

Direct Visualization of Fiber Information by Coherence

Mario Hlawitschka · Christoph Garth ·
Xavier Tricoche · Gordon Kindlmann · Gerik
Scheuermann · Kenneth I. Joy · Bernd Hamann

Received: Dec 15, 2008 / Accepted: date

Abstract

Purpose: The structure of fiber tracts in DT-MRI data presents a challenging problem for visualization and analysis. We derive visualization of such traces from a local coherence measure and achieve much improved visual segmentation.

Methods: We introduce a coherence measure defined for fiber tracts. This quantitative assessment is based on infinitesimal deviations of neighboring tracts and allows identification and segmentation of coherent fiber regions. We use a hardware-accelerated implementation to achieve interactive visualization on slices and provide several approaches to visualize coherence information. Furthermore, we enhance existing techniques by combining them with coherence.

Results: We demonstrate our method on both a canine heart, where the myocardial structure is visualized, and a human brain, where we achieve detailed visualization of major and minor fiber bundles in a quality similar to and exceeding fiber clustering approaches.

Conclusions: Our approach allows detailed and fast visualization of important anatomical structures in DT-MRI data sets.

Keywords Diffusion Tensor Imaging · Tractography · Coherent Structures

1 Introduction

The abstract mathematical concept of tensors remains a challenging field of study in scientific visualization. Visual representation of a tensor can be achieved by reducing its multidimensional information to simple, understandable geometric primitives in three-dimensional

M. Hlawitschka, C. Garth, K. I. Joy, B. Hamann,
E-mail: (hlawitschka|cgarth|kijoy|bhamann)@ucdavis.edu
Institute for Data Analysis and Visualization, University of California, Davis, CA, USA

Gordon Kindlmann, E-mail: gk@bwh.harvard.edu
Laboratory of Mathematics in Imaging, Dept. of Radiology
Brigham and Women's Hospital, Harvard Medical School, Boston, MA, USA

Xavier Tricoche, E-mail: xmt@purdue.edu
Dept. of Computer Science, Purdue University, West Lafayette, IN, USA.

G. Scheuermann, E-mail: scheuermann@informatik.uni-leipzig.de
Abteilung Bild- & Signalverarbeitung, Universität Leipzig, Germany

space. While these so-called *tensor glyphs* are powerful in conveying local properties of a tensor-valued data set [Kin04], global insight into its structure independent of the field of application, is most effectively achieved through line primitives, which constitute the most prominent tool to visualize tensor data in a continuous way. *Tensor lines* and *hyperstreamlines* [VZKL06], which represent curves that are everywhere tangent to a tensor eigenvector field, have proven to be a suitable means to globally visualize tensor fields in a number of application domains including material sciences and medical visualization. In *Diffusion Tensor Imaging* (DTI) [BML94] in particular, it is known that eigenvectors of reconstructed diffusion tensors align with the underlying local fibrous tissue structure. Its analysis provides insight into the connectivity of the brain white matter. The ability of DTI to acquire this information *in vivo* makes it a valuable tool to neuroscientists and neurosurgeons.

While single tensor lines can represent only a small subset of a considered data set, displaying large numbers of these lines suffers from visibility and occlusion problems, which significantly limits the effectiveness of visualization of this approach. Therefore, various methods have been proposed to reduce visual occlusion by placing a certain number of line representatives or by displaying boundary surfaces of tensor line bundles that exhibit a similar behavior. Their construction usually involves complex optimization or clustering algorithms (e.g. [ESM*05, STS07]), which remain computationally challenging. In contrast, an effective strategy to produce more abstract means to capture and represent tensor data and create a high-level representation usually consists of emphasizing regions where the behavior of lines changes. These regions are of primary importance in the analysis of the data. In medical imaging in particular, the seminal work by Basser et al. [BP*00] have demonstrated the anatomical significance of the structures derived from tensor lines in terms of fiber tracts. In brain white matter where the connectivity of different brain areas is a vital field of research, regions where the local geometry of tracks changes rapidly are structurally meaningful since they indicate areas where at least two neural fiber bundles touch, join, or separate. Similarly, diffusion-weighted scans of muscles support *in vivo* structural analysis even of small-animal skeletal muscle. Therefore, they can be used to analyze complex muscular structures such as found in the heart.

In this context, we propose a novel visualization method that combines dense fiber integration with local analysis of the *coherence* of the resulting fiber set to yield a structural picture of DT-MRI data sets. Specifically, the algorithm we propose effectively identifies each data set location with the fiber tract passing through it, and quantifies a coherence measure on the fiber tracts passing through a small neighborhood. We define this measure in terms of the deviation of fibers over a finite distance. It follows from this construction that regions of high cohesiveness contained in fibrous tissue corresponds to a set of points that exhibit a strong coherence with their neighbors, while the boundaries that separate such structures are associated with points where the coherence is small. An important feature of our method is the continuous nature of the coherence measure; this characteristic makes proper modeling of gradual transitions between differing fiber behaviors possible. This property is of particular importance to study the subtle changes in the tissue structure, such as, e.g., those found in the myocardium. Furthermore, because our method quantifies fiber coherence with respect to the geometric similarity of tensor lines over a certain length, this length parameter is naturally contained in our approach and provides a intuitive notion of scale in the assessment of coherence. Remark that the coherence measure presented here is not specific to medical visualization but can be applied in the general context of tensor field analysis.

In Section 2, we review previous work in tensor visualization, with a strong focus on techniques dedicated to the characterization of structures in tensor fields. Section 3 introduces the theoretical foundations underlying the proposed fiber measure, and discusses de-

tails of our implementation. In Section 4, we examine a number of visualization approaches derived from the obtained scalar coherence information and explain how it can be used to enhance commonly used visualization techniques. Specific results for the case of DTI visualization are discussed in Section 6, and we demonstrate the ability of our technique to characterize different types of structural properties in various kinds of tissue. Finally, Section 7 concludes on the presented material and provides possible directions for future research.

2 Related Work

Visualization of tensor-valued data sets using tensor lines can be traced back to the visualization of vector fields and flows using so-called integral curves that model the trajectories of particles moving through the field. The recently introduced notions of *Finite-Time Lyapunov Exponents* (FTLE) and *Lagrangian Coherent Structures* (LCS) [Hal01,LSM06] have been applied to visualize the coherence of motion among neighboring particles [GGTH07, SP07]). These *Lagrangian* visualization approaches locally express the change of the particle trajectory with the variation of the initial position, and regions of locally high variation are identified as coherent structures that represent the driving constituents of a flow. The tensor visualization approach we present is based on similar concepts applied to tensor lines.

The quantification of coherence of neighboring fibers in diffusion tensor data, specifically pertaining to the human brain, has been attempted by several authors [ESM*05, MVvW05, QRO*08]. Typically, fiber tracts that remain close to each other and have similar behavior are clustered and classified to belong to a common bundle. We do not compute a clustering but rather highlight the *Lagrangian* scalar field directly, as it provides additional information.

Especially in diffusion tensor data, these coherent structures possess a meaning on the connectivity of the underlying structure and provide valuable information for further evaluation of the data. These concepts have strongly influenced our idea of detecting coherent structures in tensor fields. In the following, we describe the theoretical considerations underlying our approach.

3 Tensor Line Coherence

Before we describe our method, we briefly revisit the concept of tensor lines.

3.1 Tensor Lines

Tensor lines were originally introduced by Dickenson et al. [Dic89] and later applied to the topological analysis of tensor fields [HLL97,TS04] and the segmentation of Diffusion-Tensor MRI data [ESM*05]. A given field of symmetric tensors T , represented by $n \times n$ -matrices induces a local eigen decomposition of the form $T(x) = B(x)^T D(x) B(x)$ for every x in the domain of definition of T . Here, $B(x)$ represents an orthogonal basis transform and D is the diagonal matrix of non-negative eigenvalues λ_i ordered by magnitude, i.e., $\lambda_1(x) \geq \dots \geq \lambda_n(x)$. If $\lambda_i \neq \lambda_j$ for $i \neq j$, this ordering is strict; in this case x is called a *regular point*, and *degenerate* otherwise.

Choosing $j \in \{1, \dots, n\}$, the decomposition gives rise to a tensor line $\varphi_x(t)$ as the uniquely defined curve that passes through x and is everywhere tangential to the j -th eigen-

vector $e_j(x)$ corresponding to $\lambda_j(x)$, i.e.,

$$\frac{d}{dt}\varphi_x = \pm e_j(x). \quad (1)$$

Hence, tensor lines are integral curves that are naturally parameterized according to arc length. In the case $n = 3$, the three possible tensor line types are called *major*, *medium*, and *minor* tensor lines depending on the choice of j . A variety of numerical schemes exist for approximations of tensor lines in symmetric tensor fields, typically based on numerical integration schemes such as the explicit Euler scheme or the family of Runge-Kutta methods [HNW93]. Remark that the orientation of φ_x is not uniquely determined by Eq. 1. This must be taken into account when working with such lines.

3.2 Coherence Measure

The central idea of our approach is a local coherence measure that describes the similarity of tensor lines passing through a small neighborhood U around a regular point x . We may assume in the following that all the tensor lines passing through U can be consistently oriented. For $t > 0$, we denote by $D\varphi_x^t$ the derivative of $\varphi_x(t)$ with respect to x . Essentially, $D\varphi_x^t$ describes the variation of $\varphi_x(t)$ as x varies slightly. Intuitively, the spectral matrix norm

$$\|D\varphi_x^t\|_2 := \sqrt{\lambda_{\max}(D\varphi_x^t D\varphi_x^t)}$$

pertaining to the maximal eigenvalue of $D\varphi_x^t$ can be interpreted as the maximum distance after a time t that any φ_y can move away from φ_x if x and y are initially arbitrarily close. Fixing $T > 0$, we now define the local *tensor line coherence measure* C by virtue of

$$C_T(x) := \max_{-T \leq t \leq T} \|D\varphi_x^t\|_2. \quad (2)$$

Essentially, $C_T(x)$ measures the maximum distance that a tensor line φ_y passing through any y arbitrarily close to x can move away from φ_x over its parameter interval $[-T, T]$. If this maximum distance (and hence C_T) is small, neighboring tensor lines stay closely together. On the other hand, a large C_T indicates significant divergence along the path.

3.3 Numerical Approximation

In the interest of simplified notation, we limit our description here to a two-dimensional regularly spaced setting. However, equivalent constructions are possible in any spatial dimension and on most types of discrete domains.

Let $T > 0$ and assume a discrete domain with points $x_{i,j} := (i \cdot h, j \cdot h)$ with $h > 0$ and $i, j \in \{0, \dots, N\}$. Furthermore, let $M > 0$ an integer defining a time discretization $t^m := \frac{k/M}{T}$ with $m \in \{-M, \dots, M\}$. Using numerical integration, we define the discrete tensor line points

$$p_{i,j}^m := \phi_{x_{i,j}}(t^m)$$

as the discrete representation of the tensor lines passing through the grid points $x_{i,j}$. We find the discrete approximation

$$\tilde{C}_T(x_{i,j}) = \max_{m=-M, \dots, M} \|D\phi_{x_{i,j}}^{t^m}\|_2 \quad (3)$$

of C_T (Eq. 2), where $D\phi_x^m$ is numerically approximated from the $p_{i,j}^m$, e.g., using finite differences. Note the discrete tensor lines passing through the grid points participating in the derivative approximation must be oriented consistently. We now examine the visualization of DTI data sets using the above coherence measure.

4 Visualization

The point-wise coherence measure defined in the previous section, applied to the entire domain of a data set, gives rise to a scalar field. Hence, typical scalar fields visualization methods are applicable. For a fast, exploratory inspection of such fields, pseudo-colored sweeping planes have proven to be a valuable, interactive tool. We have considered different types of color coding: Figure 3 (left) depicts an anteroposterior slice through a brain data set where the coherence measure is represented as a grayscale color map. The visualization can be further augmented by superimposing an anisotropy-scaled rgb color map that is prevalent in DTI visualization [KWH00]. The typical anatomical structures can be identified while additional information is provided by the coherence visualization (middle image). Fiber endpoint information can be used as an alternative color coding that enhances the structural view on the data. By mapping the direction between the endpoints using $\frac{(\delta_x, \delta_y, \delta_z)}{|\delta_x, \delta_y, \delta_z|}$ to (r, g, b) color components, generally, similar colors are assigned to areas that have similar behavior (right image). Figure 2 illustrates some of these visualization approaches on a DTI data set of a canine heart.

Regarding the three-dimensional context of a data set, the main structural components of the coherence field can be visualized using isosurfaces or volume rendering. Since low coherence values correspond to diverging fiber bundles, corresponding isosurfaces or volume renderings emphasizing such regions illustrate boundaries between coherent regions. Conversely, areas of similar behavior that are representative of the major line structures may be visualized by selecting regions of high coherence values.

Furthermore, we propose to enhance typical tensor line visualization through coherence information. Here, we make use of the coherence measure as a transparency mask for the rendered lines. This reduces in a stronger depiction of coherent structures while reducing the visual impact of regions of incoherent fiber behavior (cf. Figure 5). The use of illuminated lines [MPSS05] emphasizes the three-dimensional structure of the data and improves the perception of the topological relationships among lines, especially when interaction with the data is desired and possible. Furthermore, adding volumetric coherence information as discussed above provides spatial context to the fiber visualization.

In the next section, we focus on implementation aspects of coherence-based DTI visualization.

5 Implementation and Performance

The computation of the coherence measure for a given DTI data set is built on the observation of a large number of fiber traces traversing the domain of definition. Even for small data sets, a regular sampling of the data set domain can result in the computation of a large number of fiber traces. The resulting computational effort is significant. This problem is further aggravated by the observation that important structures in tensor fields often have sub-voxel resolution being a consequence of interpolation. To resolve these structures and gain a proper interpretation of the data set, we have found it beneficial to evaluate coherency

information with resolution exceeding that of the underlying data set. This super-sampling further increases the computational effort of our method.

We have addressed this problem in the following way: We make use of Graphics Processing Units (GPUs) to accelerate the necessary computations. Tensor line integration is a massively parallel process; there is no interdependency between lines emanating from different points and hence maps ideally to the increased computational parallelism and bandwidth available from modern GPUs. Our GPU implementation of the coherence measure is based on central finite differences. For each point at which the coherence measure must be evaluated, the six neighboring fiber traces corresponding to the finite difference stencil are computed and oriented consistently. Here, all fiber traces are advanced one step simultaneously, and after each step, $D\phi^m$ is evaluated and \tilde{C}^T is updated according to Eq. 3. This incremental algorithm permits us to approximate the coherence measure during the integration and obviates the need to store the fiber traces.

Practically, it is not always possible to advance fibers to the maximum length over which the coherence measure is to be evaluated. Fiber traces cannot continue over domain boundaries or part degenerate points in the data set. Furthermore, it is common practice to terminate fiber integration if the *fractional anisotropy* – an indirect indicator of the presence of white matter – falls below a specified threshold (cf., e.g., [Kin04]). In all these cases, we evaluate the coherence measure only over the interval on which all fibers of the derivative stencil exist.

Our implementation is straightforward since we do not expend effort to reuse fiber traces among neighboring evaluation locations. We have found that the storage and bandwidth requirements of a more elaborate implementation that computes each fiber trace exactly once are excessive in the face of the super-sampling discussed above. Hence, in our experiments, such an implementation did not perform faster than the naive implementation presented here.

To document the performance that can be achieved using the presented approach and examine its suitability for interactive visualization, we have measured timings obtained on two data sets. All timings were obtained on commodity PC (2.0GHz Intel Core 2 Duo with 2GB RAM, GeForce 8800GTX, 768MB video RAM) and are shown in Table 1. The measured computation times include all fiber integrations and evaluation of the coherence field for different resolutions of the coherence map on both slices and volumes. A further discussion of the data sets is given in the next section 6.

\tilde{C}_T resolution	Data Set	Time	Data Set	Time
512×512	Human Brain	10ms	Dog Heart	
1024×1024	($93 \times 116 \times 93$)	20ms	($64 \times 52 \times 41$)	
512^3		10s		2.4s

Fig. 1 Evaluation timings of our coherency measure implementation for two DT-MRI data sets.

For the slice-based visualization mode (cf. Fig. 3), a sweep is performed through the data set of color-coded slices using an advection approach [WKL99] at interactive frame rates up to 1024^2 pixel resolution as is shown from Table 1. For inspecting single slices and changing the color parameters, the lines are not recomputed and, therefore, the frame rates for the visualization step is much higher. All other changes that do not affect the line integration but the coherence calculation are performed at the same frame rate.




Fig. 2 Slice through a diffusion-weighted MRI scan of a canine heart data set: Even though the data is noisy and the resolution low, grayscale coherence map visualization (left) and coherence map colored by directional information clearly show the main structures of the data. The muscle fibers are organized in layers that differ in their fiber direction. Whereas the transition is gradual within a linear interpolated voxel, our approach provides sub-voxel accuracy of the barrier. Whereas previous methods ([PVS06]) have already been able to extract this behavior, our method does not rely on abrupt changes within a voxel but uses global information of line behavior. Therefore, it is superior to previously presented approaches that cannot extract results on a sub-voxel scale.

6 Results

We have applied our method to two DT-MRI data sets — a diffusion-weighted magnetic resonance scan of a canine heart, and a diffusion tensor image set of the human brain.

Canine Heart The MRI scan of the canine heart has been sampled over a regular grid with resolution $64 \times 52 \times 41$ and with anisotropic cells of size $1.4 \times 1.4 \times 2.0 \text{mm}^3$. It was acquired with the specific intent to analyze the heart’s muscular structure that is visible in the DTI images as the major tensor eigenvectors are aligned to the underlying fiber structure.

Two slice visualizations with and without additional coloring are shown in Figure 2. As the data contains a hard mask, boundary artifacts are observed in the upper part of the pictures. Nevertheless, the main apparent structures in these images are two trisector-like points, one at the upper connection between atrium and the ventricle, and the other near the bottom. The separating wall contains one major line separating the fibers going around the atrium from those following the left ventricle. More interesting is the wedge-like point near the bottom: It splits the fibers coming from the atrium in two parts. The outer layer passes (as seen in the red color coding) from left to right, while the inner layer of fibers (blue) is oriented orthogonally to the slice.

Human Brain We applied our algorithm to a DTI data set describing a human brain. It was generated from a healthy volunteer using a three-Tesla Siemens Trio MRI scanner. 60 diffusion weighted images were acquired using three times averaging and 21 baseline images. The data was converted to a second-order tensor representation using linear least-squares fit [BML94]. The complete measurement took about 20 minutes with an in-slice




Fig. 4 A combination of volume rendering of tensor coherence measure and coherence-based seeding of tensor lines. Many small-scale structures are visible, such as the region of low coherence between the corpus callosum (fibers coming out of the picture) and the singuli (fibers going left-right above the corpus callosum).




Fig. 3 GPU-based visualization of slices. The coherence map (left) is drawn as an overlay to a color map in the middle and right picture. Black lines mark low coherence while the white areas contain coherent structures, which are colored by their direction in the middle and right pictures. In the middle picture, the medical fractional anisotropy-rgb color map is used to show the direction and fractional anisotropy at the seed point. In the right picture, the areas of similar behavior are enforced using an endpoint-based directional color map to provide a segmentation-like color coding. The selected slice, its orientation, and the color coding can be changed interactively.

resolution of 128×128 voxel and 72 slices on a $1.7 \times 1.7 \times 1.7$ mm³ grid. In contrast to the dog heart, the noise in the white matter part is low, but the fiber structure is more complex.

Figure 3 shows three different visualizations of fiber coherence on a central slice of the brain, using purely the coherence measure (left), coloring by anisotropy (middle), and coloring by endpoint difference. The coherence measure was evaluated using a sampling resolution of 1024^2 . Dominant structures such as the corpus callosum and the pyramidal tract are clearly identified, and a detailed visual segmentation of the fiber tracts is observed even in regions that have almost uniquely local directional information. Both forceps major and minor are segmented in different areas depending on the fiber's global behavior. In the right image, the distinction between the tapetum (purple), posterior corona radiata (green) and superior longitudinal fasciculus (blue) is visible, similar to Fig. 10 in Kindlmann et al. [KTW07] with the difference that the coherence visualization produced by our scheme exhibits significantly reduced noise and identifies cleaner boundaries.

For fully three-dimensional analysis, we use the coherence field as a guide for the seeding and filtering of tensor lines. A set of ca. 14 000 illuminated tensor lines uniformly seeded inside the brain volume is shown in Figure 5 (left). The right image in the same figure shows the same set of tensor lines, where the transparency of the line is derived from the coherence of the regions that the fibers traverse such that more coherent regions appear more opaque. This effectively achieves a filtering of lines in noisy areas and enforces a visual emphasis of the global structure of the field. In addition to the major fiber bundles, minor bundles leaving the pyramidal tract towards the cerebral cortex can be distinguished. A combination of opacity-modulated fiber tracts with volume rendering of coherence information for ad-




Fig. 5 Seeding and transparency mapping of tensor lines depending on coherence (right) versus an unfiltered version of the same line set (left). The right image only shows the most dominant and coherent fiber bundles.

ditional context is shown in Figure 4. Whereas individual fibers that can not be associated with fiber bundles vanish, even fine-grained structures at sub-voxel resolution, such as the singuli, remain visible.

The brain data set presents a challenge with regard to visualization, as it exhibits a relatively large amount of noise and many small-scale structures. Coherence-based visualization successfully reveals the main structures of the data and boundaries of areas of different behavior can be clearly identified.

7 Discussion

The visualization method for fiber traces in DTI data sets presented here was inspired by two other successful visualization methodologies, namely coherent structures in flow fields and fiber clustering. They share the common idea of describing a data set through information defined by integral lines. Our approach uses a coherence measurement as a means to identify coherent, similarly behaving, line structures. Therefore, the visualizations derived from fiber coherence images are a direct representation of clustering distance measures on lines in a neighborhood, but it supersedes clustering because there is no need for an explicit and discrete grouping of lines. Therefore, floating point values can be used to represent similarity and due to the higher efficiency, it is possible to use far more lines for the calculation, which results in a higher resolution and a detection of finer structures compared to most clustering approaches. We have demonstrated that those measures can be used to effectively visualize global structures in DTI data and provide the user with a representation that enables a visual segmentation. Furthermore, combination with isosurfaces and volume rendering provides coherency information in spatial context for three-dimensional visualizations. This is an improvement over previous clustering, segmentation and visualization approaches that require extensive pre-computation. We have described a robust algorithm that leverages commodity GPUs to quickly measure coherency information on DTI data sets and have shown that it facilitates interactive visualizations. Furthermore, since it behaves smoothly, the coherency measure is not strongly influenced by noise such as, for example, topological methods. In addition to the direct interpretation of the coherence value and its similarity to clustering ap-

proaches, it can be used as a fully automatic approach for tensor line seeding or filtering of lines to provide appealing three-dimensional visualizations that significantly reduce clutter.

Several open questions remain as avenues for future research. Firstly, we would like to investigate the connection between our coherence measure and topological methods. Topological methods describe the coherence of fibers in an asymptotic sense, while we only measure it over a finite length interval. Secondly, we envision an automated extraction of boundaries between coherent fiber bundles by means of ridge extraction in the coherence field. Finally, we note that the scope of our work is not limited to the visualization of tensor data from medical applications, and we have already taken first steps towards applying it to other application domains. Thirdly, as our current DTI data sets are provided by neuroscientists focusing on brain connectivity, an additional study involving neurosurgeons and radiologists has to reveal possible applications for clinical data.

Acknowledgements We thank Alfred Anwander and Thomas R. Knösche from Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany, for providing the human brain data sets, and thank the members of the Visualization and Computer Graphics Research Group of the Institute for Data Analysis and Visualization, UC Davis.

Mario Hlawitschka was funded by the University of Leipzig, Germany and NSF grant CCF-0702817. Christoph Garth was supported by the Director, Office of Advanced Scientific Computing Research, Office of Science, of the U.S. Department of Energy under Contract No. DE-FC02-06ER25780 through the Scientific Discovery through Advanced Computing (SciDAC) program's Visualization and Analytics Center for Enabling Technologies (VACET). We wish to thank Edward Hsu, University of Utah, for providing access to the canine heart data set.

References

- [BML94] BASSER P., MATTIELLO J., LEBIHAN D.: Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance* 3, 103 (1994), 247–254.
- [BP*00] BASSER P. J., PJEVIC S., PIERPAOLI C., DUDA J., ALDROUBI A.: In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* (2000), 625–632.
- [Dic89] DICKENSON R. R.: A unified approach to the design of visualization software for the analysis of field problems. In *SPIE Proceedings* (Bellinham, Washington, 1989), vol. 1083, SPIE - The International Society for Optical Engineering.
- [ESM*05] ENDERS F., SAUBER N., MERHOF D., HASTREITER P., NIMSKY C., STAMMINGER M.: Visualization of white matter tracts with wrapped streamlines. In *Proceedings of IEEE Visualization 2005* (Los Alamitos, CA, USA, 2005), Silva C. T., Gröller E., Rushmeier H., (Eds.), IEEE Computer Society, IEEE Computer Society Press, pp. 51–58.
- [GGTH07] GARTH C., GERHARDT F., TRICOCHÉ X., HAGEN H.: Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transactions on Visualization and Computer Graphics TVCGs* 13, 6 (November/December 2007), 1464–1471.
- [Hal01] HALLER G.: Distinguished material surfaces and coherent structures in three-dimensional flows. *Physica D* 149 (2001), 248–277.
- [HLL97] HESSELINK L., LEVY Y., LAVIN Y.: The topology of symmetric, second-order 3D tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 3, 1 (1997), 1–11.
- [HNW93] HAIRER E., NØRSETT S. P., WANNER G.: *Solving Ordinary Differential Equations I, second edition*, vol. 8 of *Springer Series in Comput. Mathematics*. Springer-Verlag, 1993.
- [Kin04] KINDLMANN G.: Superquadric tensor glyph. *Joint EUROGRAPHICS – IEEE TCVC Symposium on Visualization* (2004).
- [KTW07] KINDLMANN G., TRICOCHÉ X., WESTIN C.-F.: Delineating white matter structure in diffusion tensor MRI with anisotropy creases. *Medical Image Analysis* 11, 5 (October 2007), 492–502.
- [KWH00] KINDLMANN G. L., WEINSTEIN D. M., HART D.: Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 6 (2) (2000), 124–137.
- [LSM06] LEKIEN F., SHADDEN S., MARSDEN J.: Lagrangian coherent structures in n-dimensional systems. *Physica D submitted* (2006).

-
- [MPSS05] MALLO O., PEICKERT R., SIGG C., SADLO F.: Illuminated lines revisited. In *Proceedings of IEEE Visualization 2005* (Los Alamitos, CA, USA, 2005), Silva C. T., Gröller E., Rushmeier H., (Eds.), IEEE Computer Society, IEEE Computer Society Press, pp. 19–26.
- [MVvW05] MOBERTS B., VILANOVA A., VAN WIJK J. J.: Evaluation of fiber clustering methods for diffusion tensor imaging. In *Proceedings of IEEE Visualization 2005* (Los Alamitos, CA, USA, 2005), Silva C. T., Gröller E., Rushmeier H., (Eds.), IEEE Computer Society, IEEE Computer Society Press, pp. 65–72.
- [PVSR06] PEETERS T., VILANOVA A., STRIJKERS G., ROMENY B. T. H.: Visualization of the fibrous structure of the heart. *Vision, Modeling, and Visualization, Aachen* (2006).
- [QRO*08] QAZI A. A., RADMANESH A., O'DONNELL L., KINDLMANN G., PELED S., WHALEN S., WESTIN C.-F., GOLBY A. J.: Resolving crossings in the corticospinal tract by two-tensor streamline tractography: method and clinical assessment using fMRI. *Neuroimage* (2008).
- [SP07] SADLO F., PEICKERT R.: Visualizing lagrangian coherent structures: A comparison to vector field topology. In *Topology-Based Methods in Visualization 2007, to appear* (2007).
- [STS07] SCHULTZ T., THEISEL H., SEIDEL H.-P.: Topological visualization of brain diffusion MRI data. *IEEE Transactions on Visualization and Computer Graphics (Proc. IEEE Visualization)* 13, 6 (2007), 1496–1503.
- [TS04] TRICOCHÉ X., SCHEUERMANN G.: *Topological Methods for Tensor Visualization*. Elsevier, Amsterdam, Dec 2004.
- [VZKL06] VILANOVA A., ZHANG S., KINDLMANN G., LAIDLAW D.: An introduction to visualization of diffusion tensor imaging and its applications. In *Visualization and Processing of Tensor Fields* (2006), Weickert J., Hagen H., (Eds.), Springer-Verlag Berlin Heidelberg, pp. 121–153.
- [WKL99] WEINSTEIN D., KINDLMANN G., LUNDBERG E.: Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *IEEE Visualization '99* (1999), pp. 249–254.