evolution. *TMI*, 25(6):685–700, 2006.
7. D. Cremers et al. A review of statistical approaches to level set segmentation: Integrating color, texture, motion and shape. *IJCV*, 72(2):195–215, 2007.
8. E. Pichon et al. A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography. In *MICCAI*, pages 180–187, 2005.
9. J. Melonakos et al. Locally-constrained region-based methods for DW-MRI segmentation. *ICCV*, pages 1–8, 2007.
10. L. Jonasson et al. White matter fiber tract segmentation in DT-MRI using geometric flows. *MEDIA*, 9(3):223–236, 2005.
11. O. Friman et al. A Bayesian approach for stochastic white matter tractography. *TMI*, 25(8), 2006.
12. P. A. Yushkevich et al. Structure-specific statistical mapping of white matter tracts using the continuous medial representation. In *ICCV*, pages 1–8, 2007.
13. S. M. Smith et al. Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31:1487–1505, 2006.
14. S. P. Awate et al. A fuzzy, nonparametric segmentation framework for DTI and MRI analysis: With applications to DTI-tract extraction. *TMI*, 26(11):1525–1536, 2007.
15. T. McGraw et al. Segmentation of high angular resolution diffusion MRI modeled as a field of von Mises-Fisher mixtures. *ECCV*, pages 463–475, 2006.
16. T.E.J. Behrens et al. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *MRM*, 50(5):1077–1088, 2003.
17. W. K. Jeong et al. Interactive visualization of volumetric white matter connectivity in DT-MRI using a parallel-hardware Hamilton-Jacobi solver. *TVCG*, pages 1480–1487, 2007.
18. Wiegell et al. Automatic segmentation of thalamic nuclei from diffusion tensor magnetic resonance imaging. *Neuroimage*, 19(2):391–401, 2003.
19. X. Bresson et al. Fast global minimization of the active contour/snake model. *Journal of Mathematical Imaging and Vision*, 28(2):151–167, 2007.
20. K. V. Mardia. Statistics of directional data. *Journal of the Royal Statistical Society. Series B.*, 37(3):349–393, 1975.
21. L. J. O’Donnell and C.F. Westin. Automatic tractography segmentation using a high-dimensional white matter atlas. *TMI*, 26(11):1562–1575, 2007.
22. D. Tschumperle and R. Deriche. Diffusion tensor regularization with constraints preservation. In *CVPR*, volume 1, pages 948–953, 2001.
23. Z. Wang and B. C. Vemuri. DTI segmentation using an information theoretic tensor dissimilarity measure. *TMI*, 24(10):1267–1277, 2005.
24. G. Watson. Equatorial distributions on a sphere. *Biometrika*, 52:193–201, 1965.