

Fast and Reproducible Fiber Bundle Selection in DTI Visualization

Jorik Blaas¹, Charl P. Botha¹, Bart Peters², Frans M. Vos^{3,4} and Frits H. Post¹

¹ Data Visualization Group, Delft University of Technology

² Psychiatric Centre, Academic Medical Centre, Amsterdam

³ Quantitative Imaging Group, Delft University of Technology

⁴ Dept. of Radiology, Academic Medical Centre, Amsterdam
The Netherlands

ABSTRACT

Diffusion Tensor Imaging (DTI) is an MRI-based technique for quantifying water diffusion in living tissue. In the white matter of the brain, water diffuses more rapidly along the neuronal axons than in the perpendicular direction. By exploiting this phenomenon, DTI can be used to determine trajectories of fiber bundles, or neuronal connections between regions, in the brain. The resulting bundles can be visualized.

However, the resulting visualizations can be complex and difficult to interpret. An effective approach is to pre-determine trajectories from a large number of positions throughout the white matter (full brain fiber tracking) and to offer facilities to aid the user in selecting fiber bundles of interest.

Two factors are crucial for the use and acceptance of this technique in clinical studies: Firstly, the selection of the bundles by brain experts should be interactive, supported by real-time visualization of the trajectories registered with anatomical MRI scans. Secondly, the fiber selections should be reproducible, so that different experts will achieve the same results.

In this paper we present a practical technique for the interactive selection of fiber-bundles using multiple convex objects that is an order of magnitude faster than similar techniques published earlier. We also present the results of a clinical study with ten subjects that show that our selection approach is highly reproducible for fractional anisotropy (FA) calculated over the selected fiber bundles.

CR Categories: I.3.8 [Computing Methodologies]: Computer Graphics—Applications I.3.6 [Computing Methodologies]: Computer Graphics—Interaction Techniques

Keywords: diffusion tensor imaging, tractography, white matter

1 INTRODUCTION

In recent years, MR diffusion tensor imaging (DTI) has emerged as a new technique for the quantification of water diffusion in living tissue [10, 9]. Diffusion is commonly represented for every voxel by a second-order symmetric positive definite tensor. These tensors, derived from a sequence of diffusion-weighted scans, each sensitive to a different diffusion direction, describe the free motion of water molecules in tissue. This makes the technique very suitable for imaging fibrous tissue, as there are large differences in diffusion speed parallel with the fiber direction, and perpendicular to it. The anisotropy is reflected in the tensor data; the amount of anisotropy and the main directions of diffusion can be made explicit by eigenanalysis of each tensor.

DTI is now widely used as a research tool for studying the micro-structure of white matter tissue in the brain. Examples of applications are studies to differentiate the global nerve structures between brains malformed by disorders and healthy brains, quantification of the strength and integrity of nerve tracts (e.g. for stroke patients), and studies to determine the connectivity structure between different parts of the brain.

Based on the diffusion tensors, the directional structure of the underlying tissue can be determined and visualized by tracing paths that follow white matter tracts. This operation is usually called fiber tracking, or white matter tractography in brain applications [6, 2]. From a number of seed points placed in a region of interest, large numbers of fibers can be tracked. The fibers generated can be used for visualization, but the resulting images are complex and difficult to interpret. An effective approach is to pre-determine trajectories from a large number of positions throughout the white matter (full brain fiber tracking) and to offer facilities to aid the user in selecting fiber bundles of interest connecting specified regions in the brain, or passing through a specified region.

Two factors are crucial for the use and acceptance of such a technique in clinical studies: Firstly, the selection of the bundles by brain experts should be interactive, supported by real-time visualization of the trajectories, registered with anatomical MRI scans. Secondly, the fiber selections should be reproducible, so that different experts will achieve the same results.

Therefore, interactive techniques are needed that allow brain experts to select the anatomically correct fiber bundles. If a number of fibers can be collected that meet the same requirements, such as connecting the same regions or passing along a certain location, the selected set can be considered to correspond to a fiber bundle. If such fiber bundles can be extracted in a reproducible way by different expert operators, then they can be used for quantification of diffusion along specific bundles to compare their characteristics between different patients, or for one patient at different times.

In this paper we describe a practical technique for fast interactive selection of fiber bundles with multiple convex objects that are defined by half-spaces. The selection technique is about an order of magnitude faster than a similar selection technique presented earlier [1], and thus selection remains interactive for far larger numbers of fibers. In addition, we present a specific selection approach based on three selection boxes and then show the validity of this approach with a group of ten subjects and two operators. The result is that this approach yields a high inter-operator reproducibility for fractional anisotropy, or FA, calculated over the selected fiber bundles.

The rest of this paper is structured as follows: Section 2 contains a summary of existing work on the visualization of DTI data. In section 3 we explain how to perform real-time fiber selection with multiple convex selection objects. We present the reproducibility study in section 4. In section 5, we give a brief overview of our DTI visualization software tool, called DTII. We present a summary and our conclusions and indicate possible avenues for future research in section 6.

2 RELATED WORK

Many techniques have been developed for the visualization of diffusion tensor data. For example, each tensor can be visualized with various kinds of glyphs that reflect the properties of the tensors that they represent [10, 4]. Direct volume rendering techniques can be applied where the color and opacity transfer functions are determined by anisotropy metrics that are derived from the tensor [5].

White matter tractography [6, 2], or fiber tracking, attempts to calculate and visualize the probable trajectories of neural pathways in the white matter of the brain. These pathways are formed by bundles of neuronal axons. This is probably one of the most popular techniques for visualizing DTI data.

Until recently, most fiber tracking implementations have focused on allowing the user to select volumes of interest and then computing the fibers that pass through these regions. However, recent work has investigated the merits of pre-calculating a large number of fibers throughout the brain and analyzing the results.

For example, fibers can be automatically clustered and visualized [11, 3]. It has also been demonstrated that pre-calculating a large number of fibers throughout the brain and enabling flexible interaction with the fibers is an effective approach [1]. Fiber bundles can be interactively selected by manipulating a number of selection boxes or ellipsoids and combining the selected fibers with logical operations. This paper emphasized the importance of interactivity as an important component of this approach.

In this work, we show that such a selection technique leads to highly reproducible FA measurements over fiber bundles and can therefore be used as a measurement tool in pre-clinical studies. In addition, we improve on the interactivity aspects by presenting a technique whereby fiber selection with multiple convex selection objects can be performed extremely rapidly.

3 REAL-TIME FIBER SELECTION

As explained above, interactive selection of pre-calculated fiber bundles with multiple selection objects is an important component of effective visualization of DTI data. However, it is crucial that this selection can be done in real-time while the user interacts with the system.

In this section, we present a technique for the real-time selection of fiber bundles based on their intersection with an arbitrary number of convex objects. The final selection of fibers is the result of an arbitrary binary relation between the different selection objects. For example, one could request all fibers that intersect with objects A and B but not with C.

3.1 Overview

Fibers are represented as collections of short line segments. The simplest way to check which fibers intersect with a given selection object, is to check all line-segments of all fibers against the selection object. However, this would be quite slow, since there is a large number of line segments: an average dataset contains 45000 fibers with on average 60 line segments each resulting in a total of 2.7 million points.

Instead, we chose to implement this test as follows:

1. During pre-processing, all points that make up the fibers are inserted into a spatial data-structure.
2. All points inside the given selection object are enumerated.
3. All fibers that have at least one point inside the selection object are marked.




Figure 1: A two-dimensional kd-tree of depth 4 containing 32 points.

Instead of intersecting the polylines that represent the fibers with the selection object, we check simply for containment of the vertices that define these polylines. Since the length of a line-segment is small compared to the size of a selection object, this does not significantly affect the results. The spatial data-structure that we use is the kd-tree. This can be used to enumerate all points inside any convex polyhedron. Each point is linked to its fiber, so that visited points can quickly be mapped to visited fibers.

In the following subsections, we explain how the kd-tree is built and how it should be traversed, and we show how the technique can be applied to combinations of multiple selection objects. We conclude with some performance figures for the three-box selection method described previously.

3.2 The kd-Tree

The kd-tree is an n-dimensional tree of spatial subdivisions [7]. Each subdivision splits a rectangular sub-space into two halves, which preferably contain an equal number of points. The splitting planes are axis oriented, and their orientation cycles with each level of recursion. For a two dimensional tree, this means that on level 1 the tree splits on the x-axis, on level 2 on the y-axis, and on level 3 on the x-axis again.

In contrast to other spatial subdivision schemes like the octree, we can freely choose the coordinate of the splitting plane, which gives us the freedom to balance the tree. If we choose each splitting location exactly at the median of the points sorted along the splitting direction, then each split space contains the same number of points, thus producing a balanced tree structure.

The fiber selection uses a 3 dimensional version of this tree, but for illustration purposes we will examine the two dimensional case.

To illustrate how this splitting technique works, 32 two dimensional points and their corresponding kd-tree are shown in Figure 1. Each split is marked with its depth, such that level 0 and 2 split along the x-axis, and level 1 and 3 along the y-axis. As we can see, each leaf node contains 2 points, but since the distribution of the nodes varies, not all leaf-nodes have the same dimensions.

3.3 Traversal

Once the kd-tree has been constructed, we can use it to determine which fiber vertices are contained in an arbitrary convex polyhedron P . We describe the polyhedron P as the intersection of a set of half-spaces $H_1 \dots H_n$:

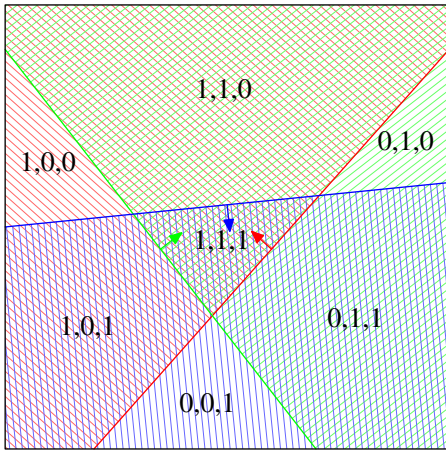


Figure 2: In two dimensions, the intersection of three half-planes yields a convex polygon. The three half-planes are hatched in red, green and blue respectively. Each area is marked with its bit vector, so that (111) forms the polygon. In three dimensions, the polygons become polyhedra, and the half-planes become half-spaces.

$$P = \bigcap_{i=1}^n H_i$$

To indicate in which half-spaces a given point x is contained, we use a binary vector v of length n .

$$v_i = \begin{cases} 1 & x \in H_i \\ 0 & x \notin H_i \end{cases} \quad i = 1 \dots n$$

These binary vectors form a way of classifying areas, as can be seen in Figure 2.

The general idea of our approach is to recursively traverse the kd-tree in order to determine which nodes of the kd-tree are contained in P , while visiting as few nodes as possible.

A pseudo-code version of the traversal procedure is shown in Algorithm 1. Refer to this during the following explanation.

We start by examining the root node of the tree, which contains all points. Each visited node is classified respectively to the given polyhedron P . A node can be either completely inside of P , completely outside of P , or partially inside (see Figure 3).

A node is classified into one of the three categories by examining all corners of the bounding box of that node. In our case, this three-dimensional box has 8 corners, denoted by $\mathbf{b}_1 \dots \mathbf{b}_8$. The collection of all eight corners is denoted B .

If all corners of the bounding box are inside of P then, as a result of the convexity of P , the entire box must be contained in the polyhedron. In this case, we can simply mark all points contained in this node as “inside”, and no further testing is necessary.

If this is not the case, then we try to establish whether the bounding box is completely outside of the polyhedron. We do this by checking whether there is any half-space in P for which all points are outside. If this is the case, then we can conclude that none of the points inside the bounding box can be inside P , and no further recursion is necessary.

When neither of the two previous cases are true, then it is possible that the points are partly inside and partly outside convex polyhedron P . This is where we exploit the kd-tree’s spatial subdivision by recursively examining both child-nodes. This recursion continues until all points contained in the kd-tree have been classified.

Algorithm 1 The recursive kd-tree search algorithm.

```

function recurse
if leafnode then
  visit()
else
  allInside  $\leftarrow$  1
  for all  $h \in H$  do // iterate the  $h$  over all half-spaces
    allOutside  $\leftarrow$  1
    for all  $x \in B$  do // iterate the  $x$  over all corners of the box
      if  $x \in h$  then //  $x$  is inside the half-space  $h$ 
        allOutside  $\leftarrow$  0
      else //  $x$  is outside
        allInside  $\leftarrow$  0
      end if
    end for
  end for
  if allOutside = 1 then // all points were outside  $h$ 
    return // completely outside: visit no points
  end if
  if allInside = 1 then // all points were inside all half-spaces
    visit_all() // completely inside: visit all points
    return
  end if
  recurse(left)
  recurse(right)
end if

```

The main kd-tree search algorithm. The auxiliary function *visit* tests all points contained in the leaf node against the polyhedron, and flags all points that are inside. The *visit_all* function does the same but assumes all points are always inside.

In short, the procedure we described is recursive and starts by examining the root node. It examines the surrounding box for the node and efficiently decides how the corners of the box relate to the convex search area. This relation can be classified in three cases: fully inside, fully outside, neither. Only in the last case further recursion is necessary. Figure 4 shows a graphical representation of all the visited nodes of the kd-tree after the cingulum has been selected.

3.4 Marking the fibers

Now that we have found all the points, we still need to find all the corresponding fiber bundles. Since each point only belongs to a single fiber we store a simple fiber-reference with each point (see Figure 5). Using this pointer, we can quickly flag all fibers that contain visited points.

3.5 Multiple Convex Regions

We can expand this idea to work with multiple regions by changing the fiber flag to a bit-vector. For each convex region we reserve a single bit.

We initially set all bit-vectors to zero, and then for each region flag all fibers that have points inside the region. After processing all regions, each fiber will have a bit-vector corresponding to which regions it intersects. This can be easily used to select all fibers that correspond to any given logical combination of selection objects.

A useful selection approach is a logical AND of three boxes, so that only the fibers that cross all three boxes are visible. This allows us to select most standard fiber tracts.

Sometimes these selection methods erroneously mark extra fibers. To get rid of these fibers we introduce extra boxes that work in a negated fashion: any fibers that pass through any of these boxes are filtered out. This allows for easy manual pruning of a selection set.




Figure 3: The three classes of containment for a box with of 4 corners $b_1 \dots b_4$, with respect to a polyhedron defined by the intersection of four half-spaces $h_1 \dots h_4$. On the left, the box is completely contained. In the middle, all points are outside the half-space h_3 , so the box is completely outside the polyhedron. In the rightmost figure neither is true, and further recursion is necessary.




Figure 4: The visited nodes of the kd-tree for the selection of the cingulum.




Figure 5: A number of points in a leaf-node, each with a reference to a fiber mask entry.

3.6 Performance

The performance of any selection method of this type is dependent on the size of the dataset, the number of fibers in the data, the size of the selection objects and the current position of the selection objects. We measured the performance of the implemented techniques by making use of a dataset with 47500 fibers and we used three selection objects as explained in section 4. By moving the selection boxes to different parts of the brain, we measured a range of selection speeds.

In all cases, the system remained completely responsive, and the fiber selection was updated in real-time during intensive interaction with the selection objects. On a first generation Pentium-M at 1.6GHz, our implementation processed between 1 and 2 million fibers per second. On a Pentium 4 at 3GHz, it processed between 1.5 and 3.5 million fibers per second. This includes selection *and* rendering of the selected fiber bundles.

This is at least an order of magnitude faster than the figures published in [1]: on a “1.6GHz Pentium laptop PC” they were able to intersect a single VOI with between 80000 and 220000 pathways per second. We attribute this improvement to the fact that our selection and rendering implementations were purpose-designed for highly-interactive DTI, whereas general-purpose visualization and collision-detection libraries were used in [1].

4 VALIDATION OF THREE-BOX SELECTION METHOD

In a DTI-based study of healthy and schizophrenic subjects, a selection method was required whereby sections of fiber bundles could be easily and reproducibly selected. DTI-derived parameters, such as for example the FA, or fractional anisotropy, can then be calculated over these sections and be compared for different test subjects.

Based on the techniques presented in section 3, we chose an approach based on three bounding boxes. All three bounding boxes can be freely positioned, rotated and resized within the DTI user interface. One or more slice-views of the data can be used as context during this navigation. All fibers intersecting with any one of the three cubes are pre-selected. One of the three bounding boxes also functions as the section selector: only sections of the pre-selected fibers that are contained within this box are added to the final selection. Figure 6 shows such a fiber selection in progress. Unselected fibers are shown with a different shading to serve as context for the selected fibers. The box in the middle is the section selector in this case: only the fiber sections contained in this box are actually selected. These are marked with red dots. The selection is performed and shown in real-time throughout the user’s interaction with the system.

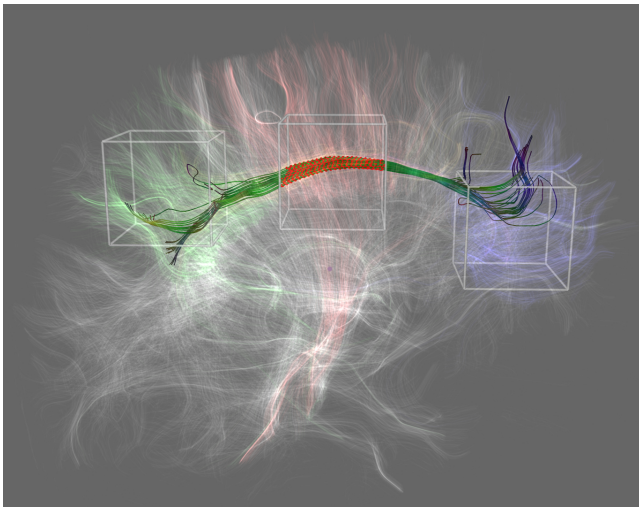


Figure 6: The three-box selection method in progress. A pre-selection is made of all fibers that intersect all three boxes. Out of this pre-selection, only fiber sections that are contained within the center box are selected for further processing.

Percentile	Left Cingulum		Right Cingulum	
	Obs. 1	Obs. 2	Obs. 1	Obs. 2
25	4359.25	4370.25	4969.00	5039.25
50	4801.00	4742.00	5505.50	5513.00
75	5196.75	5329.75	5689.50	5760.00
IQR	837.50	959.50	720.50	720.75

Table 1: 25, 50 and 75 percentiles and the interquartile ranges (IQR) for the FA measurements by each of the observers on the left and the right cingulum. The FA measurements have been scaled to the range [0, 10000].

DTI datasets were acquired of five healthy and five schizophrenic subjects on a Philips Intera 3 Tesla MRI scanner. The datasets were acquired along 16 different diffusion directions. Each direction was scanned twice and averaged to improve the signal to noise ratio. The scan time was approximately 6 minutes per patient and the final resolution of the DTI dataset was $2 \times 2 \times 2$ mm.

Two users applied the method described above to select the fiber bundles corresponding with the left and the right cingulum in all ten datasets. In other words, 40 independent selections were performed. On average, each selection took less than two minutes.

The average FA, or fractional anisotropy, was calculated over the selected fiber sections, resulting in four FA measurements per subject. The fiber tracking was performed anew by each user. We made use of FA due to the specific clinical questions that were being asked and due to its popularity in clinical DTI literature. In the future, we will continue our research by investigating other metrics, both existing and newly developed.

Due to the relatively small sample size, we chose to make use of non-parametric statistics to analyze our data.

Table 1 shows the 25th, 50th and 75th percentiles as well as the interquartile ranges (IQR) for the measurements by each of the observers on the left and on the right cingulum. The interquartile ranges show that the FA measurement varied significantly over different patients.

For the left cingulum, the non-parametric Spearman correlation between the two observers was 0.903 and for the right cingulum it was 0.976, both with a two-tailed significance of less than 0.001. This shows that the inter-observer correlation is sufficiently high

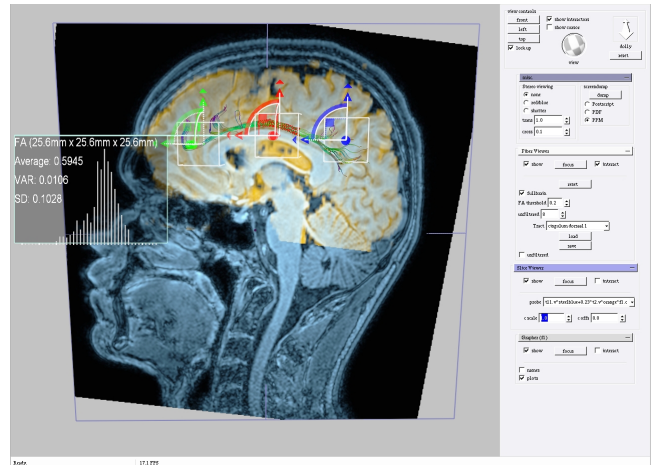


Figure 7: The main user interface of DTII. A T1-weighted and T2-weighted anatomical scan are combined using color fusion (blue/orange). This is used as a background context against which the fiber selection of the cingulum is performed using the three-box method.

for the reproducible selection of fiber bundles and subsequent measurement of FA over these selections.

5 THE DTII SYSTEM

The DTII platform was designed to provide functionality for the interactive exploration of DTI data. It was developed in cooperation with the Academic Medical Centre in Amsterdam. The software performs full brain tractography as a preprocessing step and allows the user to perform interactive selections on the pre-calculated fibers by using the techniques presented in section 3. In addition, other MRI data can be combined with the DTI visualization to provide context. It supports fusion of anatomical MRI, functional MRI (fMRI) and DTI data. This data is usually acquired in a single scanning sequence, but with proper registration multiple scans can also be combined. Figure 7 shows the the main user interface of the software. This tool was used in the reproducibility study documented in section 4.

DTII employs the the TEEM library for all data handling and fiber tracking. TEEM is an open-source image processing and visualization library that also supports tensor processing. Via this toolkit, DTII offers two types of fiber tracking: fourth-order integration based streamlines on the primary eigenvectors of the diffusion tensors or tensor lines [8].

All visualization methods in DTII support progressive updates in order to maintain high frame-rates even during complex manipulations. DTII was written in C++ and is highly portable, depending only on a few external libraries. Currently, versions for Windows and Linux are in use. It is complemented by a conversion utility which can convert both MRI and DTI DICOM series to the native DTII data format. This facilitates deployment in a medical environment.

The user interface is customizable so that only the components relevant to the task at hand are displayed. This greatly eases repetitive selection processes, as the user is no longer distracted by superfluous user interface elements.

Feedback from clinical users has been positive.

6 CONCLUSIONS AND FUTURE WORK

This paper presents two contributions to the field of DTI visualization. Firstly, we presented a technique that enables the real-time selection of fiber bundles with multiple convex selection objects. Sets of fibers that intersect with the convex selection objects can be combined with logical operations. The kd-tree datastructure combined with the halfspace representations of the selection objects is very efficient: selections with three objects in a full brain dataset with 2.7 million fiber points works very well in interactive mode. We showed that this technique performs an order of magnitude faster than similar previous techniques.

Secondly, we have tested the interactive selection of DTI-derived fibers with three manipulable boxes on 10 subjects. For each of these 10 subjects, two users independently selected the fiber bundles corresponding to the left and the right cingulum. The average fractional anisotropy (FA) was calculated over these fiber bundles. Our analysis of the test data shows both a significant variation of the FA over the different subjects and a significantly high correlation between the measurements by the two observers. This inter-observer reproducibility indicates that this type of approach can be used as a measurement tool in pre-clinical studies.

The interactive bundle selection may be combined with semi-automatic methods such as similarity clustering. We will continue developing the DTII tool and plan to continue working on its integration in a clinical setting. We will also perform more similar studies on test subjects to determine the reproducibility of other fiber selection based metrics. Ultimately, we wish to develop reliable and reproducible connectivity metrics that can be used to measure the “strength” of a neural connection between two regions of interest. In this way, diagnostic test sets can be developed to examine the main connection pathways in the brain to diagnose brain disorders. Visualization will play an important role in this type of diagnosis.

Our selection method is based on multiple convex polyhedra. At the moment there is no requirement for concave selection objects, but if this requirement arises, we plan to extend our approach by integrating methods for decomposing concave polyhedrons into multiple convex polyhedra.

ACKNOWLEDGEMENTS

DTII makes use of the `teem` toolkit, available at <http://teem.sourceforge.net>.

The involvement of Charl P. Botha and Frits H. Post in this work took place in the context of the Virtual Laboratory for e-Science project (<http://www.vl-e.nl/>). This project is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OC&W) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ).

REFERENCES

- [1] David Akers, Anthony Sherbondy, Rachel Mackenzie, Robert Dougherty, and Brian Wandell. Exploration of the Brain’s White Matter Pathways with Dynamic Queries. In Holly Rushmeier, Greg Turk, and Jarke J. van Wijk, editors, *Proceedings of IEEE Visualization 2004*, pages 377–384, 2004.
- [2] P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance Imaging*, 44(4):625–632, 2000.
- [3] Anders Brun, Hans Knutsson, Hae-Jeong Park, Martha E. Shenton, and Carl-Fredrik Westin. Clustering Fiber Traces Using Normalized Cuts. In C. Barillot, D.R. Haynor, and P. Hellier, editors, *MICCAI 2004*, volume LNCS 3216, pages 368–375. Springer-Verlag, 2004.

- [4] Gordon Kindlmann. Superquadric Tensor Glyphs. In *Joint EURO-GRAPHICS - IEEE TCVG Symposium on Visualization*, pages 147–154, 2004.
- [5] Gordon Kindlmann, David Weinstein, and David Hart. Strategies for Direct Volume Rendering of Diffusion Tensor Fields. *IEEE Transactions on Visualization and Computer Graphics*, (6(2)):124–138, 2000.
- [6] Susumu Mori, B. J. Crain, V. P. Chacko, and Peter C. M. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. In *Ann Neurol.*, pages 265–269, 1999.
- [7] Beng Chin Ooi. Spatial kd-Tree: A Data Structure for Geographic Database. In *Büro, Technik und Wissenschaft*, pages 247–258, 1987.
- [8] David Weinstein, Gordon Kindlmann, and Eric Lundberg. Tensor-lines: Advection-Diffusion based Propagation through Diffusion Tensor Fields. In *Visualization ’99, proceedings*, pages 249–530, 1999.
- [9] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualisation for diffusion tensor MRI. *Medical Image Analysis*, 6:93–108, 2002.
- [10] C.-F. Westin, S.E. Maier, B.Khidhir, P. Everett, F. A. Jolesz, and R. Kikinis. Image processing for diffusion tensor magnetic resonance imaging. In *Proceedings of Second Int. Conf. on Medical Image Computing and Computer-assisted Interventions (MICCAI’99)*, pages 441–452, 1999.
- [11] Song Zhang and David H. Laidlaw. DTI Fiber Clustering and Cross-subject Cluster Analysis. In *Proceedings of ISMRM*, Miami, FL, May 2005.