Postprint

Coherent Structures of Characteristic Curves in Symmetric Second Order Tensor Fields

Marcel Hlawatsch, Joachim E. Vollrath, Filip Sadlo, *Member, IEEE,*
and Daniel Weiskopf, *Member, IEEE Computer Society*

Abstract—This paper generalizes the concept of Lagrangian coherent structures, which is known for its potential to visualize coherent regions in vector fields and to distinguish them from each other. In particular, we extend the concept of the flow map to generic mappings of coordinates. As the major application of this generalization, we present a semi-global method for visualizing coherent structures in symmetric second order tensor fields. We demonstrate the usefulness by examples from DT-MRI, uncovering anatomical structures in linearly anisotropic regions not amenable to local feature criteria. To further exemplify the suitability of our concept, we also present its application to stress tensor fields. Lastly, an accelerated implementation utilizing GPUs is presented.

Index Terms—Lagrangian coherent structures, feature extraction, tensor field visualization, tensor field topology, general-purpose computation using graphics hardware

I. INTRODUCTION

MANY phenomena in science and engineering are amenable to an analysis by inherent tensors. Large parts in fluid dynamics and flow visualization build on the velocity gradient tensor, Maxwell's equations for electromagnetism can be described in special-relativistic covariant (form) tensors, local shape is often analyzed with the Hessian matrix, and material science and simulation often make use of tensors describing phenomena such as stress, strain, deformation, or diffusion. This work concentrates on symmetric second order tensors.

In visualization, symmetric second order tensor fields have received much attention since they are generated with the medical imaging technique known as *diffusion tensor magnetic resonance imaging* (DT-MRI). DT-MRI facilitates in-vivo measurements of the rate and orientation of water molecule diffusion within tissue. This technique is therefore particularly well suited for the analysis of fibrous tissue that displays a distinctly oriented organization as in the brain or within muscles. The visualization of symmetric second order tensors also plays an important role in other disciplines, e.g. they can be used to describe mechanical stress.

The goal of this research is to separate the domain of the tensor field into qualitatively different regions and to further structure regions perceived as coherent on a local scale by means of a scalable semi-global analysis. To this end, we contribute a formal definition and an evaluation of a method that allows one to visualize the analog of *Lagrangian*

The authors are with the Visualization Research Center, Universität Stuttgart (VISUS), Allmandring 19, 70569 Stuttgart, Germany. E-mail: {hlawatsch,vollrath,sadlo,weiskopf}@visus.uni-stuttgart.de

coherent structures in symmetric second order tensor fields. This method is derived from a generalization of the *Lyapunov exponent* from dynamical systems theory, the concept of *flow maps*, their gradients, and of Lagrangian coherent structures. We present results for diffusion tensor fields and a stress tensor field, and sketch a fully GPU-based implementation.

II. RELATED WORK

In the past, different methods were developed that aim at providing meaningful visualizations of multivariate tensor data. A diffusion tensor can be visualized geometrically with a diffusion ellipsoid [37], [30], where the axes and radii correspond directly to the eigenvectors and eigenvalues of the tensor. Similarly, boxes [56] or composite shapes [54] can be used, while Kindlmann [24] proposes the use of superquadrics to combine the benefits of both the ellipsoid and box glyphs. Another common practice is to visualize properties derived from diffusion tensor fields on cross sections of the dataset. Widely used are local scalar measures such as mean diffusivity, fractional anisotropy [1], as well as the coefficients for linear anisotropy, planar anisotropy, and spherical isotropy [55], which can also be visualized by direct volume rendering [28]. For the visualization of fiber structure, methods from classical vector field visualization have been adapted. Characteristic lines of a tensor field can be traced by following its eigenvector fields and strategies to reduce resulting visual clutter are available [50]. Hyperstreamlines [5] additionally include the other eigenvectors as thickness transversal to the chosen eigenvector. Tensorlines [53] incorporate a notion of artificial inertia to the line tracing process in order to apply some control on the direction of tracing when areas of insufficient linear anisotropy are crossed. The HyperLIC approach [58], where a noise field is filtered using primitives deformed by the tensor field, provides a dense tensor field visualization in the style of line integral convolution [4]. Further references to DT-MRI visualization techniques are given in the survey article by Vilanova et al. [51].

Tractography in diffusion tensor fields deals with the methodology of (in-vivo) reconstruction of fiber tracts in organized tissue. A technical review is given by Mori et al. [34]. Concerning cerebral tissue, the greater goal of tractographical methods is to uncover the organization of the brain, e.g. in order to understand its functional architecture [31] or to construct atlases [52]. Concerning muscle tissue, the discussion on its organization is in parts still ongoing [14], and tractographical methods based on DT-MRI [61], [36], [39] can

aid its uncovering. An effective way to gain such insight in the case of cerebral tissue is to group extracted fibers according to similarity and to connectivity with different functional areas. Brun et al. [2] propose a mapping of each fiber to an Euclidean feature space and interpret the pairwise distance of fibers in this space as a similarity measure that they represent as a weighted undirected graph. Subsequent normalized cuts in this similarity graph allow for an unsupervised segmentation of fiber bundles. Ding et al. [8] propose a different fiber similarity measure based on the ratio of fiber lengths and the average Euclidean distance of overlapping fiber segments. A *k-most-similar-fibers* algorithm then allows them to identify bundles of fibers and to classify them.

In visualization, topological analysis of tensor fields was first performed by Delmarcelle and Hesselink [6]. Such methods extract the essential structure and properties of the multivariate tensor field, thus leading to a significant reduction of information to be assessed. This is done by extracting two basic types of characteristic features: *critical* or *degenerate points*, which represent singularities in the field, and *separatrices*, which connect to some of these points and separate regions of different behavior. Scheuermann and Tricoche [42] provide an overview of topological analysis of both vector and tensor fields, and a thorough treatment is given by Tricoche [46]. The topology of 3D tensor fields is analyzed by Zheng and Pang [59] and Zheng et al. [60]. Tricoche et al. [48] use invariants for the extraction of topological features from 3D tensor fields.

Other approaches aim at extracting essential information from tensor fields by means of general feature extraction, such as interfaces [35], creases [26], [27], edges [57], and changes in tensor shape and orientation [25].

The focus of this work is the visualization of structural organization within symmetric second order tensor fields. In this respect, it can be stated that there exist certain aspects that previous approaches do not solve satisfactorily: in classical tensor field visualization (i.e. direct visualizations of scalar measures derived from tensors, or characteristic line visualization), the structure of a tensor field can only be implicitly deduced. Methods from the area of feature extraction are mostly based on local criteria (i.e. based on some kind of local gradient measure), which only capture local and distinct variation but cannot extract further structural information from regions of coherent behavior. Approaches from tensor field topology are particularly sensitive to noise, which makes them suffer from a lack of robustness. Tractographical methods successfully extract structural information from diffusion tensor fields, but require to first extract a great number of fiber trajectories and then to perform fiber similarity comparisons.

The concept of Lagrangian coherent structures (LCS) represents an alternative to vector field topology that is also well applicable and interpretable for transient vector fields because it is based on true advection (pathlines) instead of instantaneous streamlines. Coherent structures have been roughly defined as regions of coherent motion or where a flow variable exhibits significant correlation with itself or another variable [38]. More recently, Haller has defined LCS to be a time-dependent analog of separatrices [17], [15] (hence separating regions of coherent motion) and has shown that they can be obtained using the so-called finite-time Lyapunov exponent (FTLE). This definition of LCS is getting increasingly popular and is also the basis for our work. A detailed analysis of LCS in n-dimensional systems is presented by Lekien et al. [32]. In visualization, Garth et al. [13] introduce direct visualization of the FTLE whereas Sadlo and Peikert [41] extract height ridges from FTLE and compare them to vector field topology for the case of stationary vector fields. An advantage of LCS is their robustness under the effects of noise, as analyzed by Haller [16] for approximate velocity data.

In the unpublished workshop talk by Tricoche et al. [47], the concept of a structural coherence measure for DT-MRI data based on neighboring fiber pathways is sketched. A bidirectional flow map is roughly described for which neighboring fibers must be oriented consistently. The maximum spectral norm of the two Jacobians for both directions saved in the flow map is suggested for quantifying the coherence of the fibers. Our paper shares the same basic ideas; it is based on the thesis of one of us (Hlawatsch) [19], in which the application of FTLE to diffusion tensor fields is analyzed in detail. Simultaneously to our work, similar research was independently done by Hlawitschka in his thesis [20]. His local coherence measure was later published, after submission of this paper, as a journal paper [21], detailing a brief workshop abstract [22]. The coherence measure is based on the spectral norm of a derivative of fiber trajectories passing through a small region, a locally consistent orientation of these fibers is also necessary. Additionally, a GPU implementation is described to speed up the computation of this coherence measure. Only the application to diffusion tensor fields is presented, whereas we apply our concept also to stress tensor fields in this paper. We further present a formal definition of an adapted gradient tensor for a consistent treatment of 1-to-m mappings (see Section IV).

III. COHERENT STRUCTURES IN VECTOR FIELDS

In the following, a short introduction to the theoretical basis of coherent structure visualization and the Lyapunov exponent shall be given. The *Lyapunov exponent* (LE) is a concept from dynamical systems theory that measures the exponential rate of divergence of two close orbits (trajectories in phase space) of the system. An alternative interpretation of the LE is that it measures the exponential growth of an infinitesimal perturbation. For an n-dimensional dynamical system, there exists a spectrum of n LEs, of which the largest, $\sigma_1(\mathbf{x})$, is defined as

$$\sigma_1(\mathbf{x}) = \lim_{T \to \infty} \lim_{\|\boldsymbol{\delta}(t_0)\| \to 0} \frac{1}{|T|} \ln \frac{\|\boldsymbol{\delta}(t_0 + T)\|}{\|\boldsymbol{\delta}(t_0)\|}, \qquad (1)$$

where $\boldsymbol{\delta}(t)$ is the perturbation at time t, having originated at position $\mathbf{x}$ and time t_0. Since this analysis considers trajectories, it is referred to as Lagrangian to emphasize the underlying transport process.

For the computation of the Lyapunov exponent, Haller [15] uses an intermediate computation step, the *flow map* $\boldsymbol{\xi}_{t_0}^T(\mathbf{x})$: $\mathbb{R}^n \to \mathbb{R}^n$, which is a mapping from the initial position $\mathbf{x}$ at time t_0 of a massless particle in a transient vector field to its final position after advection for time T. In practice, the flow

map is typically obtained by integration of trajectories seeded on a regular grid. The direction of maximum expansion in the domain of $\boldsymbol{\xi}_{t_0}^T$ corresponds to the major eigenvector of the right Cauchy-Green deformation tensor

$$\mathbf{C}_{t_0}^T(\mathbf{x}) = (\nabla \boldsymbol{\xi}_{t_0}^T(\mathbf{x}))^\top \nabla \boldsymbol{\xi}_{t_0}^T(\mathbf{x}).$$

The maximum stretching factor can be obtained as the spectral norm of $\nabla \boldsymbol{\xi}_{t_0}^T(\mathbf{x})$, defined as the square root of the major eigenvalue $\lambda_{\max}$ of $\mathbf{C}_{t_0}^T(\mathbf{x})$.

When computing the LE for practical data, one issue arises: integration of infinite length is usually not applicable. Furthermore, the spatio-temporal domain of real world data is often finite. This led to the idea of replacing the infinite time range in Equation 1 by a finite one.

Called direct Lyapunov exponent by Haller [15], it is nowadays commonly denoted *finite-time Lyapunov exponent*. The maximum FTLE, subsequently simply referred to as FTLE, at position $\mathbf{x}$, time t_0, with advection time T, is computed as:

$$\sigma_{t_0}^T(\mathbf{x}) = \frac{1}{|T|} \ln \sqrt{\lambda_{\max}(\mathbf{C}_{t_0}^T(\mathbf{x}))}. \tag{2}$$

Haller has shown that it allows one to infer the existence and position of finite-time repelling and attracting manifolds in a transient system: the *Lagrangian coherent structures*. These structures are immanent in the form of ridges of the scalar field defined by the FTLE [15].

IV. GENERALIZED LAGRANGIAN COHERENT STRUCTURES

In this section, we generalize the concept of map-based extraction and visualization of Lagrangian coherent structures. Due to the fact that the computation of the FTLE is solely based on the flow map, we propose the following generalization: instead of restricting the derivation of the map to a vector field or dynamical system, it should be possible to extract and visualize coherent structures using the right Cauchy-Green deformation tensor from any map, irrespectively of its origin. To account for this, the general term *coordinate map* shall be used in this paper. Regarding the gradient computation on the coordinate map, the only requirement is that the map is spatially C^1 continuous, except for isolated points or null sets that may be excluded from the analysis. In the case of discrete coordinate maps, it is beneficial to require a regular sampling to account for the sampling dependency of the FTLE and to prevent artifacts from gradient estimation.

Our generalization also addresses mechanisms where a starting point corresponds to more than one end point. This is the case for undirected correspondences and for eigenvector fields, where two possible end points exist. This generalization is even more important when applying it to data from medical imaging methods with high angular resolution [49], [45], [10], where characteristic lines could be extracted e.g. with methods by Schultz and Seidel [43] or by Hlawitschka et al. [23]. There, in contrast to diffusion tensors, crossing fibers can be represented with multiple discriminable directions, which may lead to more than two end points for a given starting point.

To simplify the computation of LCS, gradient computations of such 1-to-m correspondences are avoided. Instead, these situations are modeled by a set of m maps. The generalized

map that maps a point $\mathbf{x}$ at time t_0 to one of its m possible positions at time $t_0 + T$ shall be defined as $\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})$, with $k \in \{k_1, \ldots, k_m\}$. If m equals to 1, k is omitted for clarity. Care has to be taken if local operators are applied to $\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})$, since the choice of equal k at different positions within the support of the operator may not ensure consistent selection of corresponding mappings. The consistency of mappings must be established through an appropriate choice of k.

Since the gradient of the coordinate map is the basis for the extraction of LCS, it shall be derived here. Other local operators can be derived similarly. The adapted gradient tensor operator $\tilde{\nabla}$ of the coordinate map is formulated as the following limit process

$$\begin{aligned}\left(\tilde{\nabla}\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})\right)_{ij} &= \frac{\tilde{\partial}\left(\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})\right)_i}{\partial x_j} \\ &= \lim_{h \to 0} \frac{\Delta_{j,h}\left(\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})\right)_i}{h}\end{aligned} \tag{3}$$

using index notation, with

$$\Delta_{j,h}\left(\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})\right)_i = \left(\boldsymbol{\xi}_{t_0}^{\eta,T}(\mathbf{x}+h\mathbf{e}_j)\right)_i - \left(\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x})\right)_i \tag{4}$$

where $\mathbf{e}_j$ is the j-th basis vector and η ensures that the previously formulated continuity requirement of the coordinate map is satisfied by a selection of the mapped position at location $\mathbf{x} + h\mathbf{e}_j$ such that the limit in Equation 3 exists. An explicit choice of η for the case of coordinate maps extracted from eigenvector fields of symmetric tensors is given in Section V-C.

The parameters representing time require further consideration. For datasets that consist of a single time step (stationary data), t_0 is constant and is therefore omitted in the notation. Also, there are cases where the underlying mechanism does not represent a temporal evolution. If the map is the product of a tracing procedure, T can be interpreted as "tracing time". Otherwise T is also omitted in the notation. The normalization by tracing time T in Equation 2 was originally motivated by the growth rate of a perturbation. If T is uniform over the domain, normalization only results in a scaling and hence does not change resulting features (ridges). On the other hand, if trajectories are stopped earlier, e.g. because they leave the domain, this normalization can change features. Furthermore, many maps may not reflect a temporal process. Therefore, we propose to generally leave this step out and to decide on its application depending on the respective use case.

The logarithm was introduced in Equation 2 because perturbations grow exponentially in linear vector fields or dynamical systems and the resulting measure captures the underlying linear behavior. However, since our generalization allows for the analysis of arbitrary maps, a logarithmic measure should generally not be enforced. Furthermore, since the logarithm is a monotonic function, it cannot change existing features (i.e. ridges), it only changes the level at which they are expressed. Therefore, this step is left optional and should only be applied if the goal of the analysis is the measurement of exponential growth. If the resulting data is visualized directly, it can still be beneficial to compute the logarithm in terms of a transfer

function for better exploitation of the available color map range.

All in all, the right Cauchy-Green deformation tensor of the generic coordinate map is

$$\mathbf{C}_{t_0}^{k,T}(\mathbf{x}) = (\tilde{\nabla}\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x}))^{\top}\tilde{\nabla}\boldsymbol{\xi}_{t_0}^{k,T}(\mathbf{x}). \qquad (5)$$

From this we define the generalized counterpart to the FTLE as

$$\tilde{\sigma}_{t_0}^{T}(\mathbf{x}) = \max_{k\in\{k_1,\ldots,k_m\}}\left(\sqrt{\lambda_{\max}(\mathbf{C}_{t_0}^{k,T}(\mathbf{x}))}\right), \qquad (6)$$

which we call the maximum *finite separation ratio*, short FSR in the remainder of this paper, for which the following equality holds: $\sigma_{t_0}^{T} = 1/|T| \cdot \ln \tilde{\sigma}_{t_0}^{T}$, for the special case $m = 1$. The formulation of Equation 6 contains a selection of the maximum separation over all k, in analogy to the previous definition of the Lyapunov exponent as a measure of maximum separation. Furthermore, we omit the logarithm, which has benefits in many applications where exponential growth has no significance. It is a direct measure of the maximum change of an initial distance under the considered mechanism. For the example of cerebral anatomy, this allows for a direct measurement of the maximum distance of two axonal fibers after tracing from proximate starting points in relation to their initial distance.

Also conceivable, but not applied in this work, is an FSR maximum in analogy to the FTLE maximum according to Sadlo and Peikert [40], which measures the maximum separation along an entire trajectory instead of measuring it from the final flow map for a fixed advection time T. This is equivalent to the maximum FSR over a range of advection times T, and can easily be modeled with our method by including the maps for different advection times in the set of k from which the definition of Equation 6 already selects the maximum.

The generalized counterpart to LCS are the ridges in the FSR, which we propose to call *separatrices of coherent regions* (SCR). In this new nomenclature we drop the term "Lagrangian" because our maps do not necessarily describe transport phenomena and we introduce the term "separatrix" in order to disambiguate the term LCS with respect to coherent regions and the structures that separate them from each other.

One reason why LCS got widely accepted in the last years is that they give an appropriate view to transient vector fields: as opposed to vector field topology, they allow an intuitive interpretation with respect to transient advection and come in the form of time series (of varying t_0) that allow for this interpretation at each instant of time. One question is to what extent these properties of LCS transfer to the concept of SCR. If the time steps of the data represent a process that is temporally consistent with the progression of the seeding moment t_0 and advection time T, the coordinate map can be extracted in a time-dependent way, e.g. by tracing pathlines in the case of vector fields. In this case, SCR analyze transient advection and hence lead to a transient view of the problem in question, just as LCS. However, if the mechanism represented by the coordinate maps is instantaneous, if T is not consistent with t_0 of different coordinate maps, or if the scope of the analysis is purely structural, the coordinate map should be extracted in an instantaneous manner, e.g. by tracing

streamlines in the case of vector fields. The SCR for each t_0 have then to be primarily examined on their own.

V. Coherent Structures in Symmetric Tensor Fields

We apply the computation of the FSR to symmetric second order tensor fields, in particular to diffusion tensor fields from DT-MRI and to stress tensor fields. Beside the definition of integral curves, which are required to build the coordinate map, the orientation ambiguity of eigenvectors in tensor fields has to be considered.

A. Diffusion Tensor Fields

One of the most prominent sources of symmetric second order tensor fields are medical diffusion tensor images. Since these describe the dominant direction of diffusion of water molecules in tissue, it is reasonable to base a structural analysis of such data on sets of characteristic lines, in analogy to stream- and pathlines in vector fields. In medical applications, such as tractography, it is common to use integral curves of the tensor field's major eigenvector. We will denote these as eigenvector lines in the rest of this paper. Because the major eigenvector is only significant in anisotropic regions, these lines are commonly stopped when leaving such areas. In our examples, we use these lines for the computation of the FSR from diffusion tensors and use fractional anisotropy [1] as stopping criterion. However, our concept is independent of these choices, other definitions for eigenvector-based lines are possible just as well, e.g. tensorlines introduced by Weinstein et al. [53].

In medical data, the geometric structure of tissue is likely to be independent of the actual rate of diffusivity, which is why the major eigenvalue is not incorporated in the propagation of a eigenvector line. For cases where the absolute rate of diffusivity contributes to the understanding of a diffusion tensor field, and for applications from general physics, we optionally include a step size control that scales the obtained direction vector with a factor $\delta = \lambda_1/\lambda_{\max}$, where λ_1 is the local major eigenvalue and $\lambda_{\max}$ is the maximum major eigenvalue throughout the entire dataset. Regarding neighboring eigenvector lines, this variation can create shearing (differing propagation velocities along eigenvector lines), provoking separation of their end points, which becomes visible in the FSR. Section VII-A includes an example demonstrating the utility of this variation.

B. Stress Tensor Fields

Another class of symmetric second order tensors are tensors describing mechanical stress [33]. In contrast to diffusion tensors, stress tensors can exhibit negative eigenvalues. Their eigenvalues correspond to the principal stress magnitudes and their eigenvectors to the principal stress directions. The sign of the eigenvalue classifies the stress as tension (positive eigenvalue) or compression (negative eigenvalue). In contrast to diffusion tensors, a classification into different types of anisotropy (linear, planar, spherical) is often not useful and
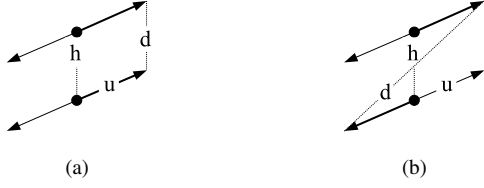
Fig. 1. Accounting for the ambiguity of eigenvector orientation for differential operators. In (a) and its symmetric counterpart, $\lim_{h \to 0} d/h = 0$ exists, whereas in (b) this is not the case, requiring local orientation of the eigenvectors.

there is typically no case of linear anisotropy, where only a single direction along the major eigenvector would be sufficient. The analysis of stress tensor fields can be based on lines following the principal stress directions as described by Dick et al. [7]. We extract coordinate maps from these lines and use them for the calculation of the FSR for all principal stress directions, resulting in three different FSR fields (Section VIII).

C. Orientation Ambiguity in Tensor Fields

Because of the orientation ambiguity of eigenvectors, there are for each starting position two end points of eigenvector-based lines. Therefore, the resulting mapping can be represented by two coordinate maps $\boldsymbol{\xi}^{k,T}(\mathbf{x})$ ($k \in \{-1, +1\}$) (Section IV), where k selects the final position according to the initial propagation direction $\mathbf{v}_0(\mathbf{x}) = k\boldsymbol{\epsilon}(\mathbf{x})$ depending on the eigenvector $\boldsymbol{\epsilon}(\mathbf{x})$. Applying Haller's approach of FTLE computation [15] to these coordinate maps directly would be inappropriate, since misaligned (opposing) starting directions can lead to overestimated separation of trajectories as well as discontinuities. Therefore, local orientation (formally motivated in Figure 1) is decided using

$$\eta = \mathrm{sgn}_1(\mathbf{v}_0(\mathbf{x}) \cdot \boldsymbol{\epsilon}(\mathbf{x} + h\mathbf{e}_j)),$$

with $\mathrm{sgn}_1(\mathbf{x} \neq 0) = \mathrm{sgn}(\mathbf{x})$ and $\mathrm{sgn}_1(0) = 1$. This definition of η makes sure that trajectories consistent with $\mathbf{v}_0(\mathbf{x})$ are chosen for the evaluation of $\tilde{\nabla}\boldsymbol{\xi}^{k,T}(\mathbf{x})$ (Equations 3–5). Similarly, local alignment of eigenvectors is also performed during line tracing. Following Equation 6, we define the FSR for symmetric tensor fields as

$$\tilde{\sigma}^T = \max\left(\sqrt{\lambda_{\max}(\mathbf{C}^{+1,T})}, \sqrt{\lambda_{\max}(\mathbf{C}^{-1,T})}\right).$$

VI. FSR Visualization

As noted by Haller [15], LCS correspond to ridges in the FTLE field. The same relation holds for SCR in the FSR field and such ridges can be obtained with appropriate ridge extraction algorithms [9], [11]. However, due to the noise sensitivity of these approaches, we preferably do not extract the geometry of ridges in 3D FSR scalar field but predominantly perform direct cross-sectional visualizations with a linear gray scale color map, where low FSR values are mapped to black and high values mapped to white. Regions where no FSR was computed because the stopping criterion (Section V-A) was already met at the starting position are

marked turquoise. Additionally, if FSR values cover a high dynamic range, one can resort to gray scale color mapping of FSR values on a logarithmic scale. For the results of the stress tensor field, high values are mapped to red or green, to allow for a combined visualization of two FSR fields. As previously shown by Garth et al. [13], a direct LCS visualization without explicit ridge extraction is often sufficient to faithfully identify coherent regions and their separating structures.

VII. Results for Diffusion Tensor Fields

In the following, the properties of our proposed method for the visualization of SCR in diffusion tensor fields shall be explored with a number of both synthetic and measured datasets and compared to other visualization techniques for diffusion tensor fields. The datasets are:

1) Alternating diffusivity: a synthetic $64 \times 64 \times 64$ tensor field consisting of unidirectional linear anisotropy with total diffusivity alternating by a factor of two every four slices. All tensors are thus linearly anisotropic and differ only in scale (see Figures 2 and 3).
2) Synthetic fiber bundle: a synthetic $64 \times 64 \times 64$ tensor field that models a simplified bundle of fibers that separates at one side (see Figures 2 and 4).
3) Canine heart: a measured $256 \times 256 \times 134$ DT-MRI image of an excised canine heart (see Figure 5).
4) Human brain: a measured $148 \times 190 \times 160$ DT-MRI image of a human brain (see Figure 6). Prior to further processing, we have smoothed the dataset with a Gaussian filter width σ of two thirds of the voxel size.

All FSR fields were computed in fourfold native dataset resolution (except Figure 11, left), normalized with the 95^{th} percentile, and visualized with a linear gray scale color map.

A. Comparison to Other DT-MRI Visualization Methods

Due to its dependence on the major eigenvector and thus on linear anisotropy, our approach is compared to established methods for analysis of linear anisotropy in diffusion tensor fields. In detail, these methods are:

1) Fractional anisotropy (FA) as introduced by Basser and Pierpaoli [1].
2) The gradient magnitude of FA, $\|\nabla\text{FA}\|$.
3) A ridge criterion based on fractional anisotropy, $|\nabla\text{FA} \cdot \boldsymbol{\epsilon}_3|^{1/3}$, as proposed by Kindlmann et al. [26], where $\boldsymbol{\epsilon}_3$ is the minor eigenvector of the Hessian of FA.

All methods used the same resolution as the FSR and gradients were computed with central differences. Drawing analogies from previous results in the visualization of LCS in vector fields, our extensions for tensor fields can be expected to be capable of highlighting or discriminating the following features in tensor fields:

- Regions with locally deviating major eigenvector directions.
- Regions of locally varying diffusivity (if step size control based on the local eigenvalue is used).
- Structuring of locally coherent regions according to semi-global divergence of trajectories (i.e. separation in the course of a fiber bundle).
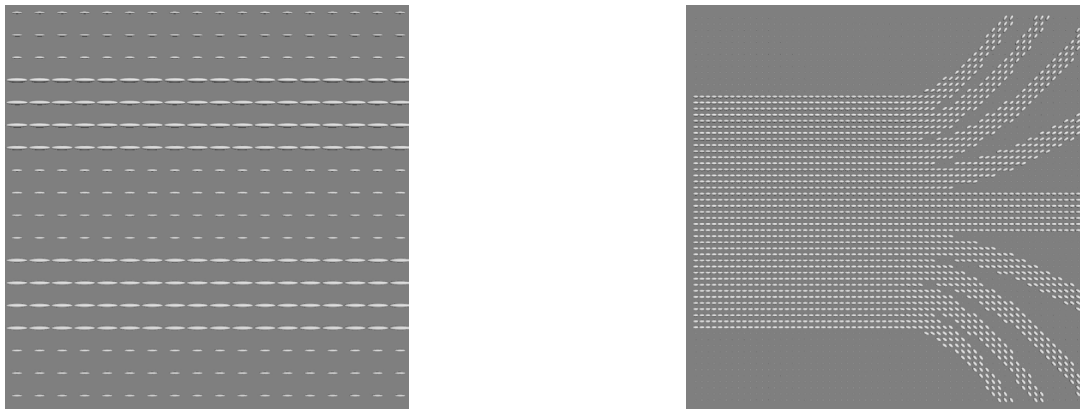
Fig. 2. Left: tensor glyph visualization of the alternating diffusivities dataset. Right: tensor glyph visualization of the synthetic fiber bundle dataset.

Section V-A introduced a variation of the eigenvector line tracing model that adds sensitivity to tensor diffusivity by performing propagation steps proportional to the major eigenvalue for the extraction of eigenvector lines. Figure 3 demonstrates results obtained by applying this variation to the alternating diffusivities dataset. It is apparent that this variation makes regions of changing diffusivity visible in the FSR field. Since the change in total diffusivity does not influence the respective ratios of eigenvalues, this feature cannot be detected with any visualization method solely derived from fractional anisotropy.

Figure 4 presents results obtained with the synthetic fiber bundle dataset. All methods delimit the fiber bundle against the background. Separating line trajectories form ridge features in the FSR, as expected. Additionally, features visible in the FSR appear to not only coincide with structures visible in the other methods, but to form a superset of features that grows with increasing integration length. The possibility to perform such a semi-global analysis is one of the major advantages of the FSR, since it allows one to structure coherent regions in accordance to distant separation in the progression of trajectories (see also Section VII-B).

In Figure 5 of the canine heart dataset and Figure 6 of the human brain, the FSR distinguish coherent regions of uniform fiber orientation (separated by ridges of high gray scale value). The myocardium of the mammalian heart consists of several layers of helical muscle fiber tracts. The angles of these fibers are changing from the epicardium to the endocardium (e.g. described and visualized in [44]). This layer structure of the myocardium is apparent in the FSR field of the heart dataset. Thin ridges appear where the angle of neighboring fiber layers is changing. The different angles of the layers are also visible in the right image of Figure 12, where eigenvector lines in the heart dataset are directly visualized. In the brain dataset, coherent structures are visible in the FSR field inside the white matter areas of the brain. The corpus callosum (with the lateral ventricles appearing in the center of the shown slice) consists of a huge number of fibers connecting the two cerebral hemispheres and is clearly visible as a region of coherent fiber structure in the FSR. Coherent fiber structures in the corona radiata and the internal capsule are visible in the FSR as well. To some extent, these structures are likewise visible with the FA-based approaches, but altogether, the FSR

field tends to represent these structures in a more coherent and robust manner. An overlay of an RGB encoding of the major eigenvector with the FSR field (Figure 7) also shows that abrupt changes of the major eigenvector direction correlate with distinct structures in the FSR field.

B. Varying Integration Length

A major benefit of our approach is that it allows for a gradual transition from local to global analysis by increasing the length of the integral curves. In doing so, propagation length effectively represents a scale space parameter that determines the scale at which separation of trajectories is analyzed. The effect is illustrated in Figure 8 with the canine heart dataset, where structures in the FSR grow and noise is suppressed with increasing line length. This effect of growing structures in the FSR also becomes apparent when comparing the respective first images of Figures 4 and 9, which were obtained with different propagation lengths.

C. Susceptibility to Noise

As already indicated in Figure 8, a semi-global analysis by means of increasing the propagation length helps reducing the impact of noise, as discussed by Haller in the context of vector fields [16]. Here, we leave the propagation length fixed and demonstrate the effect of different levels of noise to the FSR field of the synthetic fiber bundle dataset compared to the gradient magnitude of fractional anisotropy in Figure 9. As a semi-global measure, one would expect the FSR to degrade strongly due to accumulation of error along trajectories. However, while the quality of both FSR and $||\nabla\text{FA}||$ does degrade, the overall structure within the FSR remains comparatively intact. Obviously, these results can only give an initial qualitative impression. A detailed quantitative analysis is subject of future work.

VIII. RESULTS FOR STRESS TENSOR FIELDS

We applied the FSR to medical stress tensor fields by using integral curves following the principal stress directions (see Section V-B). The presented results base on a stress tensor field at a resolution of $86 \times 81 \times 226$ that resulted from
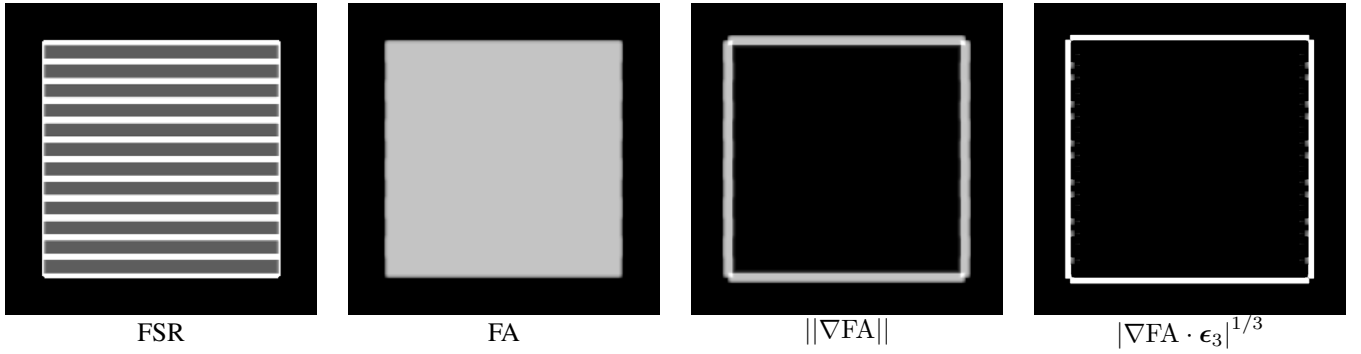
Fig. 3. A comparison of the finite separation ratio, fractional anisotropy, the gradient magnitude of fractional anisotropy, and the fractional anisotropy ridge criterion applied to the synthetic alternating diffusivity dataset. The FSR computed with eigenvalue-based step size control discriminates regions with differing total diffusivity, which cannot be extracted with the other methods.
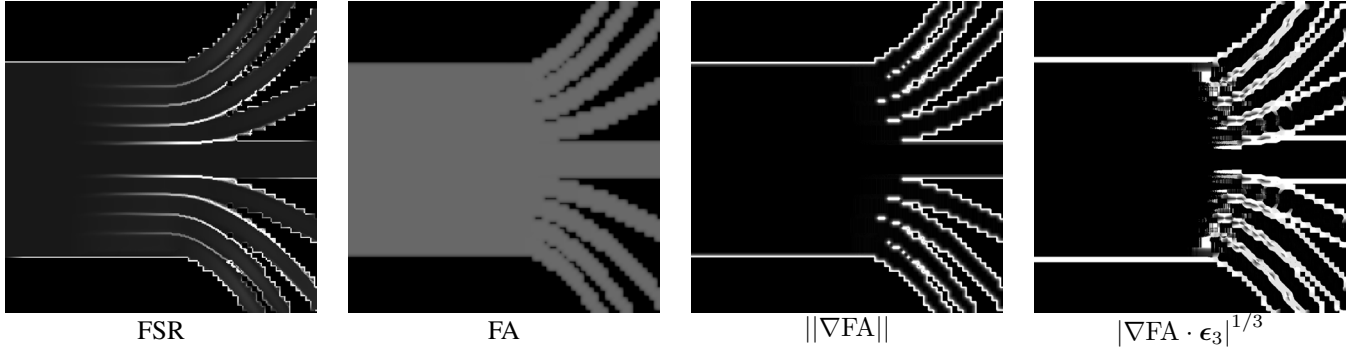

Fig. 4. A comparison of the finite separation ratio, fractional anisotropy, the gradient magnitude of fractional anisotropy, and the fractional anisotropy ridge criterion applied to the synthetic fiber bundle dataset. Features visible in the FSR form a superset of those visible in the gradient magnitude of FA and the FA ridge criterion.

simulation of a human femur under load; the tensor field is only defined inside the bone. A similar dataset is used in [7], hence we were able to confirm the resulting FSR structures. Figure 10 shows the results for the interesting area in the upper part of the bone. The upper left image shows a combined visualization of two FSR fields: green colored is the FSR for lines that follow the principal stress with highest positive eigenvalue, which corresponds to tension, red colored is the FSR for lines that follow the principal stress with highest negative eigenvalue, which corresponds to compression. The two images in the lower row, where lines along one principal stress direction are shown together with the corresponding FSR, confirm that ridges in the FSR field show the separation of the lines following the principal stress direction. In the presented images, lines can still visually cross the ridges in the FSR field or leave the dataset domain because a 2D slice of the FSR field was combined with 3D lines in this visualization. The structures indicated by the FSR field can be validated according to medical literature as shown in the illustration on the upper right. The typical distribution of principal stress directions in a human femur can be recognized in the FSR field. For example, a ridge in the tension FSR field can be found where the green lines separate on the left side in the illustration.

These results show that the FSR can help interpreting stress tensor fields. Depending on the length of the underlying lines, the FSR can be seen as a semi-global measure for coherent principal stress regions, with low FSR values indicating areas of coherent stress directions and high FSR values separating different regions of coherent stress.

IX. IMPLEMENTATION

In practice, LCS in vector fields are visualized by defining a grid of seeding positions for the computation of the flow map. Since a full 3D flow map may easily require the costly computations of hundred thousands or even millions of stream- or pathlines, adaptive acceleration methods have been proposed [12], [40]. With tensor fields, the issue of computational complexity is even aggravated. Firstly, because the coordinate map for diffusion tensor fields requires the computation of two integral curves at each of its nodes (one parallel to the local major eigenvector and another antiparallel to it). In the case of stress tensor fields, the number of computed curves per node even increases to six if the FSR is computed for all three principal stresses. Secondly, because for each of its integration steps, the propagation of these lines requires the local solution of the eigensystem of the tensor field. To make the analysis of tensor fields feasible at all, we accelerate our method using modern programmable graphics hardware inspired by the approach of Kondratieva et al. [29]. The FSR computation and visualization is implemented in C++ using OpenGL and its shading language. As in other general purpose GPU computations, we make use of the highly parallel processing in the fragment stage of modern graphics cards

FSR

FA

$||\nabla\text{FA}||$

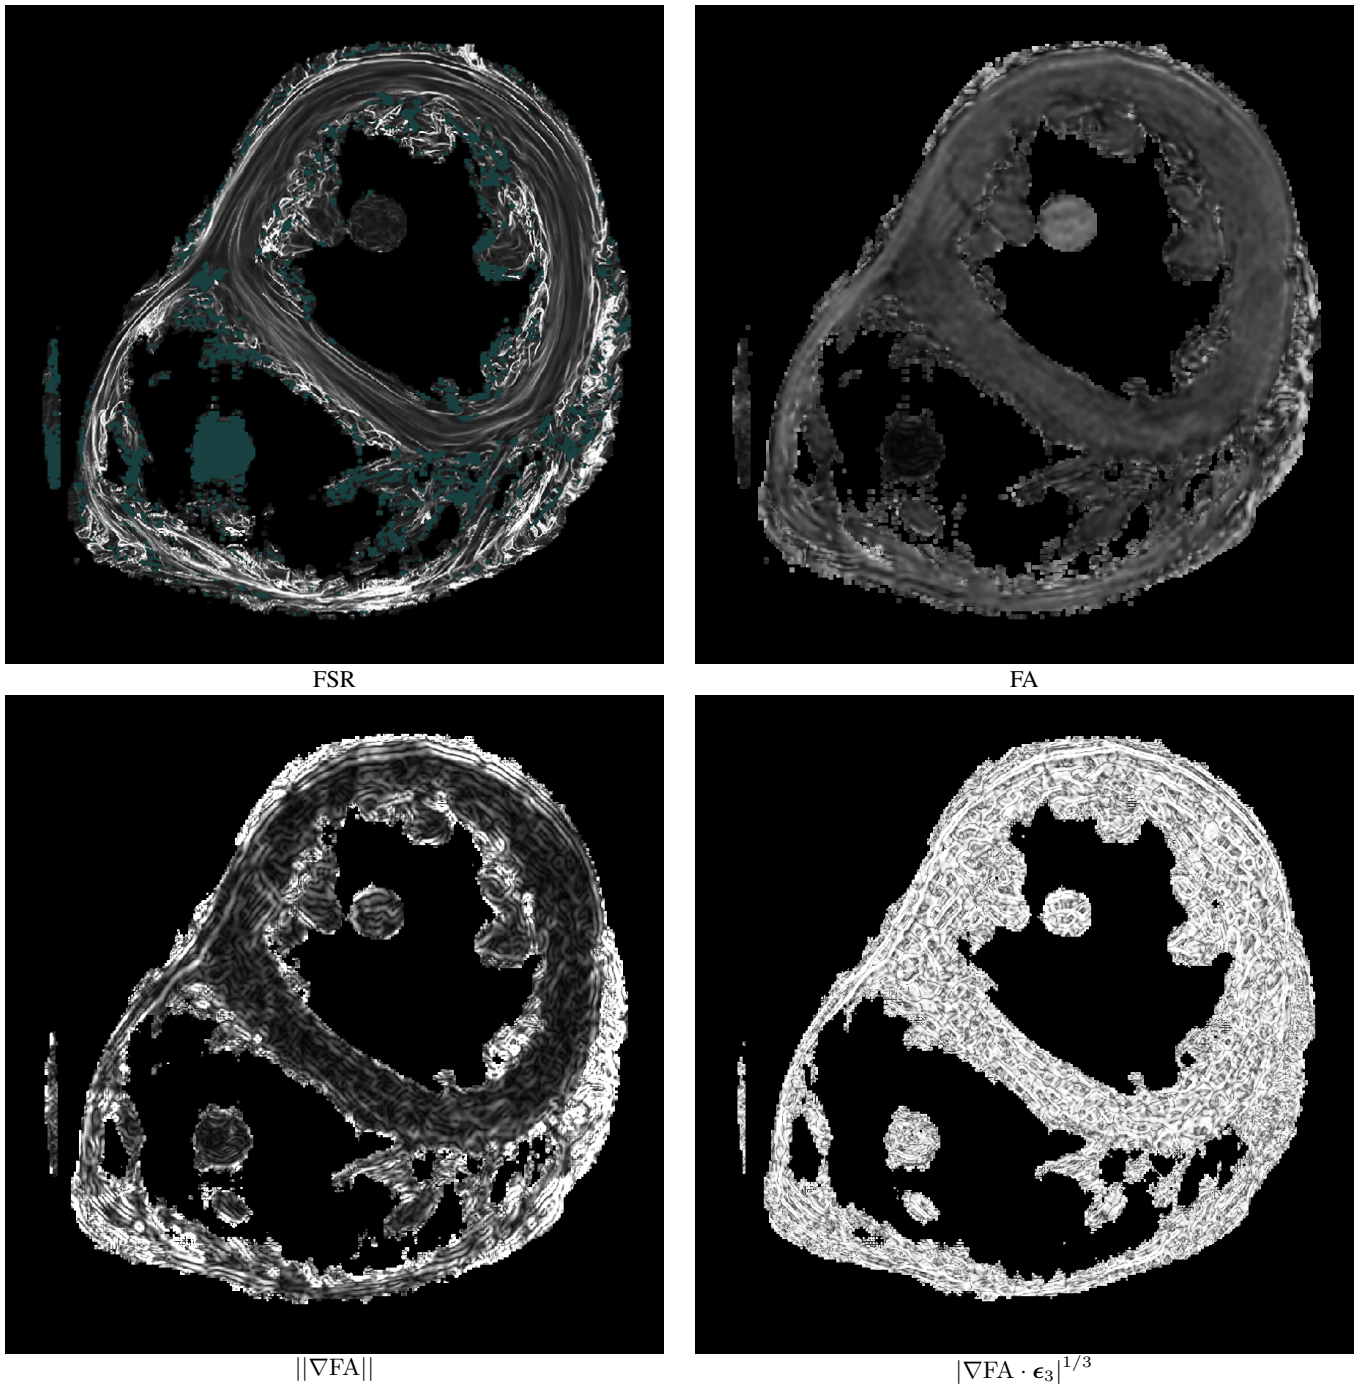$|\nabla\text{FA}\cdot\boldsymbol{\epsilon}_3|^{1/3}$

Fig. 5. A comparison of the finite separation ratio, fractional anisotropy, the gradient magnitude of fractional anisotropy, and the fractional anisotropy ridge criterion applied to the measured canine heart dataset. Although computed at the same resolution as the other approaches, the FSR reveals extensive coherent structures that are far less sensitive to noise.

by assigning rasterized fragments to the seeding positions of eigenvector lines, respectively stress lines, and propagating them in shader loops. Two 3D textures are used to store the six relevant entries of symmetric tensors and to efficiently retrieve interpolated samples of the tensor field at any position within its domain. An analytical algorithm [18] is used for computing the eigensystem of the tensors on the GPU. Similarly, both the gradient of the coordinate map and the FSR are computed with the GPU for subsequent visualization. Comparing an unoptimized single-threaded CPU reference implementation of

FSR computation to our GPU implementation, we observed an acceleration factor of around 1000. A thorough performance comparison is beyond the scope of this work. However, due to the highly parallel organization of GPUs, an average acceleration factor of 10 to 100 can be regarded as realistic.

Despite its benefits, there are also downsides to a GPU-based implementation of the FSR computation. First, the hardware accelerated interpolation of current GPUs utilizes reduced floating point precision. A rough analysis of the computational error involved with our GPU implementation

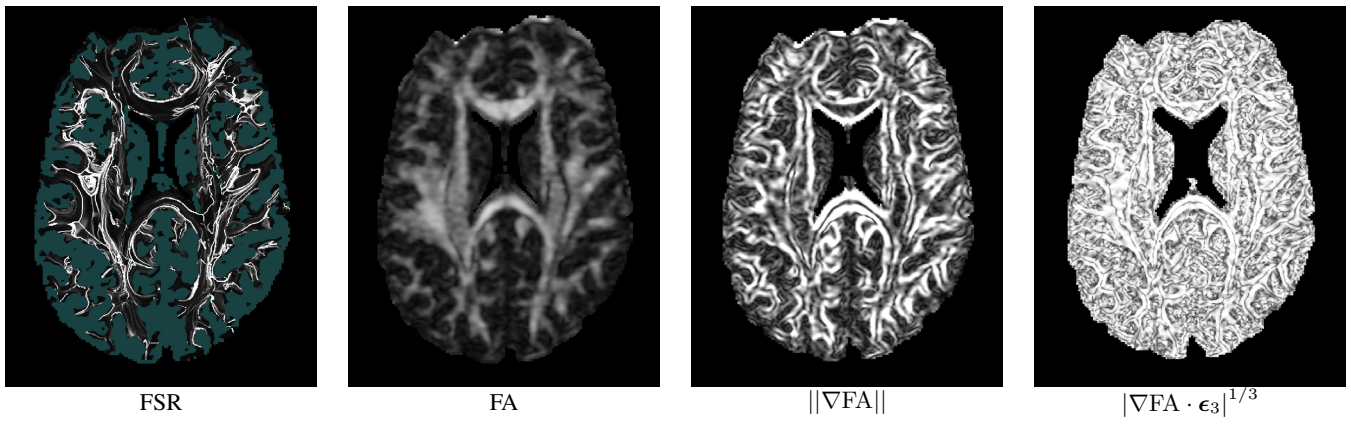|			
FSR	FA	$\|\nabla\mathrm{FA}\|$	$\|\nabla\mathrm{FA}\cdot\boldsymbol{\epsilon}_3\|^{1/3}$

Fig. 6. A comparison of the finite separation ratio, fractional anisotropy, the gradient magnitude of fractional anisotropy, and the fractional anisotropy ridge criterion applied to the measured human brain dataset. Structures visible with the other methods coincide in parts with those in the FSR. Other structures are likewise visible, due to the semi-global nature of the FSR.
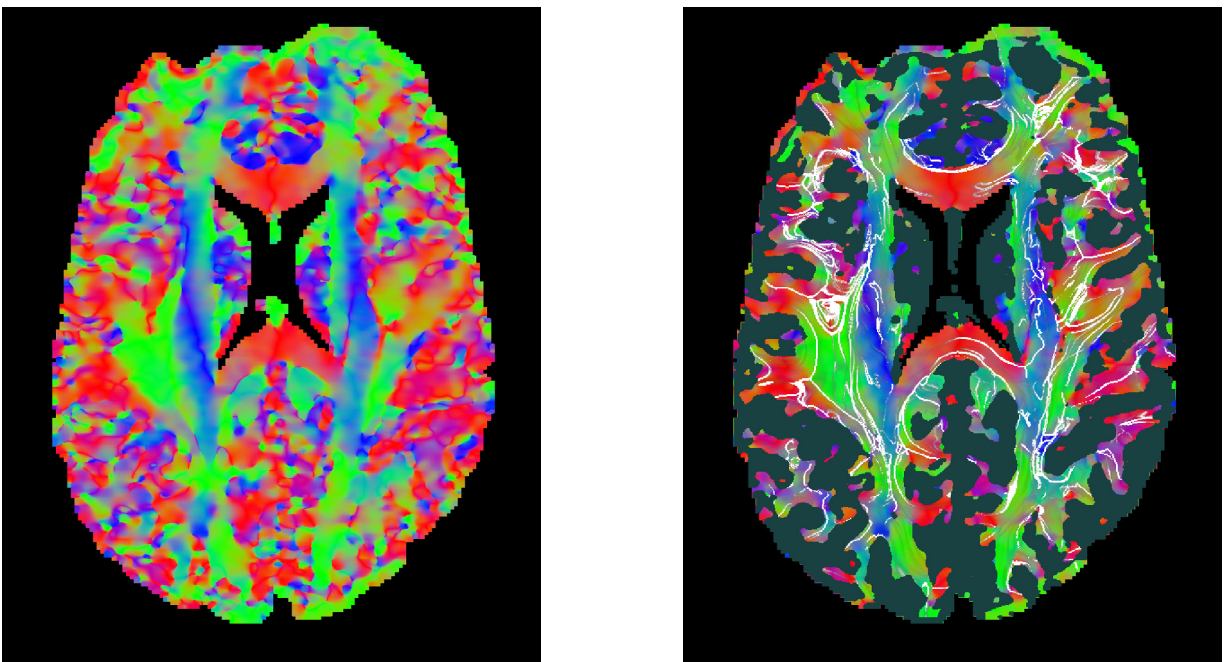

Fig. 7. Visualizations of the measured human brain dataset. Left: the major eigenvector direction mapped to RGB color values. Right: an overlay with the computed FSR field. Visible structures correlate with abrupt changes in major eigenvector direction.

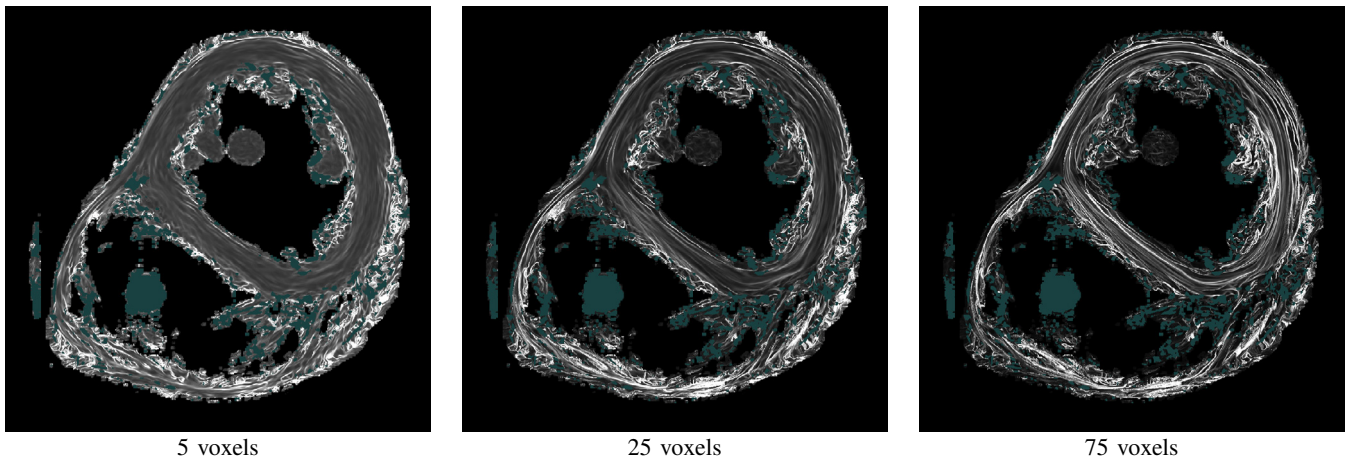|		
5 voxels	25 voxels	75 voxels

Fig. 8. Demonstration of the effect of increasing propagation length (measured in multiples of the voxel edge length) on the finite separation ratio using the measured canine heart dataset. With increasing propagation length, structures become more coherent and noise artifacts are averaged.

FSR, $A = 0.2\%$ FSR, $A = 2\%$ FSR, $A = 20\%$

$||\nabla\text{FA}||$, $A = 0.2\%$ $||\nabla\text{FA}||$, $A = 2\%$ $||\nabla\text{FA}||$, $A = 20\%$
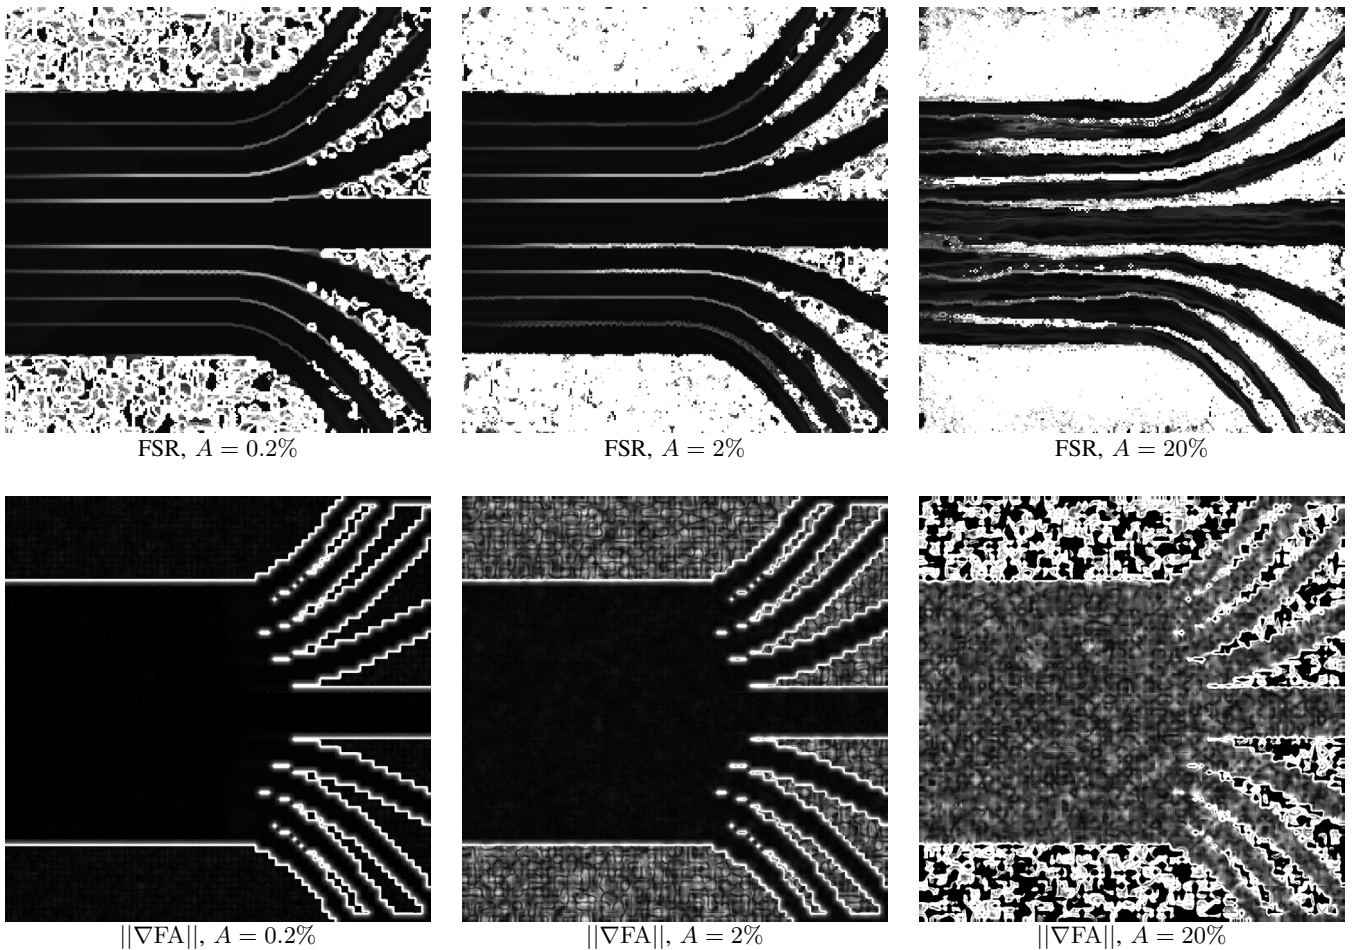
Fig. 9. Comparing the FSR and the gradient magnitude of fractional anisotropy under the effect of noise. The images are labeled with the maximum amplitude A (relative to the maximum tensor component throughout the entire dataset) of white noise with uniform distribution that has been added to all entries of the tensor. The quality of structures visible with both approaches degrades with increasing level of noise. However, even though the FSR is a semi-global measure that might potentially accumulate errors, the overall structure persists even at high noise levels.

indicates that the median of the deviation of end points of particle traces is typically at about 0.005% and the median of the deviation of the FSR at about 0.01%. Considering the speedup of three orders of magnitude obtained by our GPU implementation, this deviation may be regarded tolerable. The second limitation of the GPU-based method is that storage requirements for the tensor field, the coordinate map, the resulting FSR field, and intermediate buffers may exceed available texture memory. We follow two strategies to deal with this storage issue. The first strategy is to compute the FSR without storing the coordinate map in texture memory, since it is a mere intermediate computation step in the process of FSR computation. Instead, for each position at which the FSR should be evaluated, the two local gradient tensors $\tilde{\nabla}\boldsymbol{\xi}^{k,T}$ (see Section V-C) of the coordinate map are computed on the fly, from which the FSR is calculated directly. This is done in the FSR fragment shader by evaluating six integral curves at a time for $k = +1$ and $k = -1$ respectively, at the cost of repeated evaluations of identical curves between neighboring fragments. Thus, for each FSR value, twelve instead of two integral curves must be evaluated. Despite its obvious inefficiency, this approach performs for example in roughly 21

seconds on an NVIDIA 8800 GPU for a $256 \times 256 \times 134$ diffusion tensor field sampled at this native resolution with 20 propagation steps for each eigenvector line (about 100 million eigenvector lines in total), which we still consider sufficiently fast for explorative visualization. Furthermore, by computing the integral curves in the same call of the fragment shader, the local synchronization of the two, respective twelve, curves, e.g. regarding stopping criteria within the tensor field, becomes far more easy. Our second strategy is closely related to FTLE visualization on planar cross sections, as published by Garth et al. [12]. Initially, we compute the 3D FSR field at the same resolution as the tensor field and allow it to be navigated through volume rendering or cross sections. Additionally, we offer the possibility to perform the FSR computation on a cross section at a higher resolution, which greatly leverages the richness of details in the FSR field, as shown in Figure 11 with fourfold native resolution.

X. CONCLUSION AND FUTURE WORK

We have presented a method for the visualization of separatrices of coherent regions, a generalized analog of Lagrangian coherent structures. The SCR are based on a generalized
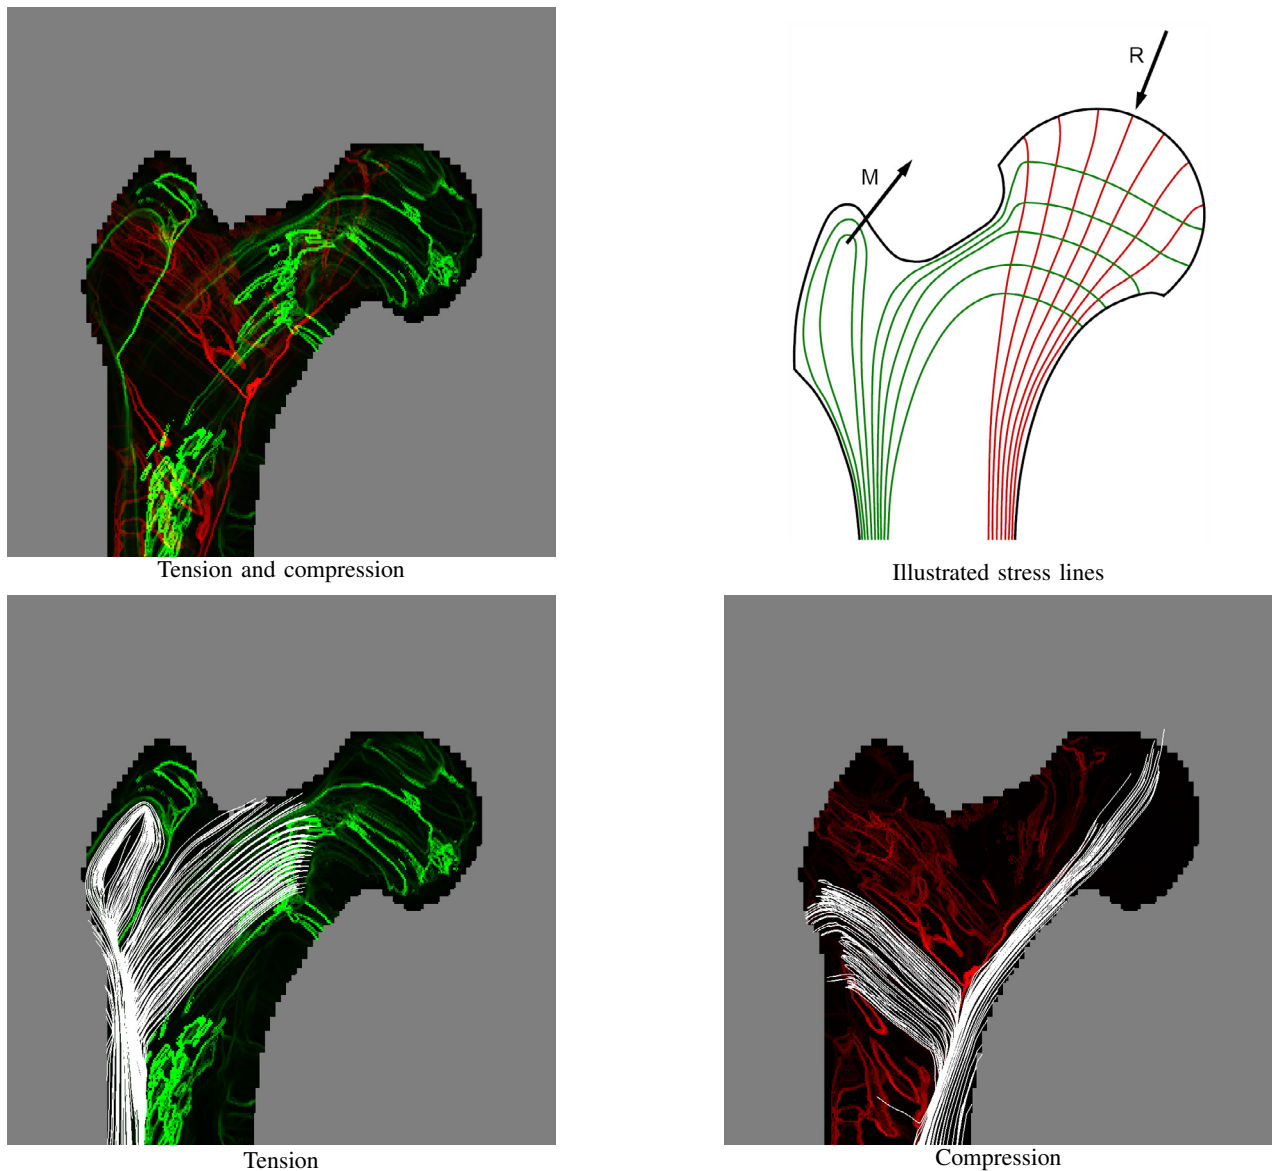
Fig. 10. Slice of stress tensor FSR in femur dataset. Upper left: combined visualization of the FSR for the principal stress direction with highest positive eigenvalue (tension - green) and with highest negative eigenvalue (compression - red). Upper right: illustration of principal stress direction with lines from medical literature (© 2009 IEEE, reprinted with permission from [7]). Lower row: tension FSR field combined with 3D line visualization of the principal stress direction (left); compression FSR field combined with 3D line visualization of the principal stress direction (right). Ridges in the FSR field indicate separating lines of principal stress direction. The revealed structures correspond to those presented in medical literature.

measure of separation, denoted as finite separation ratio, which is defined as the spectral norm of the right Cauchy-Green deformation tensor of coordinate maps, a generalization of flow maps. As an application of this method, we have visualized SCR in symmetric second order tensor fields: synthetic datasets, measured diffusion tensor images, and simulated stress tensor fields. Regarding DT-MRI data, the obtained results have been compared to those of classical methods such as fractional anisotropy as well as derived gradient and ridge measures. We have been able to exemplify that SCR tend to form a superset of the features visible with the methods derived from fractional anisotropy, and that the FSR serves for structuring areas of local linear anisotropy by means of a scalable semi-global analysis, a possibility that the other methods do not expose. Concerning measured

datasets, features in the FSR appear to be of histological relevance. Our results indicate that for the measured canine heart dataset, visible SCR correspond to interfaces between sheets of differing fiber orientation. The results for the stress tensor field prove that our method is not restricted to the analysis of diffusion tensor fields. Ridges in the FSR field can indicate separating principal stress organization and thus support the analysis of stress tensor fields. In addition, we have described a GPU-based implementation for tensor fields that provides acceleration of up to three orders of magnitude compared to a CPU implementation.

Certain aspects of our proposed generalized FSR method and its application require some discussion. First, in Equation 6 we have introduced the FSR as the maximum separation of all m possible mappings in analogy to the established

definition of the Lyapunov exponent. Here, there also may exist application-specific cases where a different measure such as a minimum or average FSR may be desired. Second, regarding diffusion tensor fields, we use eigenvector lines as characteristic lines. This implies that the analysis has highest significance in areas of predominant linear anisotropy because these trajectories in regions with planar anisotropy or spherical isotropy are likely to behave chaotically. An equivalent analysis in areas of insufficiently linear anisotropy is part of our plans for future work. In connection with this, the application of our method to medical imaging methods with high angular resolution is also of interest. They can model multiple directions of diffusion, e.g. of crossing or touching fibers, which could be handled by our generalized model of coherent structures with multiple coordinate maps (Section IV).

Besides its direct visualization and interpretation as a scalar field, further applications for the FSR are conceivable. Garth et al. [12] proposed to use the FTLE field as a probability density function that controls seeding distributions, i.e. in order to preferentially seed particles at positions of high FTLE values. This approach was later presented by Bürger et al. [3]. For diffusion tensor fields we propose a complementary approach. As showcased in Figure 12, where eigenvector lines were seeded at the positions of low FSR values in a slice, this seeding strategy allows the extraction of coherent fiber bundles. Another possible application is the use of the FSR field for segmentation of regions with locally coherent fiber orientation. With increased propagation lengths, it also appears possible to detect anatomical correspondence of coherent fiber bundles, in analogy to tractographical approaches. Also, the application of our method to tensor fields different to those presented here, as they for example arise in image processing or surface modeling, appears promising. Lastly, this work has restricted itself to the visualization of SCR in symmetric tensor fields. Naturally, it appears worthwhile to extend our method both to asymmetric tensors and to other generalized interrelations.

REFERENCES

[1] P. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, 111(3):209–219, 1996.

[2] A. Brun, H. Knutsson, H.-J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber traces using normalized cuts. In *Medical Image Computing and Computer-Assisted Intervention*, pages 368–375, 2004.

[3] K. Bürger, P. Kondratieva, J. Krüger, and R. Westermann. Importance-driven particle techniques for flow visualization. In *Proceedings of IEEE VGTC Pacific Visualization Symposium 2008*, pages 71–78, 2008.

[4] B. Cabral and L. C. Leedom. Imaging vector fields using line integral convolution. In *Proceedings of SIGGRAPH 1993*, pages 263–270, 1993.

[5] T. Delmarcelle and L. Hesselink. Visualizing second-order tensor fields with hyperstreamlines. *IEEE Computer Graphics and Applications*, 13(4):25–33, 1993.

[6] T. Delmarcelle and L. Hesselink. The topology of symmetric, second-order tensor fields. In *Proceedings of IEEE Visualization 1994*, pages 140–147, 1994.

[7] C. Dick, J. Georgii, R. Burgkart, and R. Westermann. Stress tensor field visualization for implant planning in orthopedics. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1399–1406, 2009.

[8] Z. Ding, J. C. Gore, and A. W. Anderson. Classification and quantification of neuronal fiber pathways using diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49(4):716–721, 2003.

[9] D. Eberly. *Ridges in Image and Data Analysis (Computational Imaging and Vision)*. Kluwer Academic Publishers, Doordrecht, The Netherlands, 1996.

[10] L. R. Frank. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 47(6):1083–1099, 2002.

[11] J. D. Furst and S. M. Pizer. Marching ridges. In *Proceedings of IASTED International Conference on Signal and Image Processing 2001*, pages 22–26, 2001.

[12] C. Garth, F. Gerhardt, X. Tricoche, and H. Hagen. Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1464–1471, 2007.

[13] C. Garth, G.-S. Li, X. Tricoche, and C. Hansen. Visualization of coherent structures in transient 2D flows. In H.-C. Hege, K. Polthier, and G. Scheuermann, editors, *Topology-Based Methods in Visualization II*, pages 1–13. Springer, Heidelberg, Germany, 2009.

[14] S. H. Gilbert, A. P. Bensona, P. Li, and A. V. Holden. Regional localisation of left ventricular sheet structure: integration with current models of cardiac fibre, sheet and band structure. *European Journal of Cardio-Thoracic Surgery*, 32(2):231–249, 2007.

[15] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D*, 149(4):248–277, 2001.

[16] G. Haller. Lagrangian coherent structures from approximate velocity data. *Physics of Fluids*, 14(6):1851–1861, 2002.

[17] G. Haller and G. Yuan. Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D*, 147(3-4):352–370, 2000.

[18] K. M. Hasan, P. J. Basser, D. L. Parker, and A. L. Alexander. Analytical Computation of the Eigenvalues and Eigenvectors in DT-MRI. *Journal of Magnetic Resonance*, 152(1):41–47, 2001.

[19] M. Hlawatsch. *Lagrangian Analysis of Symmetric Second Order Tensor Fields*. Diplomarbeit, Universität Stuttgart, 2008.

[20] M. Hlawitschka. *Efficient Visualization of Tensor Fields with Application to Magnetic Resonance Data*. PhD thesis, University of Leipzig, Germany, 2008.

[21] M. Hlawitschka, C. Garth, X. Tricoche, G. Kindlmann, G. Scheuermann, K. Joy, and B. Hamann. Direct visualization of fiber information by coherence. *International Journal of Computer Assisted Radiology and Surgery*, 5(2):125–131, 2010.

[22] M. Hlawitschka and G. Scheuermann. Direkte Darstellung von Faserinformation durch Kohärenzmaße. In *Tagungsband der 7. Jahrestagung der Deutschen Gesellschaft für Computer- und Roboterassistierte Chirurgie e.V.*, pages 73–75, 2008.

[23] M. Hlawitschka, G. Scheuermann, A. Anwander, T. Knösche, M. Tittgemeyer, and B. Hamann. Tensor lines in tensor fields of arbitrary order. In *Proceedings of Third International Symposium on Visual Computing (ISVC 07)*, pages 341–350, 2007.

[24] G. Kindlmann. Superquadric tensor glyphs. In *Proceedings of Eurographics/IEEE VGTC Symposium on Visualization 2004*, pages 147–154, 2004.

[25] G. Kindlmann, D. B. Ennis, R. T. Whitaker, and C.-F. Westin. Diffusion tensor analysis with invariant gradients and rotation tangents. *IEEE Transactions on Medical Imaging*, 26(11):1483–1499, 2007.

[26] G. Kindlmann, X. Tricoche, and C.-F. Westin. Anisotropy creases delineate white matter structure in diffusion tensor MRI. In *Proceedings of MICCAI 2006*, pages 126–133, 2006.
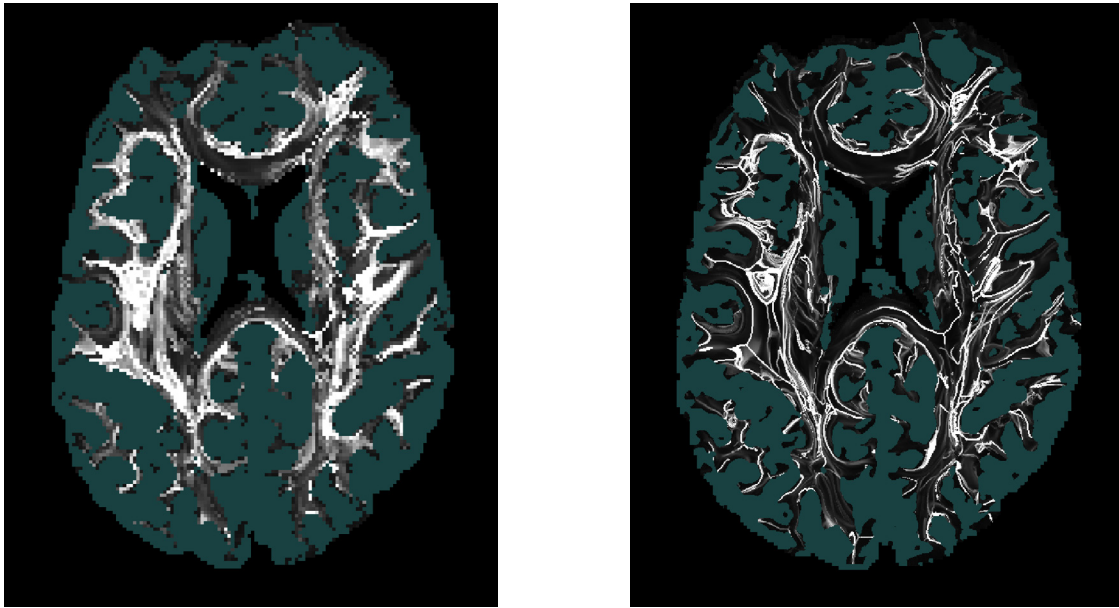
Fig. 11. A comparison of the standard (left) and high resolution (right) FSR field on a cross section of the human brain dataset. Since the integral curves on which the coordinate map is based can be efficiently computed at arbitrary seeding resolution, fine details in the FSR field can be visualized.
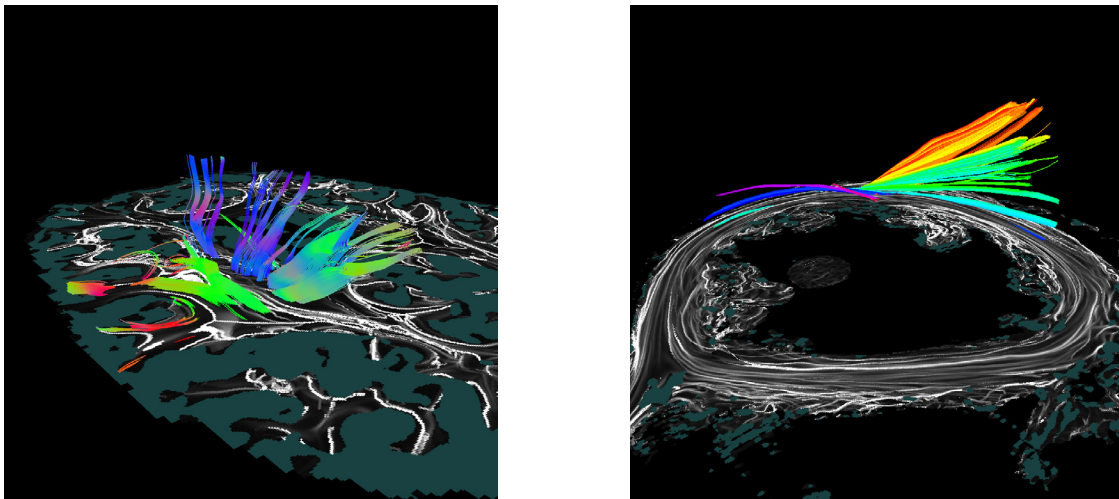

Fig. 12. FSR-guided seeding of eigenvector lines in cross-sections. Lines are seeded in regions with low FSR values within a rectangular region of interest. In the human brain dataset (left), eigenvector lines are RGB color mapped according to their tangent direction. In the canine heart dataset (right), the angle between the XY-plane and the eigenvector at the respective seeding position is mapped to hue in the HSV color model, a common approach in this context. It shows that seeding at positions of low FSR in combination with this color mapping allows one to identify coherent bundles of fibers.

[27] G. Kindlmann, X. Tricoche, and C.-F. Westin. Delineating white matter structure in diffusion tensor MRI with anisotropy creases. *Medical Image Analysis*, 11(5):492–502, 2007.

[28] G. Kindlmann, D. Weinstein, and D. Hart. Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics*, 6(2):124–138, 2000.

[29] P. Kondratieva, J. Krüger, and R. Westermann. The application of GPU particle tracing to diffusion tensor field visualization. In *Proceedings of IEEE Visualization 2005*, pages 73–78, 2005.

[30] D. H. Laidlaw, E. T. Ahrens, D. Kremers, M. J. Avalos, R. E. Jacobs, and C. Readhead. Visualizing diffusion tensor images of the mouse spinal cord. In *Proceedings of IEEE Visualization 1998*, pages 127–134, 1998.

[31] D. Le Bihan. Looking into the functional architecture of the brain with diffusion MRI. In *Proceedings of the International Symposium on Functional and Molecular Imaging of Stroke and Dementia*, pages 1–24, 2006.

[32] F. Lekien, S. C. Shadden, and J. E. Marsden. Lagrangian coherent structures in n-dimensional systems. *Journal of Mathematical Physics*, 48(6):065404, 2007.

[33] G. T. Mase and G. E. Mase. *Continuum Mechanics for Engineers, Second Edition*. CRC Press, Boca Raton, USA, 1999.

[34] S. Mori and P. C. M. van Zijl. Fiber tracking: principles and strategies - a technical review. *NMR in Biomedicine*, 15:468–480, 2002.

[35] L. O'Donnell, W. E. L. Grimson, and C.-F. Westin. Interface detection in diffusion tensor MRI. In *Proceedings of MICCAI 2004*, pages 360–367, 2004.

[36] T. Peeters, A. Vilanova, G. Strijkers, and B. ter Haar Romeny. Visualization of the fibrous structure of the heart. In *Proceedings of Vision, Modeling, and Visualization 2006*, pages 309–316, 2006.

[37] C. Pierpaoli and P. J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36(6):893–906, 1996.

[38] S. K. Robinson. Coherent motions in the turbulent boundary layer. *Annual Review of Fluid Mechanics*, 23:601–639, 1991.

[39] D. Rohmer, A. Sitek, and G. T. Gullberg. Reconstruction and visualization of fiber and sheet structure with regularized tensor diffusion MRI in the human heart. Technical report, Lawrence Berkeley National Laboratory, LBNL-60277, 2006.

[40] F. Sadlo and R. Peikert. Efficient visualization of Lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1456–1463, 2007.

[41] F. Sadlo and R. Peikert. Visualizing Lagrangian coherent structures and comparison to vector field topology. In H.-C. Hege, K. Polthier, and G. Scheuermann, editors, *Topology-Based Methods in Visualization II*, pages 15–29. Springer, Heidelberg, Germany, 2009.

[42] G. Scheuermann and X. Tricoche. Topological methods for flow visualization. In C. Hansen and C. Johnson, editors, *The Visualization Handbook*, pages 341–356. Elsevier, Burlington, USA, 2005.

[43] T. Schultz and H.-P. Seidel. Estimating crossing fibers: A tensor decomposition approach. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1635–1642, 2008.

[44] D. Sosnovik, R. Wang, G. Dai, T. Reese, and V. Wedeen. Diffusion MR tractography of the heart. *Journal of Cardiovascular Magnetic Resonance*, 11(1):47, 2009.

[45] J.-D. Tournier, F. Calamante, D. G. Gadian, and A. Connelly. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23(3):1176–1185, 2004.

[46] X. Tricoche. *Vector and Tensor Topology Simplification, Tracking and Visualization*. PhD Thesis, University of Kaiserslautern, Germany, 2002.

[47] X. Tricoche, G. Kindlmann, and C. Garth. Geometric Assessment of Structural Coherence in DTI. http://www.dagstuhl.de/Materials/Files/07/07022/07022.TricocheXavier.Slides.ppt.

[48] X. Tricoche, G. Kindlmann, and C.-F. Westin. Invariant crease lines for topological and structural analysis of tensor fields. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1627–1634, 2008.

[49] D. S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52(6):1358–1372, 2004.

[50] A. Vilanova, G. Berenschot, and C. van de Pul. DTI visualization with streamsurfaces and evenly-spaced volume seeding. In *Proceedings of Eurographics/IEEE VGTC Symposium on Visualization 2004*, pages 173–182, 2004.

[51] 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 Image Processing of Tensor Fields*, pages 121–153. Springer, Heidelberg, Germany, 2006.

[52] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230(1):77–87, 2004.

[53] D. M. Weinstein, G. L. Kindlmann, and E. C. Lundberg. Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *Proceedings of IEEE Visualization 1999*, pages 249–253, 1999.

[54] C.-F. Westin, S. E. Maier, H. Mamata, A. Nabavi, F. A. Jolesz, and R. Kikinis. Processing and visualization of diffusion tensor MRI. *Medical Image Analysis*, 6(2):93–108, 2002.

[55] C.-F. Westin, S. Peled, H. Gudbjartsson, R. Kikinis, and F. A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *ISMRM '97*, page 1742, 1997.

[56] A. Worth, N. Makris, V. Wedeen Jr, V. Caviness, and D. Kennedy. Fusion of MRI data for visualization of white matter bundles. http://neuromorphometrics.org/papers/Worth/WorthMICCAI98reject.pdf.

[57] F. Zhang and E. R. Hancock. Edge detection and anisotropic diffusion for tensor-valued images. *Proceedings of the International Conference on Image Processing 2006*, pages 229–232, 2006.

[58] X. Zheng and A. Pang. HyperLIC. In *Proceedings of IEEE Visualization 2003*, pages 249–256, 2003.

[59] X. Zheng and A. Pang. Topological lines in 3D tensor fields. In *Proceedings of IEEE Visualization 2004*, pages 313–320, 2004.

[60] X. Zheng, B. Parlett, and A. Pang. Topological structures of 3D tensor fields. In *Proceedings of IEEE Visualization 2005*, pages 551–558, 2005.

[61] L. Zhukov and A. H. Barr. Heart-muscle fiber reconstruction from diffusion tensor MRI. In *Proceedings of IEEE Visualization 2003*, pages 597–602, 2003.

Marcel Hlawatsch received the Diplom (MSc) degree in computer science from Universität Stuttgart, Germany. Since 2008, he has been a PhD student at the Visualization Research Center, Universität Stuttgart (VISUS). His research interests include scientific visualization, GPU methods, and visualization workflows.

Joachim E. Vollrath received the Diplom (MSc) degree in computer science from Universität Stuttgart, Germany, in 2006. He then joined the Visualization Research Center, Universität Stuttgart (VISUS), where his research activities included volume rendering on adaptive grids as well as feature extraction in vector and tensor fields. Since 2009 Joachim has been with FARO Technologies, where he works as 3D software developer.

Filip Sadlo received the Diplom (MSc) degree in computer science from ETH Zurich in 2003 where he also did his PhD at the Computer Graphics Laboratory. Since 2008 he has been a research associate at the Visualization Research Center, Universität Stuttgart (VISUS). His research interests include scientific visualization, 3D reconstruction, and imaging.

Daniel Weiskopf received the Diplom (MSc) degree in physics and the PhD degree in physics, both from Eberhard-Karls-Universität Tübingen, Germany, and he received the Habilitation degree in computer science at Universität Stuttgart, Germany. From 2005 to 2007, Dr. Weiskopf was an assistant professor of computing science at Simon Fraser University, Canada. Since 2007, he has been a professor of computer science at the Visualization Research Center, Universität Stuttgart (VISUS) and at the Visualization and Interactive Systems Institute (VIS), Universität Stuttgart. His research interests include scientific visualization, GPU methods, real-time computer graphics, mixed realities, ubiquitous visualization, perception-oriented computer graphics, and special and general relativity. He is member of ACM SIGGRAPH, the Gesellschaft für Informatik, and the IEEE Computer Society.