

Lawrence Berkeley National Laboratory

Scientific Data

Title

Interactive Volume Rendering of Diffusion Tensor Data

Permalink

<https://escholarship.org/uc/item/9gh7722v>

Journal

Lawrence Berkeley National Laboratory

Author

Hlawitschka, Mario

Publication Date

2009-07-01

Peer reviewed

Interactive Volume Rendering of Diffusion Tensor Data

Mario Hlawitschka¹, Gunther H. Weber², Alfred Anwander³, Owen T. Carmichael⁴, Bernd Hamann⁵, and Gerik Scheuermann¹

¹ University of Leipzig, Germany

{hlawitschka|scheuermann}@informatik.uni-leipzig.de

² Computational Research Division, Lawrence Berkeley National Laboratory,
Berkeley, CA 94720-8139, USA

ghweber@lbl.gov

³ Max Planck Institute for Human Cognitive and Brain Science, Leipzig, Germany
anwander@cbs.mpg.de

⁴ Department of Neurology, University of California, Davis, CA 95616, USA
ocarmichael@ucdavis.edu

⁵ Institute for Data Analysis and Visualization (IDAV) and
Department of Computer Science,
University of California at Davis, CA 95616, USA
hamann@cs.ucdavis.edu

1 Introduction

As 3D volumetric images of the human body become an increasingly crucial source of information for the diagnosis and treatment of a broad variety of medical conditions, advanced techniques that allow clinicians to efficiently and clearly visualize volumetric images become increasingly important. Interaction has proven to be a key concept in analysis of medical images because static images of 3D data are prone to artifacts and misunderstanding of depth. Furthermore, fading out clinically irrelevant aspects of the image while preserving contextual anatomical landmarks helps medical doctors to focus on important parts of the images without becoming disoriented. Our goal was to develop a tool that unifies interactive manipulation and context preserving visualization of medical images with a special focus on diffusion tensor imaging (DTI) data.

At each image voxel, DTI provides a 3×3 tensor whose entries represent the 3D statistical properties of water diffusion locally. Water motion that is preferential to specific spatial directions suggests structural organization of the underlying biological tissue; in particular, in the human brain, the naturally occurring diffusion of water in the axon portion of neurons is predominantly anisotropic along the longitudinal direction of the elongated, fiber-like axons

[MMM⁺02]. This property has made DTI an emerging source of information about the structural integrity of axons and axonal connectivity between brain regions, both of which are thought to be disrupted in a broad range of medical disorders including multiple sclerosis, cerebrovascular disease, and autism [Mos02, FCI⁺01, JLH⁺99, BGKM⁺04, BJB⁺03].

To date, the predominant visualization techniques for tensor images are based on pre-calculation of geometric primitives that represent tensor properties. Glyph rendering, for example, converts each tensor into an iconic shape that is elongated to match the water diffusion properties implied by the tensor entries, whereas line rendering reduces the tensor data to a discrete set of contours that trace paths along the presumed locations of axon bundles — that is, along directions of highly anisotropic water diffusion (see Delmarcelle et al. [DH92, BL92].) However, computation of these geometric primitives can be computationally expensive, renderings of them can be prone to visual clutter and they are difficult to interactively modify to accentuate varying aspects of the data in a session-specific way to meet the goals of specific users. To circumvent the cluttering problem, clustering of streamlines has been introduced, e.g., by Enders et al. [ESM⁺05], but in the same year, a user study by Moberts et al. [MVvW05] revealed that clustering neural pathways is a highly user-dependent topic and the additional computation time cannot be neglected. Others presented techniques to interactively select lines depending on different features, e.g., Blaas et al. [BBVP05], who try to maintain the context by using high quality shading of rendered lines, still these techniques used precomputed lines. Kondratieva et al. [KKW05] reduced the computation time by integrating particle lines on the GPU where a simple Euler approach of order one is used for particle integration. Even on current GPUs the computational power is limited and therefore only a limited number of lines can be rendered which restricts the number of features that can be visualized.

Volume rendering, in contrast, has the potential to present entire images without reducing them to static, discrete sets of geometric primitives. Recent volume rendering techniques utilizing commodity graphics hardware support interactive frame rates and provide a high degree of interactive manipulation. However, the question of how to convert a volume of tensor-valued voxels into color and opacity maps for volume rendering is not trivial. Previously, univariate or tri-variate summary measures of the anisotropy of water diffusion at each voxel were computed based on the tensor eigenvalues; the summary measures and direction of the major eigenvector were mapped to alpha channels and color respectively. Lit Tensors, introduced by Kindlmann [KW99], define another approach that employs shading to emphasize anisotropy. Empirical evidence suggests that physicians become disoriented by direction-encoding schemes that deviate from the traditional approach of assigning red, green, and blue to three cardinal directions (compare Pajevic et al. [PP99]). Therefore, we focus on improving other visualization parameters while keeping this well known color coding. This facilitates fast switching between pseudo-colored section planes and our volume rendering approach. Previous approaches to

volume rendering as presented by Vilanova et al. [VZKL06] focus on setting transparency depending on scalar values or modifying the shading using directional information as presented by Kindlmann et al. in [KW99]. Both approaches fail to fade out information that is currently not important to the user, thus, we employ directional parameters in addition to those presented previously to modify the opacity value of the transfer function.

Other approaches to the occlusion problem in volume rendering strive to change the location or size of features in the volume to make them visible to the user, see Correa et al. [CSC06] and Bruckner et al. [BG06]. Whereas these techniques are capable of effectively uncovering important aspects of the data and provide good illustrations of the findings, they may be inappropriate in a clinical setting because the location and size of features in brain images often relate directly to the presence or severity of medical disorders (compare Simon et al. [SDB⁺05]) and, therefore, should be preserved. In contrast, our method is inspired by Importance-Driven Volume Rendering by Viola et al. [VKG04] in that we improve visibility of important sections of the image while fading out irrelevant parts and preserving voxel location and size. Our approach differs, however, in that it does not require pre-segmented models and therefore can be applied to raw DTI and structural magnetic resonance (MR) data without requiring a pre-processing step.

Whereas previous visualization techniques for tensor data are based on manipulation of simple scalar value based alpha functions either using fractional anisotropy (FA) as single scalar value or Westin’s barycentric spherical, planar and linear anisotropy measures a_s, a_p, a_l [WPG⁺97], we introduce transfer functions that depend on more information of the tensor, especially its eigenvector directions. While local direction information does not have medical importance as, e.g., neural fibers are defined as integral curves in a certain volume, in many areas of the brain, neighboring fiber bundles have different directions but their anisotropy values only change marginally. Kindlmann et al. [KTW06] showed that FA ceases can be found after pre-processing the data set but these methods cannot be easily transferred to volume rendering. Nevertheless, these areas can be found easily by looking at the local change of direction.

2 Volume Rendering

Simple volume rendering of DTI data is based on mapping the normalized major eigenvector (x, y, z) to a color by setting $RGB = (|x|, |y|, |z|)$ and using FA or the barycentric system of linear, planar and spherical anisotropy to define 1D or 2D opacity transfer functions [EHK⁺06, VZKL06]. Kindlmann described another approach for volume rendering of diffusion tensor data. He defined transparency based on FA values and added color and shading depending on other tensor properties such as direction and shape. Being inspired by these basic approaches, our method emphasizes the directional features of the

data set by using the default color mapping as it is used in the simple approach and provides additional tools for selecting important features.

The major task for our application is reducing the occlusion problem. We do this by allowing the user to interactively manipulate the transfer function depending on the full tensor information. Previous methods represent the isosurfaces of anisotropy values and do not take into account that medical doctors are mostly interested in fiber bundles. Therefore, it should be possible to select distinct bundles and hide others that occlude important parts of the data set. While anisotropy measures and other tensor invariants have proven to provide a good estimate of white matter boundaries, they do not provide information concerning fiber bundle boundaries. However, in most cases, changes of the directional tensor information provide this information. Therefore, we propose to make use of this directional information in direct volume rendering.

2.1 The Occlusion Problem

Since visualization maps information into two spatial dimensions, occlusion is one of the major problems when handling volumetric data sets. Making surfaces more transparent is one of the basic methods, that helps provide additional information. The selection of good occlusion coefficients is a difficult problem and painting too many semitransparent surfaces leads to visual clutter produced by both mixture of color and additional geometry. We propose an interactive selection of regions of interest by considering local fiber direction. In addition to the transfer function, one or more direction vectors are used to select areas of interest—or areas that should not be shown at all.

Although the color value of the RGB color mapping scheme provides directional information, the map is not invertible, i.e., even neglecting orientation, directions cannot be recomputed from the color value. In addition to that, interpolation of eigenvalues and eigenvectors directions is not the same as interpolating tensor values and recomputing eigenvalues and eigenvectors afterwards [ZB02, Kin04]. Therefore we compute the eigenvalues and the required major eigenvector on the GPU using the non-iterative approach by Hasan et al. [HBPA01]. Different criteria are used to select transparency values including:

- anisotropy values such as fractional anisotropy (FA), relative anisotropy (RA), linear, planar and spherical anisotropy [WPG⁺97],
- local direction of major eigenvector, and
- a magic lens ⁶ that modifies the local transfer function [BSP⁺93].

We experimented with several functions for directional filtering and determined that the type of function itself is not critical. While hard drop-offs in

⁶ “Magic Lens” is a trademark of Xerox Cooperation

transparency should be avoided, it is possible to use linear step functions as well as smooth step functions, either the one provided by GLSL defined as

$$f(x) = \begin{cases} 0 & \text{when } x < 0 \\ 3x^2 - 2x^3 & \text{when } 0 \leq x \leq 1 \\ 1 & \text{when } x > 1 \end{cases} \quad (1)$$

or higher-order ones like

$$f(x) = \begin{cases} 0 & \text{when } x < 0 \\ (126 + (-420 + (540 + (-315 + 70x) * x) * x) * x) * x^5 & \text{when } 0 \leq x \leq 1 \\ 1 & \text{when } x > 1 \end{cases} \quad (2)$$

Filtering out single directions breaks the typical symmetry of DTI data sets that is provided by the standard color coding, which produces unfamiliar looking images and interferes with the common way of comparing both hemispheres of the brain. To maintain the same symmetric properties, as the color coding does, one can compare all directions that have the same color (which are up to eight vectors in general, four because of the independence of eigenvector orientation.) To provide more information about the local direction of fibers, the comparison has to be restricted to one direction on the hemisphere. This can be done by defining the vector in one hemisphere and mirroring it at a plane separating both hemispheres of the brain. Having computed the normalized direction vector v and the normalized eigenvector e in the hemisphere, the alpha value is modified by

$$\alpha \propto f(\|\langle e, v \rangle\|). \quad (3)$$

2.2 Preserving Context

It turns out that removing large parts of the image completely, i.e., making it completely transparent, may confuse the user. Therefore, our system allows one to select an opacity for “hidden” areas making them translucent, which preserves the context by fading out uninteresting parts of the data only a bit while still allowing the user to see more important parts that may lie behind the uninteresting parts.

2.3 Combined Rendering of MRI and DT-MRI Data

The previous section focused on methods used for rendering DTI data. While many neuroscientists are primarily interested in this type of data, scalar-valued MRI images provide a higher resolution and make it easier to navigate in the brain. In addition, Gyral structure of brain gray matter and thin, soft structures like blood vessels can be seen in MRI scans while they do not show up in most DTI data sets. Therefore, a combined visualization of DTI and

MRI data is useful, especially for surgical planning. We integrate visualization of conventional MRI and DTI data in our volume rendering approach in two different ways:

- We add it to a single-pass volume rendering and define more complex transfer functions based on both, local DTI and MRI data and
- implement a two-pass volume rendering approach to enforce importance-driven volume rendering.

As we aim to restrict the diffusion tensor color map to the RGB model and mapping, introduction of additional colors would be misleading. Most of the information present in MRI data can be seen in simple grayscale slices of the brain, which is the most common way of displaying them. Therefore, we use volume rendering of the MRI data only to highlight boundaries of objects by using a standard 2D grayscale transfer function for MRI data. Special material attributes and gradient-based shading as done in most volume rendering applications, see Bruckner et al. [BVG05], are applied to improve the 3D perception of the surfaces.




Fig. 1. Three components of the compositing step — from right to left: importance mask, directional-colored DTI data, final compositing with MRI data.




Fig. 2. Another view on Fig. 1.

2.4 Implementation

Our implementation is based on a standard GPU-based raycasting volume renderer with 2D transfer functions and the 3D texture based approach as described, e.g., by Engel et al. [EHK⁺06]. Gradient computation can be done either on the GPU or on the CPU, depending on the available graphics board or by the choice of the user. The transfer functions have been extended to support either 2D transfer functions based on gradient magnitude and value or 2D functions based on two scalar values, e.g., MRI signal and FA. To filter the data by direction, we need additional information for the graphics board. Modern graphics cards provide several ways for sending the data. For our approach, we use the light information present in OpenGL as it is global information for the entire data set, which can be easily manipulated between frames. Furthermore, it provides additional parameters, like the cut-off-angle, and exponent that can be used for parameterization of the smooth step function. In medical imaging, usually, no more than two light sources are used (e.g., one static light and one headlight); therefore, at least five light sources are unused that are more than enough for definition of transfer functions. While the 2D transfer function is based on a lookup-table, we evaluate the tensor transfer function analytically on the GPU as it is not computationally

extensive and can be manipulated without the need to exchange textures. The major three steps, computing the stencil, computing the tensor volume rendering, and computing the conventional MRI overlay, are shown in Figs 1 and 2.

The magic lens selection provides an interactive way to select the weighting of MRI and DTI data in the final image. All computation is done on the GPU. We use the lens to emphasize areas of interest in the DTI image while providing the context of the Gyri using MRI. In contrast to Wang et al. [WZMK05], we do not apply different transfer functions to one data set but we select different transfer functions on different modalities to build a combined image that provides all information needed by the user while the geometry is not distorted. Application of this is shown in Figs 7–9. By using these simplifications, the lens and the volume rendering of both data sets can be painted in a single rendering pass which provides interactive frame rates on a full screen application. The speed of the algorithm is independent of the size of the lens. Furthermore, we experienced no change in the frame rate for different sizes of DTI and MRI data sets.

3 Application and Results

We have applied our method to multiple data sets. Our main focus was to determine what users can learn from interactive volume rendering that cannot be seen using pseudo-colored slices or fiber tracking algorithms.

3.1 Quality Analysis for Fiber Tracking in the Gyri

Analysis of connectivity in the human brain is one of the major fields of research related to DTI where the main interest is in understanding how the brain works by analyzing the connectivity of different areas in the brain. This abstract type of research finds its application in neuro-surgery. In tumor surgery, the basic understanding of where neural fibers are situated and which of them are more important than others is one of the most important fields of research having major impact on quality and safety of brain surgeries. Most of the time, neural fibers are visualized by displaying pre-computed line structures that imply single lines with defined start and end points, which is misleading since lines are often started and stopped based on other parameters, for example, by FA. By using volume rendering for DTI data combined with displaying the cortex extracted from MRI data, we can show how far a reasonable tensor line tracking algorithm can proceed into the outer gyri.

An example on this strategy can be seen in Figs 3 and 4. It can be seen in the transparent area between the outer shell and the visualized DTI data that there is an area of about two to three millimeters in size where no tracking can be performed by streamline algorithms due to ambiguity of the tensor information at the interface between gray and white matter. Different transfer




Fig. 3. Overview of gyri clipped by an almost axial clipping plane. (RGB color coding of DTI and clipping of DTI data based on FA value of 0.3 as used in many fiber-tracking algorithms.) A semi-transparent surface was extracted from the original MRI data set visualizing the boundary of the brain.

functions can be applied to determine which values of FA, apparent diffusion coefficient (ADC) and other tensor parameters can be used as threshold for tractography. Even though there is ambiguity of orientation, the color provides a hint of the behavior of fibers in the gyri. In this case, most fibers would be oriented outwards. The small fibers connecting to the side walls of the gyri cannot be found in the data set using reasonable thresholds. We believe that it is important to transport that information showing the limits of tractography algorithms to the user, to make him understand the images better. Displaying the limits of algorithms leads to increased understanding and, therefore, better confidence in the information that is actually shown.

3.2 Quality Analysis in the Inner Brain

While simple volume rendering provides good hints on how far fiber tracking can advance into the gyri, we can use our algorithm to filter out parts of the data to see what phenomena are occurring inside the brain. It can be seen in the images shown in Fig. 6 that the pyramidal tract is “broken”. This is due to a resolution problem leading to a strong influence of the corpus callosum




Fig. 4. Close-up view of Fig. 3. The white matter in the gyri can be clearly seen. Fiber tracking without additional information stops at the colored surface. The surface color encodes the direction of fibers directly below the surface.

and the superior longitudinal fasciculus connecting the frontal and occipito-temporal part of the brain in these voxels. Tracking fiber bundles in those areas of the data set leads to false results, e.g., they follow the wrong bundles in areas where too many fibers pass closely to each other.

It is an open problem whether there is a way to apply volume rendering to higher-order tensor approaches that are usually used to avoid some of these problems. Additional information of the tensor can be used to produce a better angular structure of the tensors that makes it possible to follow different fiber tracts in these regions [TRWW03, ?].

3.3 Detection and Analysis of Diseases

Many diseases of the human brain, among them Alzheimer’s disease and multiple sclerosis, are related to a change of connectivity in the brain. One special disease is the Deletion 22q Syndrome, which is characterized by increased FA and ADC in certain areas of the brain (cf. Simon et al. [SDB⁺05] and Chapter 5 by Cook et al.). These parameters can be easily displayed with our volume rendering approach, providing a context-preserving visualization that helps to easily locate the areas of change in the brain. Despite the fact that one can obtain a good overview of this kind of data sets, 2D transfer functions that base on ADC and FA and statistical values obtained from a larger pool of subjects have to be implemented and will be subject to future research.

4 Conclusions and Future Work

We have described a novel method that allows us to interactively explore complex volumetric DTI data sets. The method has been applied to different data sets, and two major applications have been highlighted where patterns could be identified in less time than using conventional methods. We have demonstrated that all proposed methods can be easily implemented in existing medical visualization systems.

Acknowledgements

We thank the “German Academic Exchange Service” (DAAD) for partially funding this research and for making this collaboration possible (M. Hlawitschka was supported by a DAAD grant.) Furthermore, we want to thank the Max Planck Institute for Human Cognitive and Brain Sciences in Leipzig, Germany, and Cameron S. Carter, University of California, Davis, Imaging Research Center, for providing the data sets used in this research. We thank the members of the Visualization and Computer Graphics Research Group at the Institute for Data Analysis and Visualization (IDAV) at the University of California, Davis, USA, and the members of the FAnToM group at the University of Leipzig, Germany, and Xavier Tricoche at the University of Utah, Salt Lake City, USA, and Christoph Garth at the University of Kaiserslautern, Germany. Work was supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

- [BBVP05] Jorik Blaas, Charl P. Botha, Frans M. Vos, and Frits H. Post. Fast and reproducible fiber bundle selection in DTI visualization. In Cláudio T. Silva, Eduard Gröller, and Holly Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 59–64, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press.
- [BG06] Stefan Bruckner and M. Eduard Gröller. Exploded views for volume data. In Eduard Gröller, Alex Pang, Cláudio T. Silva, John Stasko, and Jarke van Wijk, editors, *Proceedings of IEEE Visualization’06*, pages 1077–1084, Los Alamitos, CA, USA, 2006. IEEE Computer Society Press.
- [BGKM⁺04] Naama Barnea-Goraly, Hower Kwon, Vinod Menon, Stephan Eliez, Linda Lotspeich, and Allan L. Reiss. White matter structure in autism: Preliminary evidence from diffusion tensor imaging. *Biological Psychiatry*, 55:323–326, 2004.
- [BJB⁺03] J Burns, D Job, ME Bastin, H Whalley, T Macgillivray, EC Johnstone, and SM Lawrie. Structural disconnectivity in schizophrenia: a diffusion tensor magnetic resonance imaging study. *British Journal of Psychiatry*, 182:439–443, May 2003.

- [BL92] Peter J. Basser and Denis LeBihannis. Fiber orientation mapping in an anisotropic medium with NMR diffusion spectroscopy. *11th Annual Meeting of the SMRM, Berlin*, page 1221, 1992.
- [BSP⁺93] Eric A. Bier, Maureen C. Stone, Ken Pier, Buxton William, and Tony D. DeRose. Toolglass and magic lenses: The see-through interface. In *Proceedings of Siggraph '93*, pages 73–80. ACM, 1993.
- [BVG05] Stefan Bruckner, Ivan Viola, and Meister Eduard Gröller. Volumeshop: Interactive direct volume illustration. In *ACM Siggraph 2005 DVD Proceedings (Technical Sketch)*, August 2005.
- [CSC06] Carlos D. Correa, Deborah Silver, and Min Chen. Feature aligned volume manipulation for illustration and visualization. In Eduard Gröller, Alex Pang, Cláudio T. Silva, John Stasko, and Jarke van Wijk, editors, *Proceedings of IEEE Visualization'06*, pages 1069–1067, Los Alamitos, CA, USA, 2006. IEEE Computer Society Press.
- [DH92] Thierry Delmarcelle and Lambertus Hesselink. Visualization of second order tensor fields and matrix data. In *Proceedings of IEEE Visualization 1992*, page 316, Los Alamitos, CA, USA, 1992. IEEE Computer Society Press.
- [EHK⁺06] Klaus Engel, Markus Hardwiger, Joe M. Kniss, Christof Rezk-Salama, and Daniel Weiskopf. *Real-Time Volume Graphics*. A K Peters, Ltd, Wellesley, MA 02482, 2006.
- [ESM⁺05] Frank Enders, Natascha Sauber, Dorit Merhof, Peter Hastreiter, Christopher Nimsy, and Marc Stamminger. Visualization of white matter tracts with wrapped streamlines. In Cláudio T. Silva, Eduard Gröller, and Holly Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 51–58, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press.
- [FCI⁺01] M Filippi, M Cercignani, M Inglese, MA Horsfield, and G Comi. Diffusion tensor magnetic resonance imaging in multiple sclerosis. *Neurology*, 56(3):304–311, February 2001.
- [HBPA01] Khader M. Hasan, Peter J. Basser, Dennis L. Parker, and Andrew L. Alexander. Analytical computation of the eigenvalues and eigenvectors in DT-MRI. *Journal of Magnetic Resonance*, 152:41–47, 2001.
- [JLH⁺99] Derek K. Jones, David Lythgoe, Mark A. Horsfield, Andrew Simmons, Steve C. R. Williams, and Hugh S. Markus. Characterization of white matter damage in ischemic leukoaraiosis with diffusion tensor MRI. *Stroke*, 30:393–397, 1999.
- [Kin04] Gordon Kindlmann. *Visualization and Analysis of Diffusion Tensor Fields*. PhD thesis, School of Computing, University of Utah, Salt Lake City, UT 84112, USA, September 2004.
- [KKW05] Polina Kondratieva, Jens Krüger, and Rüdiger Westermann. The application of GPU particle tracing to diffusion tensor field visualization. In Cláudio T. Silva, Eduard Gröller, and Holly Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 73–78, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press.
- [KTW06] Gordon Kindlmann, Xavier Tricoche, and Carl-Fredrik Westin. Anisotropy creases delineate white matter structure in diffusion tensor MRI. In *Ninth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'06)*, Lecture Notes in

- Computer Science 4190, pages 126–133, Copenhagen, Denmark, October 2006.
- [KW99] Gordon Kindlmann and David Weinstein. Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *VIS '99: Proceedings of the conference on Visualization '99*, pages 183–189, Los Alamitos, CA, USA, 1999. IEEE Computer Society Press.
 - [MMM⁺02] Elias R. Melhem, Susumu Mori, Govind Mukundan, Michael A. Kraut, Martin G. Pomper, and Peter C. M. van Zijl. Diffusion tensor MR imaging of the brain and white matter tractography. *American Journal of Roentgenology*, 178(1):3–16, January 2002.
 - [Mos02] Michael Moseley. Diffusion tensor imaging and aging — a review. *NMR In Biomedicine*, 15:553–560, 2002.
 - [MVvW05] Bart Moberts, Anna Vilanova, and Jarke J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In Cláudio T. Silva, Eduard Gröller, and Holly Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 65–72, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press.
 - [PP99] Sinisa Pajevic and Carlo Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42 (3):526–540, 1999.
 - [SDB⁺05] T. J. Simon, L. Ding, J. P. Bish, D. M. McDonald-McGinn, E. H. Zackai, and J. Gee. Volumetric, connective, and morphologic changes in the brains of children with chromosome 22q11.2 deletion syndrome: an integrative study. *NeuroImage*, 25:169–180, 2005.
 - [TRWW03] David Solomon Tuch, Timothy G. Reese, Mette R. Wiegell, and Van J. Wedeen. Diffusion MRI of complex neural architecture. *Neuron*, 40:885–895, December 2003.
 - [VKG04] Ivan Viola, Armin Kanitsar, and Meister Eduard Gröller. Importance-driven volume rendering. In Holly Rushmeier, Greg Turk, and Jarke J. van Wijk, editors, *Proceedings of IEEE Visualization'04*, pages 139–145, Los Alamitos, CA, USA, 2004. IEEE Computer Society Press.
 - [VZKL06] A. Vilanova, S. Zhang, G. Kindlmann, and D Laidlaw. An introduction to visualization of diffusion tensor imaging and its applications. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, pages 121–153. Springer-Verlag Berlin Heidelberg, 2006.
 - [WPG⁺97] Carl-Fredrik Westin, Sarel Peled, Hakon Gudbjartsson, Ron Kikinis, and Ferenc A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM '97*, page 1742, Vancouver Canada, April 1997.
 - [WZMK05] Lujin Wang, Ye Zhao, Klaus Mueller, and Arie Kaufman. The magic volume lens: An interactive focus+context technique for volume rendering. In Cláudio T. Silva, Eduard Gröller, and Holly Rushmeier, editors, *Proceedings of IEEE Visualization 2005*, pages 65–72, Los Alamitos, CA, USA, 2005. IEEE Computer Society, IEEE Computer Society Press.
 - [ZB02] Leonid Zhukov and Alan H. Barr. Oriented tensor reconstruction: Tracing neural pathways from diffusion tensor MRI. In *Proceedings of IEEE Visualization '02*, pages 387–394, Los Alamitos, CA, 2002. IEEE Computer Society.




Fig. 5. Original image and x-direction (left-right) removed from DTI data to allow to produce a cleaner view of the singuli and the pyramidal tract. The lasting red component is the corpus callosum as seen in the MRI data and is shown here to provide context.




Fig. 6. Pyramidal tract (blue) highlighted — other directions are faded out. The pyramidal tract is broken in this data set by the influence of the corpus callosum (red) and the green bundle passing by.




Fig. 7. Left: Top/frontal view on the prefrontal lobe of a healthy subject. The cingulum (green) and the corpus callosum (red) are clearly visible. Right: Top left view highlighting the pre- and postcentral sulcus. The underlying superior longitudinal fasciculus (green) is also visible.




Fig. 8. Left: Side view on the parietal and superior temporal lobe showing the superior longitudinal fasciculus (green, middle of the loop) and the inferior fronto-occipital fasciculus (green, lower left of the lens). Right: Dorsal view on medial brain regions showing also the cingulum (green) and the corpus callosum with two different shadings of the magic lens. Two different shading models of the lens are shown.