

Topological Visualization of Brain Diffusion MRI Data

Thomas Schultz, Holger Theisel, and Hans-Peter Seidel

Abstract—Topological methods give concise and expressive visual representations of flow fields. The present work suggests a comparable method for the visualization of human brain diffusion MRI data. We explore existing techniques for the topological analysis of generic tensor fields, but find them inappropriate for diffusion MRI data. Thus, we propose a novel approach that considers the asymptotic behavior of a probabilistic fiber tracking method and define analogs of the basic concepts of flow topology, like critical points, basins, and faces, with interpretations in terms of brain anatomy. The resulting features are fuzzy, reflecting the uncertainty inherent in any connectivity estimate from diffusion imaging. We describe an algorithm to extract the new type of features, demonstrate its robustness under noise, and present results for two regions in a diffusion MRI dataset to illustrate that the method allows a meaningful visual analysis of probabilistic fiber tracking results.

Index Terms—Diffusion tensor, probabilistic fiber tracking, tensor topology, uncertainty visualization.

1 INTRODUCTION

In diffusion-weighted MRI (magnetic resonance imaging), signal intensity is modulated by the Brownian motion of water molecules [28]. In the human brain, this allows conclusions about tissue microstructure, since it restricts molecular motion [21]. A popular variant of such diffusion imaging is diffusion tensor MRI (DT-MRI), which derives a second-order tensor field to model the apparent diffusivities from a series of measurements. Since only some of the methods discussed in this paper employ a tensor model, we prefer the more general term “diffusion MRI” and only talk about DT-MRI when referring to a second-order tensor model as suggested by Basser et al. [2].

Since their introduction by Helman and Hesselink [14], the concise representations generated by topological methods have become a powerful tool for the visualization of vector fields describing fluid flows (cf. [26] for an overview). These methods partition the domain into regions in which all streamlines connect the same source to the same sink, or in other words, into regions in which the flow exhibits the same asymptotic behavior. The resulting topological skeleton reduces the flow to the structurally significant information.

There exists a variety of fiber tracking techniques to investigate connectivity within the brain based on diffusion MRI data (cf. the review in [22]). The fact that connectivity is a fundamental topological notion suggests that a topological visualization of diffusion MRI data may be beneficial. It is the goal of the present paper to find such a method.

A natural starting point for our investigation is the work by Delmarcelle and Hesselink [10], who have generalized the concepts of topological vector field visualization to second-order tensor fields. Based on their fundamental definitions, further research has been conducted on 3D tensor topology [15, 30, 31]. More specifically, Zheng et al. [32] have argued that applying tensor topology to DT-MRI is likely to prove beneficial. While they expect noise artifacts to dominate a naïve topological visualization of DT-MRI data sets, they suggest that additional selection of the most important features would produce a “simple yet powerful representation” [31]. However, no results from applying tensor topology to DT-MRI have been published so far. In Section 2, we discuss the interpretation of the features from tensor topology and present both experimental results and theoretical argu-

ments which suggest that topological features are, unfortunately, not useful in the context of DT-MRI.

After reviewing more related work in Section 3, we introduce a new paradigm for transferring the basic notions of topological flow visualization to diffusion MRI data in Section 4. In contrast to the existing “tensor topology”, we refer to it as “diffusion MRI topology” to reflect the fact that we neither restrict ourselves to a second-order diffusion tensor model, nor do we expect our approach to be useful for tensor fields that describe different phenomena (e.g., stress tensors [30]). In particular, we do not question the fact that tensor topology holds the potential to extract interesting features from generic tensor fields. In Section 5, we propose a method that can be used to extract the new type of features. In Section 6, we demonstrate the robustness of our method under noise and present additional experimental results to illustrate that the novel features allow a meaningful interpretation of the data. Finally, in Section 7, we conclude the paper and discuss possible directions for future work.

2 DEGENERATE LINES IN DT-MRI FIELDS

In topological flow visualization, critical points play a central role. They are the points at which the vector field magnitude vanishes and the only locations at which streamlines intersect. The expressive power of topological flow visualizations is owed to the clear physical meaning of the critical points: they can be classified as sinks, sources, and saddles, which are of distinct importance in flow fields.

Tensor topology as defined by Delmarcelle and Hesselink [10] is the topology of hyperstreamlines, i.e., the integration lines of eigenvectors. The analogs of critical points are now “degenerate” locations in which at least two of the tensor’s eigenvalues are equal. At these loci, the corresponding eigenvectors become ill-defined and hyperstreamlines intersect. At type P (planar) degeneracies, the larger two eigenvalues are equal (major and medium hyperstreamlines intersect), while type L (linear) features involve the smaller eigenvalue pair. Zheng et al. [30] have proven that in generic 3D tensor data, type L and type P features form stable lines, and they have presented several algorithms to extract them.

Unfortunately, the interpretation of these features differs from the critical points in flow fields: connectivity in DT-MRI can only be inferred in a probabilistic sense. The major hyperstreamlines of a diffusion tensor field can be interpreted as maximum likelihood pathways [6], but type P features are merely locations in which no single direction has maximum likelihood, not locations in which the pathway “ends” as does a streamline at a sink.

Even if this means that degenerate lines have limited relevance for the topology of neuronal fiber pathways, they may still constitute an interesting tool for the analysis of DT-MRI data if they provide stable features in practice. The following subsection explores this potential.

- Thomas Schultz is with MPI Informatik, Saarbrücken, Germany, E-mail: schultz@mpi-inf.mpg.de.
- Holger Theisel is with BieGraph Group, Bielefeld University, Germany, E-mail: theisel@techfak.uni-bielefeld.de.
- Hans-Peter Seidel is with MPI Informatik, Saarbrücken, Germany, E-mail: hpseidel@mpi-inf.mpg.de.

Manuscript received 31 March 2007; accepted 1 August 2007; posted online 2 November 2007.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org.

2.1 Experimental Setup

For our experiments, we have implemented the prediction-correction scheme based on discriminant constraint functions and Hessian factorization, as described in [30]. To obtain the best possible results even under difficult conditions, we allow a large number of Newton-Raphson iterations in the correction phase and repeat failed steps with an extremely small stepsize.

Our dataset consists of diffusion-weighted images (DWIs) acquired on a Siemens 3T Trio Scanner at $b = 1000\text{s/mm}^2$ in 60 isotropically distributed gradient directions (3 averages each), plus one non-diffusion weighted T_2 image (7 averages), voxel size 1.72mm (isotropic). We received the images pre-registered to compensate motion and imaging artifacts. We estimated diffusion tensors via a simple least-squares fit on the logarithm of signal intensities [2].

To avoid visual clutter, we limit our analysis to a region of interest which spans $21 \times 29 \times 14$ voxels at the center of the corpus callosum. As suggested in [32], we only consider regions of sufficient anisotropy ($FA \geq 0.2$). FA is the fractional anisotropy [4], defined for a diffusion tensor \mathbf{D} as

$$FA(\mathbf{D}) = \sqrt{\frac{3}{2} \frac{\|\mathbf{D} - \frac{1}{3}\text{tr}(\mathbf{D})\mathbf{I}\|}{\|\mathbf{D}\|}} \quad (1)$$

The tractography in Figure 1 (a), obtained by major eigenvector integration and standard XYZ-RGB coloring, shows the corpus callosum (in red) from a superior point of view, the cingulum bundles (in green) and a small part of the pyramidal tract (in dark blue, at the right and left image boundaries).

In our experiment, we examine the line features and their robustness under noise. While type L features are not part of the major hyper-streamline topology, Zheng et al. [32] have suggested that they may be of particular interest for DT-MRI, so we include them in our analysis.

At low noise levels (with a signal-to-noise ratio $SNR \geq 3$), noise in magnitude MR images can be approximated with a Gaussian distribution [13], so we obtain noisy datasets by adding Gaussian noise to the DWI and T_2 images and re-estimating the tensors. The standard deviation is chosen as $\sigma = A/SNR$ with $SNR \in \{12, 8\}$, where A is the average of signal intensities within the white matter mask.

2.2 Practical Results

The degenerate lines in Figure 1 are colored using the same XYZ-RGB scheme as the tractography to facilitate orientation. For type P/L, the color indicates the minor/major eigenvector direction of the tensor field. The degenerate lines themselves are not in general aligned with any eigenvector direction, so the color coding does not indicate the direction of the degenerate features. Rather, red type L features are located within the corpus callosum, green ones in the cingulum bundle, and blue ones are in the pyramidal tract.

Since our dataset has ten times the minimum number of DWIs required to estimate the tensors, the noise can be considered low and moderate, which is reflected by the fact that the major features remain discernible in the tractography at all noise levels. Still, the degenerate features change significantly, especially those of type L. The results may not rule out the possibility of selecting a set of type P features which remain recognizable under noise. However, we used a smooth (C^2) B-spline approximation of the tensor data [23], which stabilizes feature extraction. Figure 2 illustrates the effect of using trilinear interpolation instead and exhibits significant differences, even when compared to the results from exactly the same data in Figure 1.

In order to assert that the encountered instabilities neither indicate a general flaw in the concept of tensor topology, nor an error in our implementation, we finally present results on a randomly generated dataset similar to the one used by Zheng et al. [30]. Figure 3 shows both type L (cool colors) and type P features (warm colors). In this case, changing the interpolation scheme alters the exact shape of the features slightly, but generally leaves them well-recognizable.

2.3 Interpretation

From these experiments, we conclude that DT-MRI fields cannot be regarded as generic second-order tensor fields in the sense that tensor

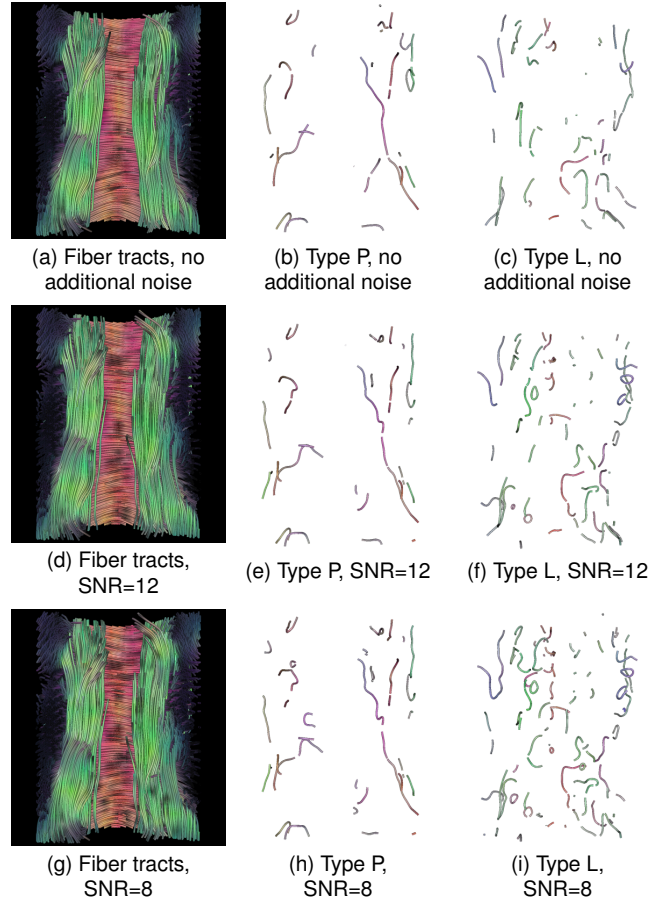


Fig. 1. A comparison of type P and type L features under Gaussian noise shows significant changes for even moderate noise levels.




Fig. 2. Type P features with linear interpolation instead of B-spline approximation as in Figure 1. Feature lines depend significantly on the choice of interpolation.




Fig. 3. In a generic dataset, the degenerate lines are far less affected by the choice of interpolation.

topology requires. This also explains the high number of short and broken feature lines in Figures 1 and 2, which indicate that the loci of degeneracy do not in general form stable lines in DT-MRI data. Moreover, the extracted features do not correlate with any known structures in the data, which makes their interpretation difficult.

An explanation for this discouraging result may be found in an approach by Behrens et al. [6]. Instead of modeling the apparent diffusivities (like in DT-MRI), they create a fiber model which predicts a diffusivity profile from a set of fiber parameters and estimate a posterior distribution of these parameters within a Bayesian framework. In our context, the crucial aspect is that the predicted profiles from a single-fiber model always correspond to a linear degeneracy. In other words, we can expect regions where the model applies to be densely filled with type L features whose exact location will depend on factors outside the model (like artifacts from noise and interpolation). On the other hand, we cannot expect hyperstreamline topology to be beneficial in regions where the single-fiber model breaks down, since considering major hyperstreamlines implicitly assumes such a model.

3 RELATED WORK

In the previous section, we reviewed tensor topology and its potential in the context of DT-MRI. Before we proceed, we will now discuss works which are related to our alternative paradigm for topological diffusion MRI visualization.

Our new method depicts fiber pathways as a whole. While Enders et al. [11] have followed a similar goal by wrapping clustered streamlines, we use a completely different approach. Rather than clustering streamlines from a deterministic fiber tracking method, we first partition grey matter voxels based on the results of a probabilistic method and infer the pathways that connect them only in a second step. Moreover, our final visualization does not involve any streamlines.

Jonasson et al. [18] have segmented fiber tracts as a whole, but aim more at the interactive segmentation of specific structures than at the visualization of the dataset. Their approach relies on the placement of an initial seed for a surface growing algorithm, which is driven by the similarity of diffusion tensors in adjacent voxels and does not determine connectivity explicitly.

In their work on anisotropy creases, Kindlmann et al. [19] have demonstrated that bounding surfaces between fiber bundles, which could be considered a complement of the topological faces in our method, can often be found from anisotropy alone without considering the connectivity that underlies the topological notions. However, this analogy breaks down when fiber bundles are only distinguished by their connectivity. For example, both our approach and streamline clustering partition the *corpus callosum* into several sections (cf. Figures 6 and 7), while anisotropy creases do not reflect this subdivision.

Our paradigm for topological diffusion MRI visualization draws on methods which have recently been introduced for connectivity-based cortex parcellation studies in the neuroscience community [17, 1]. These works show that changes in connectivity profiles allow the partitioning of grey matter into functionally distinct regions; our focus is to construct a novel visualization method based on this insight. Moreover, existing approaches do not consider the asymptotic behavior of the employed tractography methods, so they do not constitute a topological analysis.

Finally, as a result of using probabilistic tractography, the topological features we suggest are fuzzy, which reflects the uncertainty inherent in the inferred connectivity. In flow visualization, uncertainty has not yet played a major role. Salzbrunn and Scheuermann [24] have recently introduced “fuzzy” streamline predicates as a means to define characteristic sets of predicates for which it is algorithmically difficult to calculate them directly. However, they do not use them to visualize uncertainty. To the best of our knowledge, a fuzzy topology which conveys the confidence level of region boundaries to the user, has not yet been considered.

4 TOPOLOGICAL FEATURES IN DIFFUSION MRI DATA

Previous research on tensor topology has started from mathematical analogies [10], which is appropriate to define stable features in generic

tensor data. In this work, we are concerned with finding features that have a meaningful interpretation in the context of our particular type of data, so we choose brain anatomy as the starting point of our reasoning. Axons, which form the white matter pathways whose connectivity we would like to investigate, have an orientation: they start at a cell soma and end in a synapse. However, diffusion imaging does not reveal this polarity, so we cannot distinguish if a connection endpoint is a source or a sink; note that tensor topology does not make this distinction either.

Critical points in flow topology are an instance of the more general notion of limit sets: They are locations in which a streamline integration starts or ends. In general, such limit sets do not necessarily form points. For example, the degenerate locations in 3D tensor topology form lines. Within the scope of diffusion images of the brain, neuronal pathways end at surfaces, namely, at the interfaces between grey and white matter or between white matter and the boundary of the domain. Recently, so-called cortex parcellation studies have shown that to a certain extent, functionally distinct regions within grey matter can be found by considering changes in their connectivity profile [17, 1]. We will call connected regions of uniform connectivity, which are likely to represent anatomically meaningful units, *critical regions*, and identify them as the suitable limit sets for our diffusion MRI topology.

As discussed in Section 2, the endpoints of streamlines that result from deterministic fiber tracking methods [3] do not necessarily coincide with endpoints of the underlying neuronal pathways, so we do not consider them appropriate for defining a diffusion MRI topology. Instead, we base our analysis on the asymptotic behavior of a probabilistic fiber tracking approach [20] that employs the widely used diffusion tensor model and will be summarized in Section 5.1. Alternative methods, which may or may not depend on diffusion tensors (e.g., [6]), could be plugged into our framework, making its use independent of the preferred choice of diffusion and fiber models.

4.1 Critical Regions and Basins

The fact that the selected fiber tracking method provides a probabilistic connectivity measure has to be reflected in the definition of topological features from its results. In topological flow visualization, the α -basin of a source is the union of all streamlines that emerge from it. Accordingly, the ω -basin of a sink is the union of streamlines that end in it [25]. Analogous to these notions, we define the p -basin of a critical region as the set of points from which a probabilistic tractography reaches that region with probability $P \geq p$. For a point that connects two regions, we expect that around half of the particles end in each region, so we typically consider p -basins with $p < 0.5$.

To clarify these basic notions visually, we present some examples obtained with the method from Section 5, on the same region of interest as in Figure 1. Since it is taken from the center of the brain, the critical regions segment the domain boundaries rather than the cortex. Figure 4 (a) shows the deterministic tractography from a posterior/left viewpoint and a sample critical region as a yellow surface. It corresponds to the left endpoints of the fibers that pass through the central part of the corpus callosum and extends to a portion of the internal capsule. This is understandable, since fibers from both structures intermingle in this region and are not cleanly separated anatomically.

Figure 4 (b) presents the same critical region with its 0.4-basin instead of the tractography. To provide a confidence interval, the 0.25-basin is rendered transparently. As expected, the basin extends over the central part of the corpus callosum and down towards the internal capsule. For the XYZ-RGB color coding of basins and faces, a weighted average is computed from the tensors within the corresponding structure, with the local probabilities as weights. Thus, the purple color of the basin indicates the mixture of fibers that run through the corpus callosum (red) and the internal capsule (blue).

4.2 Faces

In flow visualization, one is typically interested in the *faces* which result from all intersections of α - and ω -basins. These are regions of uniform asymptotic flow behavior, i.e., regions in which all streamlines emerge from the same source and end in the same sink. Taken

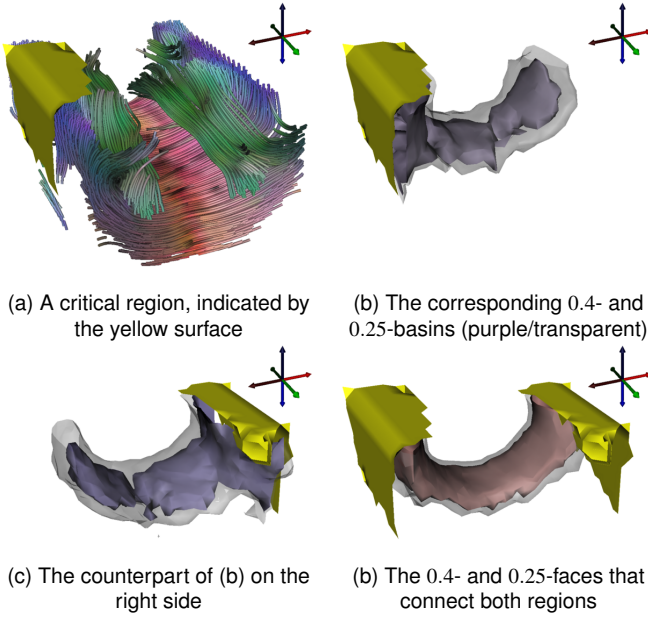


Fig. 4. The basin of a critical region consists of the voxels from which a probabilistic tractography reaches the region. A face of two regions consists of the voxels that connect them.

together, they form the topological skeleton of a flow field. In diffusion MRI topology, the corresponding notion is the p -face of a pair of critical regions, consisting of the set of points which connect both regions with probability $P \geq p$. The derivation of this probability is left to Section 5.5. To illustrate the notion beforehand, Figure 4 (c) presents the counterpart of the basin in (b) on the right side of the corpus callosum. Figure 4 (d) shows the common 0.4- and 0.25-faces of the two critical regions, which clearly depict the part of the corpus callosum that connects both sides.

5 EXTRACTION OF TOPOLOGICAL FEATURES

Figure 5 gives an overview of the processing pipeline that will be described in this section. It comprises a preprocessing step in which the fiber tracking is performed (Section 5.1), a clustering step which forms the critical regions (Sections 5.2 and 5.3), as well as algorithms for the extraction and ranking of faces for examination by the user (Sections 5.4 and 5.5). We expect that a topological visualization of diffusion MRI data will be of specific interest to researchers in neu-

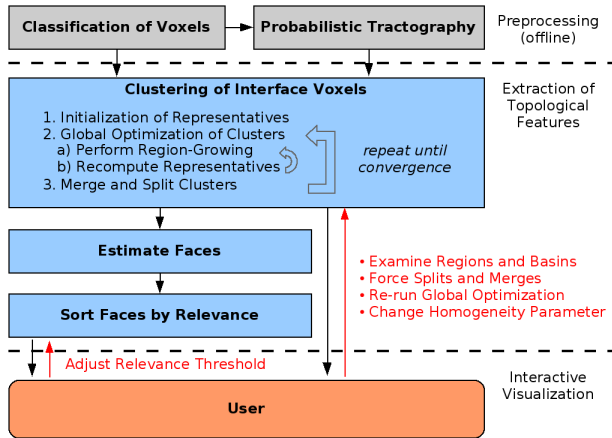


Fig. 5. An overview of the processing pipeline. The user can interact with it in a number of ways to test specific hypotheses.

rosience, so the proposed method aims to provide a sensible initial visualization, then allows the user interaction for formation and testing of specific hypotheses.

5.1 Preprocessing

As a first step in finding our topological features, we perform a probabilistic fiber tracking, using a 3D variant of the algorithm proposed by Koch et al. [20]. First, it classifies voxels as white matter, grey matter, and cerebrospinal fluid (CSF). The volume outside the brain is masked previously during tensor estimation, based on low signal values. Voxels with a tensor trace $\text{tr}(\mathbf{D})/3 > 10^{-9} \text{m}^2/\text{s}$ are marked as cerebrospinal fluid and grey matter is distinguished from white matter based on an anisotropy threshold (white matter: $\text{FA} > 0.2$). Isolated white or grey matter voxels are caused by fluctuations around the FA threshold, and removed in a post-processing step. Consequently, we call non-white matter voxels adjacent to white matter *interface voxels* and we add dummy voxels around the region of interest when white matter reaches the boundary of the domain.

The tractography itself is based on a random walk of particles at voxel resolution. Let \mathbf{r}_{mn} be the unit vector pointing from voxel m to a voxel n in its 26-neighborhood \mathcal{N} and $d_m(\mathbf{r}_{mn}) = \mathbf{r}_{mn}^T \mathbf{D}_m \mathbf{r}_{mn}$ be the apparent diffusivity in that direction, derived from the diffusion tensor \mathbf{D}_m . Then, Koch et al. define the transition probability $p(m \rightarrow n)$ from voxel m to n as

$$p(m \rightarrow n) = \frac{[d_m(\mathbf{r}_{mn}) + d_n(\mathbf{r}_{mn})]^a}{\sum_{n' \in \mathcal{N}} [d_m(\mathbf{r}_{mn'}) + d_n(\mathbf{r}_{mn'})]^a} \quad (2)$$

where the exponent a is empirically fixed at $a = 7$. Taking the exponent focuses the diffusivity profile to its major direction, which is likely to align with an actual fiber direction, while allowing for a certain surrounding spread. Some authors have used the product of diffusivities instead of the sum to adapt this method. In this modified form, Equation (2) has produced plausible cortex parcellations [1] and results that agreed with findings from fMRI [12].

After the first step, Koch et al. restrict the probability distribution to directions that deviate less than 90° from the previous step. We make two small improvements to this: First, we additionally set the transition probabilities to CSF voxels to zero, because it is anatomically impossible that fiber tracts end in the CSF-filled ventricles. Second, we do not simply truncate the distribution at 90° , but rather weight the probabilities in forward direction with $\cos \phi$, where ϕ is the angle between \mathbf{r}_{mn} and the current tracking direction \mathbf{t} , calculated from the direction \mathbf{r}'_{mn} of the previous step as $\mathbf{t} = \mathbf{D}_m \mathbf{r}'_{mn}$. This definition of \mathbf{t} accounts for the fact that the fiber direction changes from voxel to voxel and is analog to the “outgoing” direction in the tensorline propagation by Weinstein et al. [29]. Section 6.2 presents an example where these modifications are necessary to obtain correct results.

The random walk is terminated when the particle reaches an interface voxel. For each white matter voxel, we trace 10 000 particles and record the percentage that goes to the individual interface voxels. Similar to previous methods that pre-compute a deterministic tractography [7], this step is performed offline. For the region of interest in Figure 6 (4 948 white matter voxels), it takes more than six minutes on a 2 GHz Athlon 64 processor. The computations required by our modifications to the original algorithm account for 25% of the total time.

5.2 Clustering Criteria for Critical Regions

Cortex parcellation studies have computed and clustered a correlation matrix for the interface voxels in the region of interest, either manually [17] or with k-means [1]. However, forming critical regions within a topological visualization method requires that the number of clusters is chosen automatically, based on the data. Moreover, we cannot ensure connectivity of the critical regions when considering only the correlation matrix, since it does not contain any information about voxel adjacency. Consequently, a novel approach is required for the clustering of critical regions. This subsection introduces some notation and formalizes suitable clustering criteria. A custom algorithm which fulfills these requirements will then be presented in the following subsection.

Let \mathcal{W} be the set of white matter voxels w , $W = |\mathcal{W}|$. Similarly, \mathcal{I} is the set of interface voxels i , $I = |\mathcal{I}|$. Then, the tractography result for voxel w can be written as a vector $\mathbf{t}(w)$ of dimension I , where $t_i(w)$ is the percentage of particles originating from w that reached i . From this, we define the footprint $\mathbf{f}(i)$ of an interface voxel as a vector of dimension W :

$$f_w(i) = \text{FA}(\mathbf{D}_w) \cdot t_i(w) \quad (3)$$

Weighting particles with the fractional anisotropy at the originating voxel w has not been done by previous authors and is not strictly necessary to get sensible results. However, it helps to stabilize the clustering in the presence of noise (cf. Section 6.1), where the principal direction in regions of low FA may be unreliable.

A clustering Γ of the interface voxels is a partition of \mathcal{I} into C clusters $\Gamma_1, \dots, \Gamma_C$, where we require that each Γ_c is connected. The number of clusters C is not known a priori and changes as part of the clustering process. For each cluster c , the footprint $\mathbf{F}(c)$ is defined as the accumulated footprint of its members:

$$\mathbf{F}(c) = \sum_{i \in \Gamma_c} \mathbf{f}(i) \quad (4)$$

The similarity $\psi_c(i)$ between a cluster c and an interface voxel i is defined as

$$\psi_c(i) = \frac{\mathbf{f}(i) \cdot \mathbf{F}(c)}{\|\mathbf{f}(i)\| \cdot \|\mathbf{F}(c)\|} \quad (5)$$

Since none of the involved vectors have any negative components, the range of $\psi_c(i)$ is $[0, 1]$. From this, the homogeneity Ψ_c of a cluster c is defined as

$$\Psi_c = \frac{\sum_{i \in \Gamma_c} \|\mathbf{f}(i)\| \psi_c(i)}{\sum_{i \in \Gamma_c} \|\mathbf{f}(i)\|} \quad (6)$$

Since the total number of particles that arrive at an interface voxel can vary significantly with the number of white matter voxels in its neighborhood, it is appropriate to normalize ψ_c by the product of footprint magnitudes in Equation (5). In contrast, the weighting in Equation (6) reflects the fact that interface voxels with only a small number of particles should contribute less to the overall homogeneity of a cluster.

Let γ be a function that maps each interface voxel i to its cluster c (i.e., $\gamma(i) = c$ if $i \in \Gamma_c$). Then, a clustering is appropriate with respect to the data if the total homogeneity Ψ is high:

$$\Psi = \frac{\sum_{i \in \mathcal{I}} \|\mathbf{f}(i)\| \psi_{\gamma(i)}(i)}{\sum_{i \in \mathcal{I}} \|\mathbf{f}(i)\|} \quad (7)$$

If we leave the problem unconstrained, Ψ reaches its optimum at the trivial clustering, in which each interface voxel has its own cluster ($C = I$). Thus, we are interested in a clustering that is optimal under the additional condition that the homogeneity of each individual cluster c should approximately equal a parameter h ($\Psi_c \approx h$).

In our experiments, values around $h \approx 0.2$ generally gave useful results. However, part of the insight in [1] has been gained by trying various values of k for the k-means clustering, so leaving h as a user-defined parameter is useful for allowing an interactive exploration of the data. Also, the authors of [1] try to discover whether the data supports further subdivision of specific clusters, so we allow interactive splitting and merging of user-selected clusters. A subsequent global optimization of Ψ indicates if a split resulted in valid sub-clusters: In that case, surrounding clusters should not change significantly.

5.3 Clustering Algorithm

To find a clustering according to the criteria of the previous section, our method proceeds in two steps: The first step follows a greedy local strategy to create an initial clustering Γ . The second step globally optimizes both the clustering and the number C of clusters with respect to Ψ , preserving the conditions of connectivity and cluster homogeneity $\Psi_c \approx h$. Similar two-step methods have previously been used in computer vision to reduce the complexity of segmenting images into an unknown number of regions (e.g., [8]).

A common building block of both steps is a variant of the k-means algorithm that uses a fast-marching region-growing scheme to ensure connectivity of the resulting clusters. Like k-means, it iteratively computes new cluster footprints \mathbf{F}^{n+1} from a given clustering Γ^n and subsequently uses them to re-assign all interface voxels to new clusters Γ^{n+1} . Convergence is assumed when only a small percentage (e.g., 2%) of the voxels is re-assigned to a different cluster.

The footprints \mathbf{F}^{n+1} are determined by evaluating Equation (4). Consequently, for each cluster c , the voxel $i \in \Gamma_c^n$ with the highest similarity $\psi_c^{n+1}(i)$ is selected as a seed point. Starting from these seeds, voxels which have not yet been assigned to Γ^{n+1} are added to an adjacent cluster c . In order to optimize Ψ , voxels are added in descending order of their similarity $\psi_c^{n+1}(i)$. Thus, good-matching voxels are assigned early on, while dissimilar voxels are initially left free, which gives more suitable clusters the chance to become adjacent to them. This scheme is efficiently implemented using a priority queue.

The initial clustering ignores interface voxels i with $\|\mathbf{f}(i)\| < 0.2$. Because of their low weight in Equations (6) and (7), their influence on the final result is small. However, many of the final clusters are separated by regions of small footprint magnitude, so a connected component analysis of the interface voxels that fulfill this condition is an extremely simple and cheap way to identify some of the relevant clusters. The final global optimization, however, takes into account all voxels with $\|\mathbf{f}(i)\| > 0$. Experiments have indicated that using this heuristic does not affect the final result significantly, but nearly doubles the speed of the clustering process.

To make the algorithm more stable, we replace the parameter h with two parameters, h^+ and h^- , where h^+ is slightly larger than h^- . If the average similarity Ψ_c of a cluster is smaller than h^- , the cluster is split, to allow a more precise adaptation to the data. On the other hand, if merging two adjacent clusters would lead to a cluster homogeneity which is still larger than h^+ , the merge is performed.

Initially, each connected component is treated as a cluster and subdivided until h^- is reached. At this stage, the region-growing only acts locally on the voxels of the two newly formed sub-clusters. When a cluster is split, one half of its members are assigned arbitrarily to each of the two new clusters. After the first iteration of the region-growing algorithm, the results are again connected and converge to an optimum. In rare cases, this ‘‘careless’’ initialization causes very small sub-clusters to split off. However, this is acceptable, since such clusters will be re-merged later.


When the initial clusters have been found, the region-growing is used to extend them to all interface voxels and to refine them until global convergence. After merging and splitting clusters as appropriate, this procedure is iterated until no more merges or splits are necessary. Since the initial clustering is usually quite good, convergence is quickly reached.

In our implementation, we exploit the fact that most interface voxels only connect to a small fraction of the white matter, i.e., the footprint vectors \mathbf{f} are sparse. Thus, we store them as lists of <voxel index, value> pairs rather than as full-length arrays, which significantly reduces the cost of evaluating Equations (4) and (5). On the region of interest shown in Section 4 ($I = 3408$, $W = 4948$), the full clustering took 1.3 s on a 2 GHz Athlon 64 processor. Afterwards, small modifications to h , or user-specified splits and merges, followed by a global optimization, take around half a second, making these operations appropriate for interactive exploration of the data.

5.4 Definition of Faces

According to the definition in Section 4.2, we must determine the probability that a given voxel connects any two critical regions to find the faces in diffusion MRI topology. This information can be collected in the tractography step by using particle pairs that leave the starting voxel in opposite directions. Pairs of interface voxels that are reached this way are connected through the starting voxel.

Even though the target space of such pairs is of order I^2 , only a few pairs are actually connected, so for small enough regions of interest, a sparse representation makes this approach feasible. For example, in the region discussed above, probabilistic tractography from a single



(a) Exact faces from tracking particle pairs (b) Faces from a simple heuristic

Fig. 6. A visual comparison suggests that our simple heuristic provides a suitable approximation of the exact faces.

white matter voxel reaches 285 individual interface voxels on average, but only 1 135 voxel pairs. Still, this exact solution may become prohibitively expensive on larger regions of interest. Already in the case of Figure 9, the probabilities of more than $4 \cdot 10^7$ voxel pairs have to be stored. Thus, we additionally present a simple heuristic that estimates the face probabilities from the cluster footprints \mathbf{F} alone. Its fundamental idea is to divide the particles reaching a given region according to the ratio of particles that went to the remaining regions and to let them vote for a connection to these regions.

Let $T_c(w)$ be the percentage of particles from a white-matter voxel w that reach cluster c . If a single region collects more than 50% of the particles ($T_c(w) > 0.5$), we assume that the voxel w connects that region to itself with probability $P_{cc}(w) = 2 \cdot T_c(w) - 1$. To compute the connectivity between different regions, we remove these probabilities from T_c : Let $T'_c(w) = T_c(w) - P_{cc}(w)$ be the reduced percentages for which $T'_c(w) \leq 0.5$ and let $\bar{P}(w) = \sum_c P_{cc}(w)$ be the probability that w connects any critical region to itself. Then, the estimated probability $P_{ab}(w)$ that w connects clusters a and b , $a \neq b$, is given by

$$P_{ab}(w) = T'_a(w) \cdot \frac{T'_b(w)}{1 - T'_a(w) - \bar{P}(w)} + T'_b(w) \cdot \frac{T'_a(w)}{1 - T'_b(w) - \bar{P}(w)} \quad (8)$$

Equation (8) is an ad hoc definition, designed to satisfy the requirements that the resulting probabilities are non-negative and form a partition of unity. Its symmetry $P_{ab}(w) = P_{ba}(w)$ reflects the fact that we cannot distinguish between sources and sinks.


Figure 6 provides a comparison of faces obtained from tracking particle pairs (in (a)) and faces computed with our heuristic (in (b)). Visually, the approximation appears adequate. To allow an objective comparison, we averaged the absolute difference between exact and approximated $P_{ab}(w)$ over \mathcal{W} for the nine displayed faces. The absolute deviation ranged between 0.0007 and 0.0095, with an overall average of 0.0057. Over the voxels relevant for display ($P_{ab}(w) > 0.33$, corresponding to the confidence bounds in Figure 6), the average relative deviation of the heuristic from the exact method was between 0.5% and 10%, the overall average being 5%.

5.5 Selection of Relevant Faces

Similar to 3D flow topology, diffusion MRI topology is sensitive to the fact that three-dimensional faces may occlude each other, leading to visual clutter that is difficult to interpret. To alleviate this problem, we define a metric of face relevance, helping the user to select only the important faces for display.

A face is relevant if the voxels it contains belong to it with a high probability. Thus, a suitable formalization of the relevance ρ_{ab} of a face between clusters a and b is its summed probability P_{ab} over \mathcal{W} , normalized by the magnitude of the joint footprint:

$$\rho_{ab} = \frac{\sum_{w \in \mathcal{W}} P_{ab}(w)}{\|\mathbf{F}_a + \mathbf{F}_b\|} \quad (9)$$



(a) Clustered streamlines in the same data as Figure 6 (b) Faces in less-ideal data of the same subject

Fig. 7. A comparison with clustered streamlines and the faces in a noisier dataset validate the results and the reliability of our method.

Equation (9) is evaluated for all possible pairs of clusters and the resulting faces are ordered according to their value of ρ . The user can then add faces until she feels that cluttering occurs or less important faces start to appear. In Figure 6, the nine most relevant p -faces have been selected this way and rendered with $p = 0.66$ (confidence bounds at $p = 0.33$). Computing and ordering the faces took 0.27s for exact faces, and around 10 ms with the heuristic.

Figure 7 (a) gives a visual comparison of our results to a deterministic tractography, pseudo-colored according to a clustering with the method by Brun et al. [9], which is based on normalized cuts and was also used by Enders et al. [11]. Both clusterings agree on the main features: The *corpus callosum* (CC) is subdivided into multiple regions and separated from the *cingulum bundles* (Cing). Only a few streamlines were seeded inside the *internal capsule* (IC, truncated by the region of interest), making it more recognizable in Figure 6.

Note that in Figure 7 (a), some streamlines of the *corpus callosum* are clustered as part of the *cingulum bundle* and vice versa. This is due to the fact that Brun et al. project the fibers to a low-dimensional feature space in which these streamlines are not well-separated. In contrast, our clustering works directly on the high-dimensional voxel footprints. Also, clustering the 2290 displayed streamlines using an unmodified NCut algorithm took 90s, which made fine-tuning of the parameters more time intensive than with our method.

6 EXPERIMENTAL RESULTS

6.1 Robustness under Noise

We tested the robustness of our features under both real and synthetic measurement noise. Figure 7 (b) presents results on a second dataset from the same subject as in Figure 6. It also uses the setup described in Section 2.1, but includes only one, rather than three, measurements per direction. This reduces the measurement time for a full-brain scan to 15 minutes, at the cost of stronger physical noise.

For direct comparison with the results in Section 2.2, Figure 8 additionally shows faces in the datasets which were corrupted with additive Gaussian noise. Since the noise causes higher variability in the voxel footprints, we have selected a lower homogeneity target h than in Figure 6 to obtain a comparable number of critical regions ($h = 0.24$ without noise, $h = 0.22$ with artificial noise, $h = 0.21$ with physical noise). In all examples, we show the 0.66- and 0.33-faces with a relevance value $\rho > 5$.

The results suggest that even though the exact clustering changes, most notably in the slightly different subdivision of the corpus callosum, all major structures remain well recognizable at all noise levels. Timings were around 2s in all cases.

6.2 Results on a Larger Region of Interest

For additional validation, Figure 9 presents results on a second, larger region of interest, which spans $49 \times 39 \times 25$ voxels in the center of

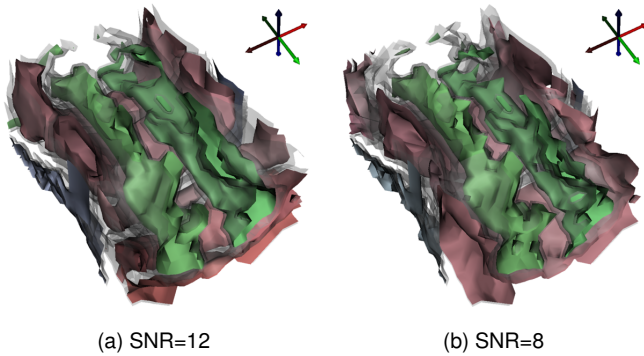


Fig. 8. The major faces ($\rho > 5$) remain recognizable when adding synthetic noise.




Fig. 9. The most important structures are found also in this larger region. However, the probabilistic tractography makes it difficult to trace thin fibers over long distances.

the brain, right below the *corpus callosum*. It is shown from a posterior/superior point of view, with additional data context provided by FA slices. As can be seen in the streamline tractography in Figure 10, this region contains more white matter than could be shown occlusion-free in a single rendering. Thus, we sorted the 1 176 possible faces using the ρ -criterion and manually selected twelve out of the 23 highest-ranking faces for display. They depict the *inferior fronto-occipital fasciculus* (IF), the *internal capsule* (IC), the *anterior thalamic radiation* (ATR) and parts of the *corpus callosum* (CC, truncated by region of interest), which also can be partly recognized in the streamline clustering in Figure 10 (b).

In this experiment, it becomes apparent that the probabilistic tractography method used for pre-processing is not well-suited for tracing thin fibers over long distances. At some point, most random paths end in the walls of such structures rather than traversing them fully. The notable deviation between the core ($p = 0.4$) and the confidence bounds ($p = 0.2$) of IF and ATR in the rendering reflects the high amount of uncertainty that results from this.

Even though the modifications suggested in Section 5.1 improve on this problem to some degree (with the original formulation by Koch

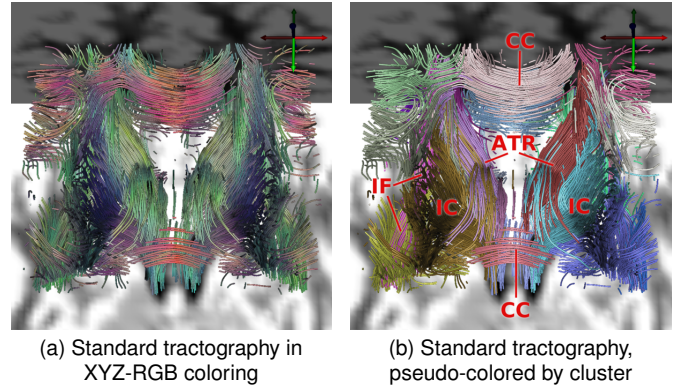


Fig. 10. Results from a deterministic tractography and a streamline clustering for comparison.

et al. [20], we were not able to reproduce the IF and ATR at all), we consider this a major obstacle in analyzing larger regions of interest using our method. We find it likely that a reliable estimate of long-range connectivity will require the introduction of prior knowledge into the tracking process. However, such an undertaking is outside the scope of this paper.

Finally, the following table summarizes the timings from our experiments, which indicate that our clustering method remains feasible for larger regions of interest.

	I	W	time	time/ I	WM/ I
Figure 6	3 408	4 948	1.3 s	0.4 ms	413
1/2 of Figure 9	7 803	12 765	10.8 s	1.4 ms	771
Figure 9	14 633	24 787	22.2 s	1.5 ms	852

For comparison, we have included a region of interest that covers only the left half of Figure 9. When doubling the input size from the second to the third row, the time spent per interface voxel remains almost constant. However, the clustering for Figure 6 is much cheaper. This is partially explained by the sparse representation of voxel footprints, which exploits the fact that the number of white matter voxels connected to an average interface voxels (WM/ I) is significantly lower.

7 CONCLUSION AND FUTURE RESEARCH

The motivation for this work was a lack of methods for connectivity-aware feature extraction from diffusion MRI data that would resemble the expressiveness of topological methods for flow visualization. We have closed this gap by defining suitable, anatomically meaningful topological features in brain diffusion MRI and proposing a method for their extraction.

At the same time, we have contributed a method for visual analysis of results from probabilistic tractography. In the past, visualization research focused on deterministic streamline techniques. However, researchers interested in quantitative connectivity studies have deemed confidence intervals provided by probabilistic methods indispensable for their work and are lacking appropriate methods for visualizing their results. In recent papers, slice projections [5] or volume renderings [1] of scalar connectivity values derived from the tractography constitute the state of the art.

While our work solves some open issues, it also leaves a number of questions to future research. To reduce the complexity of our approach, we have neglected the uncertainty in the critical regions themselves. It could be worthwhile to investigate if a probabilistic clustering further improves visualization. Also, more work could be done on the rendering of features. Currently, we only assign uniform colors to the basins and faces. Textures could add information relevant for interpreting the probabilistic tractography, like local fiber orientations and their variance, or the local density of particles.

Out of the need to reflect the uncertainty inherent in our data, we have defined probabilistic versions of some basic topological features,

leading to expressive visualizations. While it is outside the scope of the present work, it would be interesting to derive a rigorous mathematical framework for fuzzy topological visualization that may rest on existing fuzzy set theory (e.g., cf. [16]), and to apply it to other cases in which uncertainty visualization may play a role.

Finally, our approach integrates two topics of active research in the neuroscience community, namely, finding probabilistic fiber tracking methods that reliably reproduce fiber tracts known from anatomy (e.g., cf. [5]), and clustering grey matter voxels in a way that reflects functional units (e.g., cf. [1]). Our work has both benefited from this research and leads to a method that could help neuroscientists to better explore their data. It is our hope that having these tasks as part of our visualization pipeline will continue to create synergies between the two exciting fields of visualization and neuroscience.

ACKNOWLEDGEMENTS

We would like to thank Lawrence Frank and Greg Balls for a kind invitation to the University of California at San Diego (UCSD) and for inspiring discussions on probabilistic fiber tracking.

We are grateful to Alfred Anwander, who is with the Max Planck Institute for Human Cognitive and Brain Science in Leipzig, Germany, for stimulating discussions and suggestions, and for providing the datasets used in this paper.

Our implementation uses the BioTensor framework [27], developed at the Scientific Computing and Imaging Institute at the University of Utah, and the team libraries by Gordon Kindlmann, <http://teem.sf.net/>. Natascha Sauber kindly provided us with code for streamline clustering.

The thoughtful and constructive comments given by the anonymous reviewers were of substantial help in improving this paper. Kaleigh Smith has helped by proofreading the final manuscript; any remaining errors are our own fault.

This research has partially been funded by the Max Planck Center for Visual Computing and Communication (MPC-VCC).

REFERENCES

- [1] A. Anwander, M. Tittgemeyer, D. von Cramon, A. Friederici, and T. Knösche. Connectivity-based parcellation of broca's area. *Cerebral Cortex*, 17(4):816–825, 2007.
- [2] P. J. Basser, J. Mattiello, and D. L. Bihan. Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance*, B(103):247–254, 1994.
- [3] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
- [4] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, B(111):209–219, 1996.
- [5] T. Behrens, H. Johansen-Berg, S. Jbabdi, M. Rushworth, and M. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34:144–155, 2007.
- [6] T. Behrens, M. Woolrich, M. Jenkinson, H. Johansen-Berg, R. Nunes, S. Clare, P. Matthews, J. Brady, and S. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50:1077–1088, 2003.
- [7] J. Blaas, C. P. Botha, B. Peters, F. M. Vos, and F. H. Post. Fast and reproducible fiber bundle selection in DTI visualization. In C. Silva, E. Gröller, and H. Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 59–64, October 2005.
- [8] T. Brox and J. Weickert. Level set segmentation with multiple regions. *IEEE Transactions on Image Processing*, 15(10):3213–3218, October 2006.
- [9] A. Brun, H. Knutsson, H. J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber tracts using normalized cuts. In C. Barillot, D. Haynor, and P. Hellier, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI'04)*, volume 3216 of *Lecture Notes in Computer Science*, pages 368–375. Springer, 2004.
- [10] T. Delmarcelle and L. Hesselink. The topology of symmetric, second-order tensor fields. In R. Bergeron and A. Kaufman, editors, *Proceedings of IEEE Visualization 1994*, pages 140–147, 1994.
- [11] F. Enders, N. Sauber, D. Merhof, P. Hastreiter, C. Nimsky, and M. Stamminger. Visualization of white matter tracts with wrapped streamlines. In C. Silva, E. Gröller, and H. Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 51–58, 2005.
- [12] A. D. Friederici, J. Bahlmann, S. Heim, R. I. Schubotz, and A. Anwander. The brain differentiates human and non-human grammars: Functional localization and structural connectivity. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 103(7):2458–2463, February 2006.
- [13] H. Gudbjartsson and S. Patz. The rician distribution of noisy MRI data. *Magnetic Resonance in Medicine*, 36(2):910–914, 1995.
- [14] J. Helman and L. Hesselink. Representation and display of vector field topology in fluid flow data sets. *Computer*, 22(8):27–36, 1989.
- [15] L. Hesselink, Y. Levy, and Y. Lavin. The topology of symmetric, second-order 3D tensor fields. *IEEE Transactions on Visualization and Computer Graphics*, 3(1):1–11, 1997.
- [16] U. Höhle and E. P. Klement, editors. *Non-classical logics and their applications to fuzzy subsets*, volume 32 of *Theory and decision library. B : mathematical and statistical methods*. Kluwer, 1995.
- [17] H. Johansen-Berg, T. Behrens, M. Robson, I. Drobnjak, M. Rushworth, J. Brady, S. Smith, D. Higham, and P. Matthews. Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 101(36):13335–13340, 2004.
- [18] L. Jonasson, X. Bresson, P. Hagmann, O. Cuisenaire, R. Meuli, and J.-P. Thiran. White matter fiber tract segmentation in DT-MRI using geometric flows. *Medical Image Analysis*, 9:223–236, 2005.
- [19] G. Kindlmann, X. Tricoche, and C.-F. Westin. Anisotropy creases delineate white matter structure in diffusion tensor MRI. In R. Larsen, M. Nielsen, and J. Sparring, editors, *Medical Image Computing and Computer-Assisted Intervention (MICCAI'06)*, volume 4190 of *Lecture Notes in Computer Science*, pages 126–133, Copenhagen, Denmark, October 2006. Springer.
- [20] M. A. Koch, D. G. Norris, and M. Hund-Georgiadis. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*, 16:241–250, 2002.
- [21] D. Le Bihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet. MR imaging of intravoxel incoherent motions: Application to diffusion and perfusion in neurologic disorders. *Radiology*, 161(2):401–407, 1986.
- [22] S. Mori and P. C. van Zijl. Fiber tracking: principles and strategies – a technical review. *NMR in Biomedicine*, 15:468–480, 2002.
- [23] S. Pajevic, A. Aldroubi, and P. J. Basser. A continuous tensor field approximation of discrete DT-MRI data for extracting microstructural and architectural features of tissue. *Journal of Magnetic Resonance*, 154:85–100, 2002.
- [24] T. Salzbrunn and G. Scheuermann. Streamline predicates. *IEEE Transactions on Visualization and Computer Graphics*, 12(6):1601–1612, 2006.
- [25] G. Scheuermann, B. Hamann, K. I. Joy, and W. Kollmann. Visualizing local vector field topology. *Journal of Electronic Imaging*, 9(4):356–367, October 2000.
- [26] G. Scheuermann and X. Tricoche. Topological methods in flow visualization. In C. Johnson and C. Hansen, editors, *The Visualization Handbook*, pages 341–356. Academic Press, 2004.
- [27] Scientific Computing and Imaging Institute (SCI). BioTensor: A scirun power app for processing and visualizing diffusion tensor images, 2006.
- [28] E. Stejskal and J. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *Journal of Chemical Physics*, 42:288–292, 1965.
- [29] D. Weinstein, G. Kindlmann, and E. Lundberg. Tensorlines: advection-diffusion based propagation through diffusion tensor fields. In *Proceedings of the Conference on Visualization '99*, pages 249–253, Los Alamitos, CA, USA, 1999. IEEE Computer Society Press.
- [30] X. Zheng, B. Parlett, and A. Pang. Topological lines in 3D tensor fields and discriminant hessian factorization. *IEEE Transactions on Visualization and Computer Graphics*, 11(4):395–407, July 2005.
- [31] X. Zheng, B. Parlett, and A. Pang. Topological structures of 3D tensor fields. In C. Silva, E. Gröller, and H. Rushmeier, editors, *Proc. IEEE Visualization 2005*, pages 551–558, 2005.
- [32] X. Zheng, X. Tricoche, and A. Pang. Degenerate 3D tensors. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, pages 241–256. Springer, 2006.