

Fabric-like Visualization of Tensor Field Data on Arbitrary Surfaces in Image Space

Sebastian Eichelbaum, Mario Hlawitschka, Bernd Hamann, and Gerik Scheuermann

Abstract

Tensors are of great interest to many applications in engineering and in medical imaging, but a proper analysis and visualization remains challenging. It already has been shown that, by employing the metaphor of a fabric structure, tensor data can be visualized precisely on surfaces where the two eigendirections in the plane are illustrated as thread-like structures. This leads to a continuous visualization of most salient features of the tensor data set.

We introduce a novel approach to compute such a visualization from tensor field data that is motivated by image-space line integral convolution (LIC). Although our approach can be applied to arbitrary, non-selfintersecting surfaces, the main focus lies on special surfaces following important features, such as surfaces aligned to the neural pathways in the human brain. By adding a postprocessing step, we are able to enhance the visual quality of the of the results, which improves perception of the major patterns.

1 Motivation and Related Work

Since the introduction of tensor lines and hyperstreamlines [6], there have been many research efforts directed at the continuous representation of tensor fields, including research on tensor field topology [11, 23, 22]. Zheng and Pang introduced HyperLIC [30], which makes it possible to display a single eigendirection of a tensor field in a continuous manner by smoothing a noise texture along integral lines, while

Sebastian Eichelbaum · Gerik Scheuermann

Abteilung für Bild- und Signalverarbeitung, Institut für Informatik, Universität Leipzig, Germany,
e-mail: {eichelbaum | scheuermann}@informatik.uni-leipzig.de

Mario Hlawitschka · Bernd Hamann

Institute for Data Analysis and Visualization (IDAV), Department of Computer Science, University of California, Davis, CA, e-mail: {hlawitschka | bhamann}@ucdavis.edu

neglecting secondary directions. Recent approaches to visualize Lagrangian structures on tensor fields [12] provide information on one chosen tensor direction and are especially useful for diffusion tensor data, where the main tensor direction can be correlated to neural fibers or muscular structures, whereas the secondary direction only plays a minor role. More recently, Dick et al. [7] published an interactive approach to visualize a volumetric tensor field for implant planning.

Hotz et al. [13] introduced Physically Based Methods (PBM) for tensor field visualization in 2004 as a means to visualize stress and strain tensors arising in geomechanics. A positive-definite metric that has the same topological structure as the tensor field is defined and visualized using a texture-based approach resembling LIC [4]. Besides other information, eigenvalues of the metric can be encoded by free parameters of the texture definition, such as the remaining color space. Whereas the method's implementation for parameterizable surfaces that are topologically equivalent to discs or spheres is straightforward, implementations for arbitrary surfaces remains computationally challenging. In 2009, Hotz et al. [14] enhanced their approach to isosurfaces in three-dimensional tensor fields. A three-dimensional noise texture is computed in the data set and a convolution is performed along integral lines tangential to the eigenvector field. LIC has been used in vector field visualization methods to imitate *Schlieren* patterns on surfaces that are generated in experiments where a thin film of oil is applied to surfaces, which show patterns caused by the air flow. In vector field visualization, image-space LIC is a method to compute *Schlieren*-like textures in image space [27, 28, 17, 9], intended for large and non-parameterized geometries. Besides the non-trivial application of image-space LIC to tensor data, image-space LIC has certain other drawbacks. Mainly because the noise pattern is defined in image space, it does not follow the movement of the surface and, therefore, during user interaction, the three-dimensional impression is lost. A simple method proposed to circumvent this problem is animating the texture pattern by applying randomized trigonometric functions to the input noise. Weiskopf and Ertl [26] solved this problem for vector field visualization by generating a three-dimensional texture that is scaled appropriately in physical space.

We developed and implemented an algorithm similar to the original PBM but for arbitrary non-intersecting surfaces in image space. Our algorithm can perform at interactive frame rates for large data sets on current desktop PCs. We overcome the drawbacks present in image-space LIC implementations by defining a fixed parameterization on the surface. Thus, we do not require a three-dimensional noise texture representation defined at sub-voxel resolution for the data set. Our approach is capable of maintaining local coherence of the texture pattern between frames when (1) transforming, rotating, or scaling the visualization, (2) changing the surface by, e.g., changing isovalues or sweeping the surface through space, and (3) changing the level of detail. In addition, we implemented special application-dependent modes to ensure our method integrates well with existing techniques. Besides this, we also apply several postprocessing steps to further increase the visual quality and clarity of the shown structures.

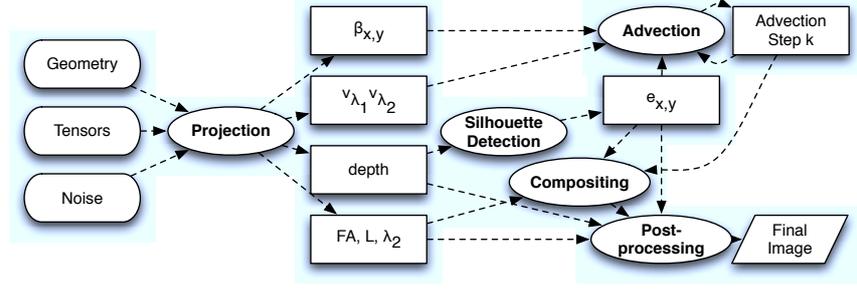


Fig. 1 Flowchart indicating the four major steps of the algorithm: **projection**, which transforms the data set in an image-space representation and produces the initial noise texture on the geometry; **silhouette detection**, required for the advection step and the final rendering; **advection**, which produces the two eigenvector textures; **compositing**, which combines intermediate textures; and the **postprocessing**, which adds additional shading and improves the perceptual quality of the final visualization. Between consecutive steps, the data is transferred using textures.

2 Method

We employ a multi-pass rendering technique that consists of four major rendering passes as outlined in Figure 1. After generating the basic input textures once, the first pass projects all required data into image space. Pass two performs a silhouette detection that is used to guarantee integrity of the advection step computed by multiple iterations of pass three. Eventually, pass four composes the intermediate textures in a final rendering.

2.1 Projection into Image Space

First, we project the data into image space by rendering the surface using the default OpenGL rendering pipeline. Notably, the surface does not need to be represented by a surface mesh. Any other representation that provides proper depth and surface normal information works just as well (e.g., ray-casting methods for implicit surfaces, cf. Knoll et al. [16]). In the same rendering step, the tensor field is transformed from world space to object space, i.e., each tensor T , that is interpolated at the point on the surface from the surrounding two- or three-dimensional tensor field is projected onto the surface by

$$T' = P \cdot T \cdot P^T, \quad (1)$$

with a matrix P defined using the surface normal n as

$$P = \begin{pmatrix} 1 - n_x^2 & -n_y n_x & -n_z n_x \\ -n_x n_y & 1 - n_y^2 & -n_z n_y \\ -n_x n_z & -n_y n_z & 1 - n_z^2 \end{pmatrix}. \quad (2)$$

The camera viewing system configuration and the available screen resolution imply a super- or sub-sampling of the data. We obtain an interpolated surface tensor in every pixel which is decomposed into the eigenvector/eigenvalue representation using a method derived from the one presented by Hasan et al. [10]. These eigenvectors, which are still defined in object space, are projected into image space using the same projection matrices M_M and M_P used for projecting the geometry to image space, usually the standard *modelview* and *projection* matrices OpenGL offers:

$$v'_{\lambda_i} = M_P \times M_M \times v_{\lambda_i}, \text{ with } (i \in 1, 2). \quad (3)$$

Even in the special case of symmetric second-order tensors in \mathbb{R}^3 , which, in general, have three real-valued eigenvalues and three orthogonal eigenvectors in the non-degenerate case, the projected eigenvectors are not orthogonal in two-dimensional space. To simplify further data handling, we scale the eigenvectors as follows:

$$\|v\|_\infty = \max\{|v_x|, |v_y|\} \quad (4)$$

$$v''_{\lambda_i} = \frac{v'_{\lambda_i}}{\|v'_{\lambda_i}\|_\infty} \text{ with } i \in \{1, 2\}, \quad \text{and} \quad \|v'_{\lambda_i}\|_\infty \neq 0 \quad (5)$$

The special case $\|v'_{\lambda_i}\|_\infty = 0$ only appears when the surface normal is perpendicular to the view direction and, therefore, can be ignored. The maximum norm (L_∞ -norm) ensures that one component is 1 or -1 and, therefore, one avoids numerical instabilities arising when limited storage precision is available, and can use memory-efficient eight-bit textures.

2.2 Initial Noise Texture Generation

In contrast to standard LIC approaches, to achieve a proper visual representation of the data, high-frequency noise textures, such as white noise, are not suitable for the compositing of multiple textures. Therefore, we compute the initial noise texture using the reaction diffusion scheme first introduced by Turing [24] to simulate the mixture of two reacting chemicals, which leads to larger but smooth “spots” that are randomly and almost uniquely distributed (cf. Figure 2, right). For the discrete case, the governing equations are:

$$\begin{aligned} \Delta a_{i,j} &= F(i,j) + D_a \cdot (a_{i+1,j} + a_{i-1,j} + a_{i,j+1} + a_{i,j-1} - 4 \cdot a_{i,j}), \\ \Delta b_{i,j} &= G(i,j) + D_b \cdot (b_{i+1,j} + b_{i-1,j} + b_{i,j+1} + b_{i,j-1} - 4 \cdot b_{i,j}), \text{ where} \quad (6) \\ F(i,j) &= s(16 - a_{i,j} \cdot b_{i,j}) \text{ and } G(i,j) = s(a_{i,j} \cdot b_{i,j} - b_{i,j} - \beta_{i,j}). \end{aligned}$$

Here, we assume continuous boundary conditions to obtain a seamless texture in both directions. The scalar s allows one to control the size of the spots where a smaller value of s leads to larger spots. The constants D_a and D_b are the diffusion

constants of each chemical. We use $D_a = 0.125$ and $D_b = 0.031$ to create the input textures.

2.3 Noise Texture Transformation

Mapping the initial texture to the geometry is a difficult and application-dependent task. Even though there exist methods to parameterize a surface, they employ restrictions to the surface (such as being isomorphic to discs or spheres), require additional storage for texture atlases (cf. [19, 15]) and, in general, require additional and often time-consuming pre-processing.

Another solution, proposed by Turk et al. [25], calculates the reaction diffusion texture directly on the surface. A major disadvantage of this method is the computational complexity. Even though these approaches provide almost distortion-free texture representations, isosurfaces, for example, may consist of a large amount of unstructured primitives, which increases the pre-processing time tremendously.

Whereas previously published approaches for image space LIC either use parameterized surfaces to apply the initial noise pattern to the surface or use locally or globally defined three-dimensional textures [26], we define an implicit parameterization of the surface that provides an appropriate mapping of the noise texture to the surface.



Fig. 2 Illustration of the reaction diffusion texture used (right) and the noise texture mapped to geometry (left).

We start by implicitly splitting world space in voxels of equal size, filling the geometry's bounding box, i.e., we define a regular grid. Each voxel i is described by its base coordinate b_i and a constant edge length l . The seamless reaction diffusion texture is mapped to the surface of each of these voxels. To assign a texture coordinate to each vertex, the object space coordinate is transformed to the voxel space that is described by a minimum and maximum coordinate whose connecting line is the bounding box' diagonal. Points v_g on the geometry are transformed to v_{voxel} using

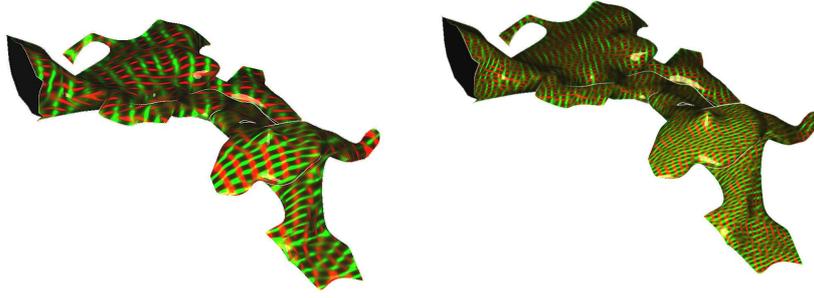


Fig. 3 Comparison of two different values of l to demonstrate the possibility for dynamic refinement of the input noise to achieve different levels of detail.

$$v_{\text{voxel}} = v_g \cdot \begin{pmatrix} l & 0 & 0 & -b_{\text{min}_x} \\ 0 & l & 0 & -b_{\text{min}_y} \\ 0 & 0 & l & -b_{\text{min}_z} \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (7)$$

and transformed into the local coordinate system by

$$v_{\text{hit}} = v_{\text{voxel}} - \lfloor v_{\text{voxel}} \rfloor. \quad (8)$$

The two texture coordinates are chosen to be those two components of v_{hit} that form the plane that is closest to the tangential plane of the surface in this point.

$$t = (v_{\text{hit}_i}, v_{\text{hit}_j}), \text{ with } i \neq j \neq k \wedge (n_k = \max\{n_i, n_j, n_k\}). \quad (9)$$

In other words, this method transforms the world coordinate system to a system defined by one voxel, ensuring that every component of every point v_{hit} is in $[0, 1]$. The texture coordinate is determined by the surface's normal and, in particular, by the voxel side-plane whose normal is most similar to the surface's normal (in terms of angle between them). The use of the term “voxel” is for illustration purposes only; those voxels are never created explicitly.

Regardless of its simplicity, this method supports a continuous parameterization of the surface space that only introduces irrelevant distortions for mapping the noise texture (cf. Figure 2). The mapping is continuous but not C^1 -continuous, which is not required for mapping the noise texture as discontinuities in the first derivatives automatically vanish in the advection step.

Another positive aspect of this mapping is the possibility of a change of scale that is not available in the approaches of, e.g., Turk et al. [25]. By changing the size of voxels during the calculation, different frequencies of patterns can easily be produced and projected onto the geometry. This capability allows one to change the resolution of the texture as required for automatic texture refinement when zooming. A comparison of two different levels of detail is shown in Figure 3.

2.4 Silhouette Detection

To avoid advection over geometric boundaries, a silhouette of the object is required to stop advection in these areas [17]. Otherwise, tensor advection would lead to a constant flow of “particles” across surface boundaries which makes the surface’s geometry and topology unrecognizable.

A standard three-by-three Laplacian filter, defined by the convolution mask

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix} \quad (10)$$

applied to the depth values followed by thresholding, has proven to be suitable for our purposes. The silhouette image $e_{x,y}$ for each pixel (x,y) is stored in the red color channel of its output texture.

2.5 Advection

We have discussed how to project the geometry and the corresponding tensor field to image space. With the prepared image space eigenvectors and the input noise texture, mapped to geometry, advection can be done. We use a simple Euler integration applied to both vector fields. With Euler’s method a particle can be followed along a stream. In our case, we do not calculate streamlines at each position of both vector fields, as normally done in LIC. We directly advect the noise input texture with the given vector fields, which has the same result as locally filtering the data along pre-computed streamlines. This decision was based on the fact that massively parallel architectures like modern GPUs are able to perform this task in parallel for each pixel several hundred times per second. Formally, the advection step can be summarized as follows: First, we assume an input field P to be a continuous function, defined over a two-dimensional domain:

$$f_P : (x,y) \rightarrow p, \text{ with } x,y,p \in [0,1], \quad (11)$$

i.e., it is a function returning the input value of the field P at a given position. Continuity is ensured with interpolating values in-between. With this in mind, the iterative advection on each point (x,y) on the image plane can now be described by

$$\begin{aligned} & \forall x,y \in [0,1] : \forall \lambda \in \{\lambda_1, \lambda_2\} : \\ & p_0^\lambda = \beta_{x,y}, \\ & p_{i+1}^\lambda = k \cdot \beta_{x,y} + (1-k) \cdot \frac{f_{p_i^\lambda}((x,y) + v'_\lambda) + f_{p_i^\lambda}((x,y) - v'_\lambda)}{2}. \end{aligned} \quad (12)$$

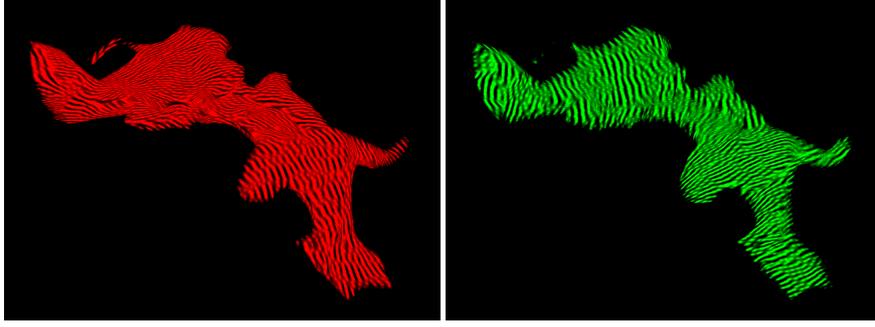


Fig. 4 Advection texture after ten iterations. Left: red channel containing $p_9^{\lambda_1}$, the advection image of eigenvector field 1; right: green channel containing $p_9^{\lambda_2}$, the advection image of eigenvector field 2.

The iterative advection process has to be done for each eigenvector field separately with separate input and output fields p_i as shown in Figure 4. The value at a given point is a combination of the input noise and the iteratively advected input noise from the prior steps. Since the eigenvectors v'_{λ_j} do not have an orientation, the advection has to be done in both directions. The iteration can be stopped when the value change exceeds a threshold.

2.6 Compositing

In a subsequent rendering pass, an initial fabric-like texture is composed. For the sake of simplicity and the limitations of some graphics board, we split the compositing in an initial compositing of four textures followed by a postprocessing step described in the next section, which mainly improves visual quality. Whereas pixels that are not part of the original rendering of the geometry are discarded using the information from the depth buffer, the color values for all other pixels (x, y) in image space after k iterations is defined as:

$$\begin{aligned}
 R(x, y) &= \frac{r \cdot f_{p_k^{\lambda_2}}(x, y)}{8 \cdot f_{p_k^{\lambda_1}}^2(x, y)} + e_{x, y} + \text{light}(\mathcal{L}_{x, y}), \\
 G(x, y) &= \frac{(1 - r) \cdot f_{p_k^{\lambda_1}}(x, y)}{8 \cdot f_{p_k^{\lambda_2}}^2(x, y)} + e_{x, y} + \text{light}(\mathcal{L}_{x, y}), \\
 B(x, y) &= e_{x, y} + \text{light}(\mathcal{L}_{x, y}),
 \end{aligned} \tag{13}$$

where $p_k^{\lambda_1}$ and $p_k^{\lambda_2}$ are the fields generated from the eigenvector advection and e is

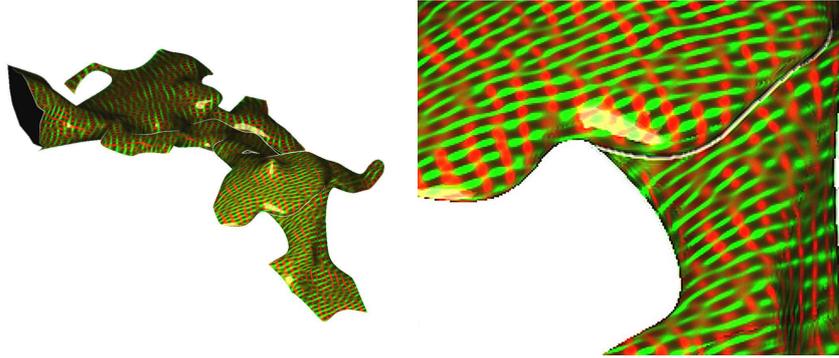


Fig. 5 The composed image produced by the compositing shader with lighting. Left: the whole geometry. Right: a zoomed part of the geometry to show the still blurry fabric structure on the surface.

the silhouette image. The scalar factor r is used to blend between the two chosen tensor directions. Equation 13 is a weighting function that defines the influence of the terms $f_{p_k}^{\lambda_1}(x, y)$ and $f_{p_k}^{\lambda_2}(x, y)$, which influence the green and red color channels, respectively. By squaring the denominator, we emphasize the difference in both fields. This approach creates a mesh resembling the tensor field's structure. To reduce the effect of light sources on the color coding, we use a separate lighting function $light$ that, while manipulating the intensity, does not affect the base color of the mesh structure. Even though Blinn-Phong shading [2] provides the required depth cues, additional emphasis of the third dimension using depth-enhancing color coding has proven to provide a better overall understanding of the data [5].

2.7 Postprocessing

Additional filters can be applied to the composed image, such as contrast enhancement or sharpening filters, which are commonly used in vector field LIC [26, 9]. Figure 5 shows the result of Equation 13 combined with Blinn-Phong shading after applying a sharpening filter.

Bump mapping, first introduced by Blinn [3] to simulate three-dimensionality in planar surfaces, can be used to improve spatial perception of the fabric surface. As bump mapping is normally computed in world space, where the three-dimensional tangent space is known, the textured surface would be required in world space, whereas our texture is parameterized for use in image space. Transforming the modified normal, which is required for bump mapping and, in fact, depends on the gradient information on the surface, from image-space back to world-space is not a trivial task, especially when using a perspective projection. Therefore, we use a modified approach that can be applied in image space only.

Bump mapping requires the surface normal at each point (x, y) of the surface, which can be obtained using the gradient information on each pixel of the surface in image space:

$$g(x, y) = \|\nabla(R + G)(x, y)\|. \quad (14)$$

The resulting two-dimensional vector $g(x, y)$ describes the gradient on the image plane using each pixel's intensity. The blue color channel is not used as it does not contain relevant information besides lighting and edges. It is also worth noting that we exclude the light $\mathcal{L}_{x,y}$ and edge information $e_{x,y}$ from gradient calculation, as we do lighting using bump mapping. Using this gradient, the new surface normal is a weighted sum of the surface normal and the gradient, and is used for calculating Phong lighting $\mathcal{B}_{x,y}$ as seen in Figure 6.

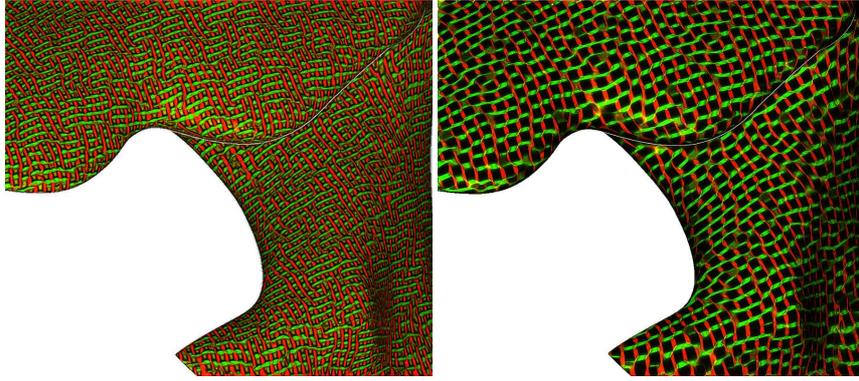


Fig. 6 The final image produced by the postprocessing shader in combination with bump mapping, the geometry's Phong shading and combined edges. Left: standard bump mapping. Right: the same zoomed part of the original geometry to show the effect of weighting the resulting Phong intensities by the original $R(x, y)$ and $G(x, y)$ intensities. This approach creates a more fabric-like impression that can be misunderstood as rotating ribbons similar to stream ribbons.

Figure 6 (right) shows the additional scaling of the red and green color channels by the original color intensities, to lead to a more fabric-like impression of the lines. Equation 15 shows this in more detail:

$$\begin{aligned} R_b(x, y) &= \mathcal{B}_{x,y}(R(x, y)G(x, y) + R^2(x, y)) + e_{x,y} + \text{light}(\mathcal{L}_{x,y}), \\ G_b(x, y) &= \mathcal{B}_{x,y}(R(x, y)G(x, y) + G^2(x, y)) + e_{x,y} + \text{light}(\mathcal{L}_{x,y}). \end{aligned} \quad (15)$$

With the help of bump mapping, we achieve a better spatial impression of the fabric-like pattern. Besides this, postprocessing filters can help to avoid blurry structures.

A further visual improvement can be achieved by interpreting the structure on the surface as streamtubes [29] along the surface. Therefore an approach similar to the ones in [18, 20] is appropriate to create the visual effect of streamtubes on the geometry's surface, without actually creating tubes. First, we need to have some

kind of tangential coordinate system, similar to the one needed for bump mapping. Since our bump mapping is done in image space, the normal of the image plane is $(0, 0, 1)^T$. The eigenvectors v'_{λ_1} and v'_{λ_2} in Equation 3 from the tensor field in image space can be used as the tangent for each field. These tangents denote the direction of the tube along the surface and, together with the normal, define the binormal vector, which is nearly equal to the gradient vector. In practice, it normally is not exactly the same. The binormal b for each eigenvector field i is defined as:

$$b_i = (0, 0, 1)^T \times v'_{\lambda_i}, \text{ with } i \in \{0, 1\}. \quad (16)$$

These binormals can be calculated for each point on the surface. To finally determine the point's actual position on the tube, described by the eigenvectors v'_{λ_i} , one has to find the border of the fabric structure that has been created by the compositing step. Mathematically, this can be expressed in this way:

$$\begin{aligned} B &= \{s | R(sb_1) < \varepsilon \wedge s \in \mathbb{R}\} \wedge (a_p, a_n) \in \mathcal{P}(B), \text{ with} \\ a_p &= \min\{s | s \in B \wedge s \geq 0\} \\ a_n &= \max\{s | s \in B \wedge s < 0\}. \end{aligned} \quad (17)$$

In other words, we find the smallest scaling factors a_n and a_p which scale the binormal vectors b_1 and b_2 in both directions, so that they point on a area below a given threshold ε in the composited image from the prior step, therefore pointing to the border of the tube. As the mapping functions R and G , from Equation 13, only need two-dimensional positions x and y , the binormal's x and y -components are used and the z -component is ignored, as it is always zero. The same factors a_p and a_n for the second eigenvector field are calculated using the green color channel of the composited image, which are used in the same way as described below to render the tubes for the second eigenvector field. The width of the tube at a given point is defined by $a_p + a_n$. The width of the tube is set in relation with the factor a_p to find the actual position of the current point (x, y) on the tube by using:

$$p = \frac{a_p}{a_p + a_n} \in [0, 1], \quad (18)$$

which finally is squared to describe the round surface of a tube:

$$r = \begin{cases} (1 - 2p)^2 & \text{if } p \geq 0 \\ -(1 - 2p)^2 & \text{if } p < 0 \end{cases} \text{ with } ratio \in [-1, 1]. \quad (19)$$

The value of r describes the ratio between the normal completely on the plane (with a zero z -component) and the normal completely pointing towards the camera (with a z -component of one):

$$n = (1 - r)(0, 0, 1)^T + rb_0 \quad (20)$$

The normal n is used to calculate the Phong shading on the surface and produces the tube-like effect with proper spatial impression on the surface, as can be seen in Figure 7.

The artifacts seen in Figure 7 result from the local approach we are using to calculate the tubes. As we do not integrate along the eigenvector-field, there may be discontinuities along a tube in the produced image. There are also artifacts caused by a blurry input field, where borders cannot be found clearly. But, since the frequency of the fabric structure is normally much higher, these effects are not visible anymore, as can be seen in Figure 7, left.

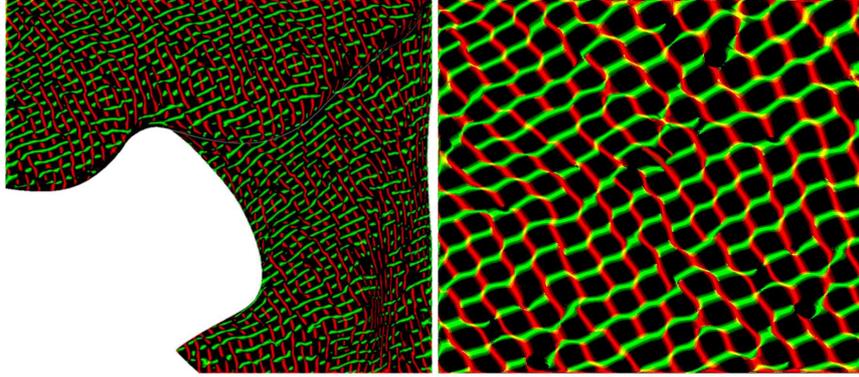


Fig. 7 Left: Interpreting the final image from Figure 5 as streamtubes along the geometry’s surface, and lighting them accordingly, results in a less blurry surface. Right: zoomed part of the left geometry to show the tube effect. Although there are plenty of artifacts in the zoomed image, they do not influence the overall impression of images not zoomed as much. Especially, such strongly zoomed images are not useful for gathering an overview over the tensor field’s structure.

2.8 Implementation

Our implementation is not limited to a special kind of geometry. It is able to handle almost every tensor field defined on a surface. It is, for example, possible to calculate an isosurface on a derived scalar metric, like fractional anisotropy, or on a second data set to generate a surface in a three-dimensional data domain. Other methods include hyper-stream surfaces [6], wrapped streamlines [8], or domain-dependent methods like dissection-like surfaces presented in [1]. The only requirement for the surface is that it is non-selfintersecting and that smooth normals are provided as they are required for the projection step and for proper lighting. The noise texture can be pre-calculated in a pre-processing step or stored in a file as it is independent of the data.

The first step projects the geometry into image space, simply by rendering the geometry and pre-calculating the Phong light intensity $\mathcal{L}_{x,y} \in [0, 1]$ at every rendered fragment with the coordinates x and y . In the same step the tensors are projected as well. When the tensors are symmetric, it is sufficient to transfer six floating-point

values per vertex to the GPU. In our case, two three-dimensional texture coordinates are used per vertex to upload the tensor information along with the geometry. Assuming the tensor T is available on the GPU, it is possible to map the two main directions to the surface described by the normal n at the current vertex using Equation 1. This projection is implemented in a per-vertex manner in the vertex shader. In contrast, to ensure proper interpolation, eigenvalue decomposition and eigenvector calculation together with image space projection need to be done in the fragment shader. Since the eigenvectors are without orientation, it is possible to have sign flips between adjacent vertices. If the interpolation takes place after the eigenvector decomposition, these sign changes can render the interpolation useless. The projection step also includes mapping the noise texture to the geometry. Calculating each vertex' position in one voxel, using the equations from Section 2.3, can be done along the tensor projection. The GPU interpolates those values for each fragment, where it can be used to determine the noise texture element to use.

Since texture space is limited on our hardware, an NVIDIA GeForce 8800 GTS, to four bound textures per rendering pass and active framebuffer object, it is important to store as much information as possible in each texture. Most data are calculated during the projection step and need to be stored in, at most, four textures. Since $\lambda_1 \geq \lambda_2$, it is not necessary to transfer both values. Normalizing the vector (λ_1, λ_2) using the maximum norm as defined in Equations 4 and 5, it suffices to transfer the smaller λ_2 because λ_1 always is one. Also, the eigenvectors v'_{λ_1} and v'_{λ_2} need to be scaled since textures are used for transportation where each value must be in the interval $[0, 1]$.

All intermediate calculations are done on the GPU using the OpenGL shading language (GLSL) in an offscreen *framebuffer object* (FBO). Thus, we avoid the need for rendering the geometry multiple times or even doing those calculations on the CPU, which leads to a large speed gain.

The next step applies an edge detection filter to the depth map, in our case a discrete Laplacian filter kernel, which we have implemented on the GPU using GLSL shaders as a separate offscreen rendering pass. We merge several input data into one texture to decrease the total number of textures. channels of the edge detection shader output texture.

We store the result of the edge detection filter, namely the depth buffer value used to calculate the edges, the unprocessed input noise field, and the noise, which has been mapped to the geometry's surface, in the green, blue, and alpha channel, respectively. By combining multiple data values into a single texture, we reduce the number of texture look-up operations required in the consecutive steps. The used input textures are the values calculated during the projection step, the input noise texture and the depth buffer. The advection step, directly following the edge detection, can "grab" its data directly from a single texture. During advection, it is not advisable to have all values in one texture, since just the noise mapped onto the geometry $\beta_{x,y}$ and the depth buffer are used. We also transfer the unprocessed input noise in this texture, since our implementation is able to toggle which noise field is used during advection because near-planar geometry does not require any complex geometry mapping.

The advection step advects the previously mapped noise texture using the eigenvector color map from the projection step.

Figure 4 shows the red and green channels of this output texture after several iterations. Note that, due to data merging in the prior step, there need to be just three input texture bindings per iteration, instead of four, for depth buffer, edge detection texture, the eigenvector color map and the last iteration’s output texture. Together, both algorithms perform the advection iteration described in Section 2.5. After some iterations, the geometry-mapped noise gets advected more and more and is ready to be processed for output.

The final output processing includes blending both advected eigenvector fields with light information and the geometry edges as well as clipping fragments using the depth buffer and possibly an additional postprocessing step. The last output texture created during advection iteration is used together with the color map containing lighting information, and the edge detection output texture for clipping and edge blending to create the final visualization.

As mentioned in Section 2.6 several modifications are possible. Since the information available in this step also contains the fractional anisotropy, λ_2 and implicitly λ_1 , it is possible to blend-in those values to emphasize additional tensor field features. In Figure 5 both eigenvector fields are blended with the silhouettes and with Phong luminance, calculated earlier in the projection step. Figures 6 and 7 show the post processing possible in a further rendering pass using another postprocessing shader.

3 Results

We have introduced a method to create a fabric-like surface tensor LIC in image space, similar to the one introduced in [13]. We used ideas from [17] to transform the algorithm into image space. Our implementation, using this method, is able to reach frame rates high enough for real-time user interaction. The only bottleneck is the hardware’s ability to render large and triangle-rich geometry. All further steps can be done in constant time, see Table 1.

3.1 Artificial Test Data Sets

We first applied our method to artificial test data sets that have complex topology: a torus, the Bretzel5, and the Tangle data set (cf. [16]), defined as implicit surfaces:

$$(1 - \sqrt{x^2 + y^2})(1 - \sqrt{x^2 + y^2}) + z^2 - 0.125 = 0, \quad (21)$$

$$((x^2 + .25 * y^2 - 1) * (.25 * x^2 + y^2 - 1))^2 + z^2 - 0.1 = 0, \text{ and} \quad (22)$$

$$x^4 - 5 * x^2 + y^4 - 5 * y^2 + z^4 - 5 * z^2 + 11.8 + w = 0. \quad (23)$$

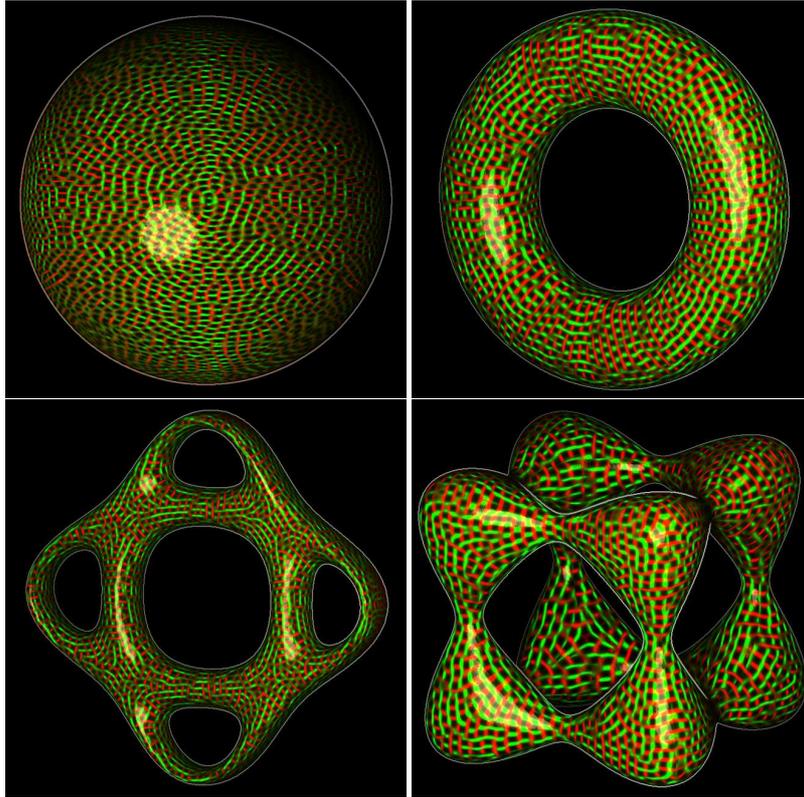


Fig. 8 Analytic test data sets. We applied our method to isosurfaces and the scalar field’s Laplacian to demonstrate the suitability for complicated surfaces. Shown are the final images using our method for a sphere, torus, Tangle, and Bretzel5 data set (Equations 21–23).

We used the Laplacian on the surfaces as tensor fields. The results displayed in Figure 8 show that neither the topology nor our artificial parameterization of the input noise texture influences the quality of the final rendering.

3.2 Modification for Medical Data Processing

Even though many higher-order methods have been proposed, due to scanner, time, and cost limitations, second-order tensor data is still dominant in clinical application. Medical second-order diffusion tensor data sets differ from engineering data sets because they indicate one major direction whereas the secondary and tertiary directions only provide information in areas where the major direction is not well-defined, i.e., the fractional anisotropy—a measure for the tensor shape—is low. Almost spherical tensors, which indicate isotropic diffusion, occur in areas where mul-

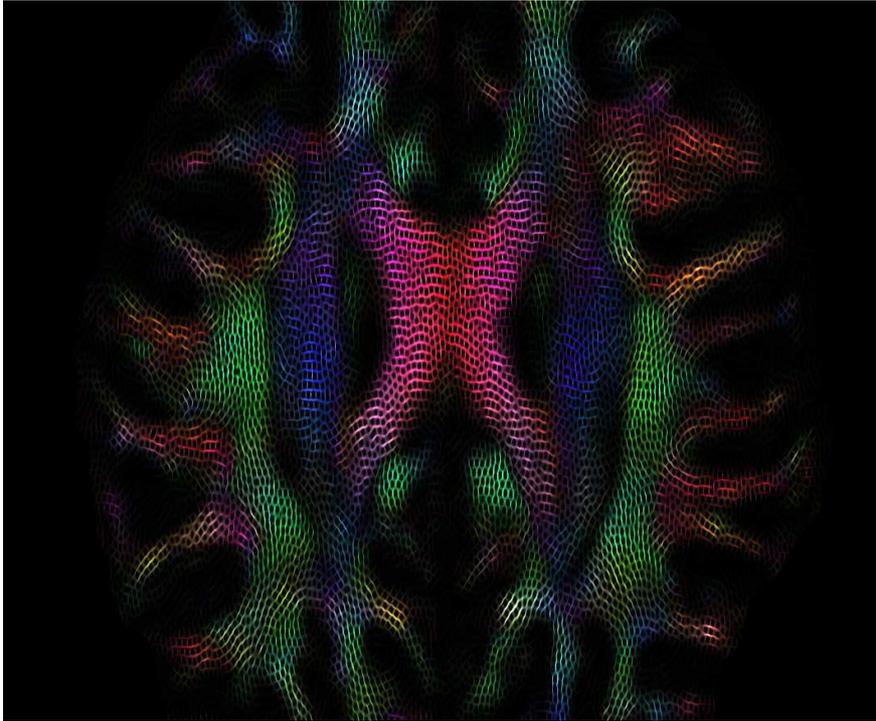


Fig. 9 An axial slice through a human brain: Corpus callosum (CC) (red), pyramidal tract (blue), and parts of the cinguli (green in front and behind the CC) are visible. The main direction in three-dimensional space is indicated by the RGB color map, where red indicates lateral (left–right), green anterior–posterior, and blue superior–inferior direction. The left–right structure of the CC can clearly be seen in its center, whereas color and pattern indicate uncertainty towards the outer parts. The same is true for the cinguli’s anterior–posterior structure. As seen from the blue color, the pyramidal tract is almost perpendicular to the chosen plane and, therefore, secondary and ternary eigenvectors dominate the visualization. Alternatively, we could easily fade out those out-of-plane structures in cases where they distract the user.

multiple fiber bundles traverse a single voxel of the measurement or when no directional structures are present. Therefore, we modulate the color coding using additional information: In areas where one fiber direction dominates, we only display this major direction using the standard color coding for medical data sets, where x , y , and z alignment are displayed in red, green, and blue, respectively. In areas where a secondary direction in the plane exists, we display this information as well but omit the secondary color coding and display the secondary direction in gray-scale rendering mode and always below the primary direction (cf. Figure 10). We use the method of Anwander et al. [1] to extract surfaces that are, where possible, tangential to the fiber directions. Hence, we can guarantee that the projection error introduced by our method in the surface’s domain remains small. Even in areas where the fractional anisotropy is low and the color coding does no longer provide directional informa-

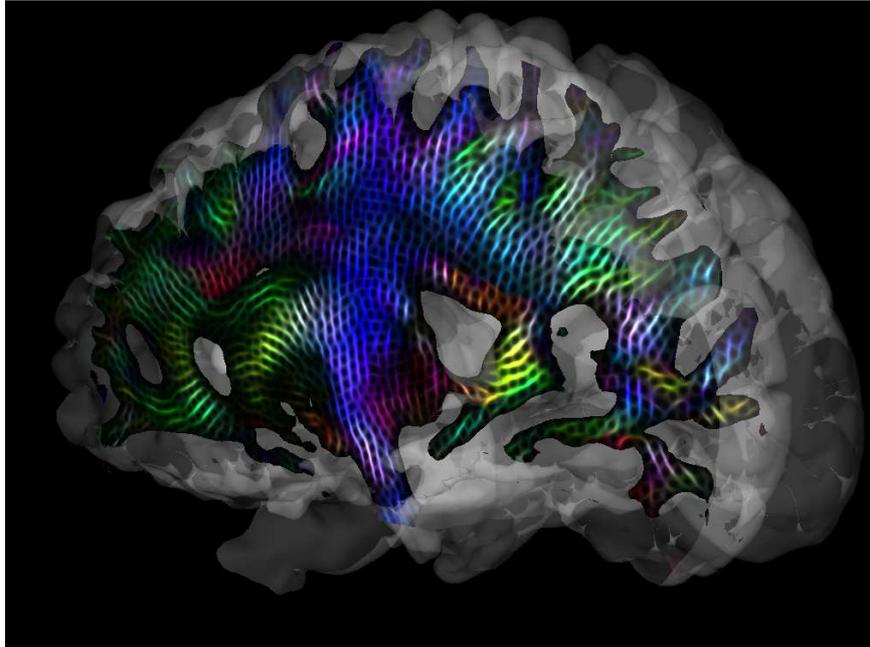


Fig. 10 Diffusion tensor data set of a human brain. We employed the method by Anwander et al. [1] to extract a surface following neural fibers and applied our method with an alternative color coding that is more suitable and can be incorporated more easily into medical visualization tools.

tion, such as in some parts of the pyramidal tract in Figure 10, the texture pattern still provides this information.

3.3 Mechanical Datasets

Our approach is not only applicable to medical datasets, but it can also be applied to many other tensor data sets. Figures 11 and 12 show a slice in an earthquake dataset and an analytical strain tensor field. The analytical data set is the well-known single point load data set, where a single infinitesimally small point source pushes on an infinite surface. The forces and distortions inside the object are represented by stress and strain tensors, which are symmetric, second-order tensors. The earthquake data set is a simulation of a single concrete pile in solid ground excited by a measured earthquake pattern from the Kyoto earthquake (cf. Figure 12). As shown, the data, material stress tensors, are defined on an irregular grid. We extracted a plane perpendicular to the pile and show the tensor information in that plane. Due to the time-dependent nature of the simulation, static images are quite complex.

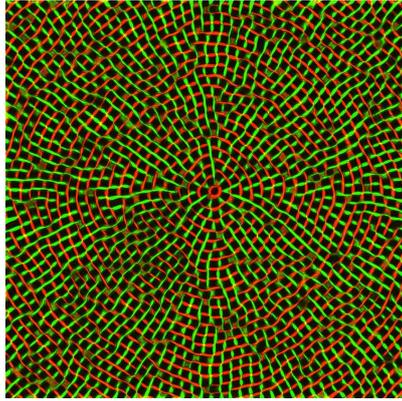


Fig. 11 A slice in the well-known single point load data set, showing the symmetric strain tensor at the surface of the slice.

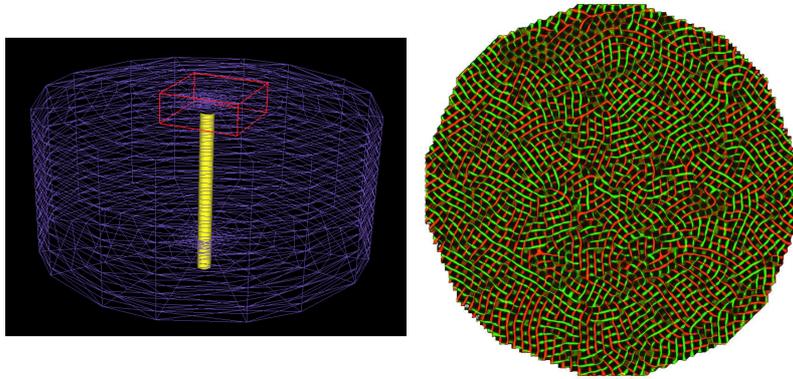


Fig. 12 A concrete pile in solid ground. Left: the original grid shown in purple. Right: a slice of the dataset showing the symmetric part of the tensor field.

3.4 Performance

As indicated before, the only “bottleneck” in the visualization pipeline that is strongly geometry-dependent is the projection step. Since the surface needs to be rendered repeatedly in case of user interaction, the performance measures of our method consider repeated rendering of the geometry. The frame rate with geometry not being moved and, therefore, making the projection step and the edge detection step unnecessary, is considerably higher. Our implementation requires only few advection iterations per frame, which ensures high frame rates and smooth interaction. To make the frame rates in the following tables comparable, user interaction is assumed and, therefore, rendering a single frame always consists of

- one projection step, including geometry rendering;

- one edge detection pass;
- three advection iterations; and
- one output processing pass.

As seen in the previous sections, fragments not belonging to the geometry are discarded as soon as possible without using deferred shading. This also leads to performance gain in advection and output processing. In Table 1, a selection of data sets with their corresponding number of triangles and tensors are listed. The frame rates shown were obtained on an AMD Athlon(tm) 64 X2 Dual Core Processor 3800+ (512K L2 Cache) with a NVIDIA G80 GPU (GeForce 8800 GTS) and 640MB of graphics memory at a resolution of 1024×768 pixels.

| Figure | Nb Triangles | Nb Tensors | fps | fps (Phong only) | ∅ Geometry Share |
|--------|--------------|------------|-----|------------------|------------------|
| 10 | 41472 | 63075 | 32 | 61 | 72% |
| 5 | 58624 | 88803 | 30 | 60 | 69% |
| 9 | 571776 | 861981 | 14 | 16 | 90% |

Table 1 Frames per second (fps) for different data sets with given number of triangles and numbers of tensors. The frame rates are compared to simple rendering of the geometry using Phong shading. The frame rates were obtained for an AMD Athlon(tm) 64 X2 Dual Core Processor 3800+ (512K L2 Cache) with an NVIDIA G80 GPU (GeForce 8800 GTS) and 640MB of graphics memory at a resolution of 1024×768 pixels. The geometry share relates the time used by the GPU to rasterize the geometry to the overall rendering time, which contains all steps of the pipeline. The time used to render the geometry clearly dominates the rendering times and reaches up to 90 percent of the overall rendering time even for medium-sized geometries.

The assumption that geometry rendering with projection is the weakest component in this pipeline and that edge detection, advection, and output processing perform at a data-independent frame rate is confirmed by the frame rates shown in Table 1. It confirms that for large geometries, rendering the geometry alone is the dominant component. Since the vertex-wise calculations during projection are limited to tensor projection (Equation 1) and vertex projection (Equation 7), the most expensive calculations during projection are executed per fragment. This means that the expensive eigenvalue decomposition and eigenvector calculations are only required for fragments (pixels). To further decouple the calculation effort from the geometry’s size, the depth test should be performed before performing the eigendecomposition. This goal can be achieved by first rendering the projected tensors to a texture, and computing the decomposition for visible fragments only. Nevertheless, this is not necessary for our current data set and screen sizes where the time required to render the geometry itself clearly dominates the time required to compute the texture pattern in image space. This can be seen in the increasing values in Table 1 with increasing size of vertices rendered.

4 Conclusions and Possible Directions for Future Research

We have presented a novel method for rendering fabric-like structures to visualize tensor fields on almost arbitrary surfaces without generating three-dimensional textures that span the whole data set at sub-voxel resolution. Therefore, our method can be applied to complex data sets without introducing texture memory problems common to methods relying on tree-dimensional noise textures. As major parts of the calculation are performed in image space, the performance of our algorithm is almost independent of data set size, provided that surfaces can be drawn efficiently, e.g., by using acceleration structures to draw only those parts of the geometry that intersect the view frustum or using ray tracing methods.

Whether the surface itself is the domain of the data, a surface defined on the tensor information (e.g., hyperstream surfaces), or a surface defined by other unrelated quantities (e.g., given by material boundaries in engineering data or anatomical structures in medical data) is independent from our approach. Nevertheless, the surface has to be chosen appropriately because only in-plane information is visualized. To overcome this limitation, information perpendicular to the plane could be incorporated in the color coding, but due to a proper selection of the plane that is aligned with our features of interest, this has not been necessary for our purposes.

Especially in medical visualization, higher-order tensor information is becoming increasingly important and different methods exist to visualize these tensors, including local color coding, glyphs, and integral lines. Nevertheless, an extension of our approach is one of our major aims. In brain imaging, experts agree that the maximum number of possible fiber directions is limited. Typically, a maximum of three or four directions in a single voxel are assumed (cf. Schultz et al. [21]). Whereas the number of output textures can easily be adapted, the major remaining problem is a lack of suitable decomposition algorithms on the GPU. Image-space techniques, by their very nature, resample the data and, therefore, require one to use such proper interpolation schemes. In addition, maintaining orientations and assigning same fibers in higher-order data to the same texture globally is not possible today and, therefore, is a potential topic for further investigation.

Acknowledgements

We thank Alfred Anwander and Thomas R. Knösche from the Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany, for providing the human brain image data sets, and for fruitful discussions and comments, and Boris Jeremić, Department of Civil and Environmental Engineering, UC Davis, for providing the earthquake data set. We thank the members of the Visualization and Computer Graphics Research Group of the Institute for Data Analysis and Visualization, Department of Computer Science, UC Davis, and the members of the Abteilung für Bild- und Signalverarbeitung des Instituts für Informatik der Universität Leipzig,

Germany.

This work has been supported by NSF grant CCF-0702817.

References

1. Anwander, A., Schurade, R., Hlawitschka, M., Scheuermann, G., Knösche, T.R.: White matter imaging with virtual Klingler dissection. *Human Brain Mapping* 2009 (2009)
2. Blinn, J.F.: Models of light reflection for computer synthesized pictures. In: SIGGRAPH '77: Proceedings of the 4th annual conference on Computer graphics and interactive techniques, pp. 192–198. ACM, New York, NY, USA (1977). DOI <http://doi.acm.org/10.1145/563858.563893>
3. Blinn, J.F.: Simulation of wrinkled surfaces. *SIGGRAPH Comput. Graph.* **12**(3), 286–292 (1978). DOI <http://doi.acm.org/10.1145/965139.507101>
4. Cabral, B., Leedom, L.C.: Imaging vector fields using line integral convolution. In: SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, pp. 263–270. ACM, New York, NY, USA (1993). DOI <http://doi.acm.org/10.1145/166117.166151>
5. Chu, A., Chan, W.Y., Guo, J., Pang, W.M., Heng, P.A.: Perception-aware depth cueing for illustrative vascular visualization. In: BMEI '08: Proceedings of the 2008 International Conference on BioMedical Engineering and Informatics, pp. 341–346. IEEE Computer Society, Washington, DC, USA (2008). DOI <http://dx.doi.org/10.1109/BMEI.2008.347>
6. Delmarcelle, T., Hesselink, L.: Visualization of second order tensor fields and matrix data. In: VIS '92: Proceedings of the 3rd conference on Visualization '92, pp. 316–323. IEEE Computer Society Press, Los Alamitos, CA, USA (1992)
7. Dick, C., Georgii, J., Burgkart, R., Westermann, R.: Stress tensor field visualization for implant planning in orthopedics. *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 1399–1406 (2009). DOI <http://doi.ieeeecomputersociety.org/10.1109/TVCG.2009.184>
8. Enders, F., Sauber, N., Merhof, D., Hastreiter, P., Nimsky, C., Stamminger, M.: Visualization of white matter tracts with wrapped streamlines. In: C.T. Silva, E. Gröller, H. Rushmeier (eds.) *Proceedings of IEEE Visualization 2005*, pp. 51–58. IEEE Computer Society, IEEE Computer Society Press, Los Alamitos, CA, USA (2005)
9. Grabner, M., Laramée, R.S.: Image space advection on graphics hardware. In: SCCG '05: Proceedings of the 21st spring conference on Computer graphics, pp. 77–84. ACM, New York, NY, USA (2005). DOI <http://doi.acm.org/10.1145/1090122.1090136>
10. Hasan, K.M., Basser, P.J., Parker, D.L., Alexander, A.L.: Analytical computation of the eigenvalues and eigenvectors in DT-MRI. *Journal of Magnetic Resonance* **152**(1), 41 – 47 (2001). DOI [10.1006/jmre.2001.2400](http://doi.org/10.1006/jmre.2001.2400)
11. Hesselink, L., Levy, Y., Lavin, Y.: The topology of symmetric, second-order 3d tensor fields. *IEEE Transactions on Visualization and Computer Graphics* **3**(1), 1–11 (1997). DOI <http://dx.doi.org/10.1109/2945.582332>
12. Hlawitschka, M., Garth, C., Tricoche, X., Kindlmann, G., Scheuermann, G., Joy, K.I., Hamann, B.: Direct visualization of fiber information by coherence. *International Journal of Computer Assisted Radiology and Surgery, CARS, CUARC.08 Special Issue* (2009)
13. Hotz, I., Feng, L., Hagen, H., Hamann, B., Joy, K., Jeremic, B.: Physically based methods for tensor field visualization. In: VIS '04: Proceedings of the conference on Visualization '04, pp. 123–130. IEEE Computer Society, Washington, DC, USA (2004). DOI <http://dx.doi.org/10.1109/VIS.2004.80>
14. Hotz, I., Feng, Z.X., Hamann, B., Joy, K.I.: Tensor field visualization using a fabric-like texture on arbitrary two-dimensional surfaces. In: T. Möller, B. Hamann, R.D. Russel (eds.) *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Springer-Verlag Heidelberg, Germany (2009)

15. Iwakiri, Y., Omori, Y., Kanko, T.: Practical texture mapping on free-form surfaces. In: PG '00: Proceedings of the 8th Pacific Conference on Computer Graphics and Applications, p. 97. IEEE Computer Society, Washington, DC, USA (2000)
16. Knoll, A., Hijazi, Y., Hansen, C., Wald, I., Hagen, H.: Interactive ray tracing of arbitrary implicits with simd interval arithmetic. In: RT '07: Proceedings of the 2007 IEEE Symposium on Interactive Ray Tracing, pp. 11–18. IEEE Computer Society, Washington, DC, USA (2007). DOI <http://dx.doi.org/10.1109/RT.2007.4342585>
17. Laramee, R.S., Jobard, B., Hauser, H.: Image space based visualization of unsteady flow on surfaces. In: VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03), p. 18. IEEE Computer Society, Washington, DC, USA (2003). DOI <http://dx.doi.org/10.1109/VISUAL.2003.1250364>
18. Merhof, D., Sonntag, M., Enders, F., Nimsky, C., Hastreiter, P., Greiner, G.: Hybrid visualization for white matter tracts using triangle strips and point sprites. *IEEE Transactions on Visualization and Computer Graphics* **12**(5), 1181–1188 (2006). DOI <http://doi.ieeecomputersociety.org/10.1109/TVCG.2006.151>
19. Purnomo, B., Cohen, J.D., Kumar, S.: Seamless texture atlases. In: SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing, pp. 65–74. ACM, New York, NY, USA (2004). DOI <http://doi.acm.org/10.1145/1057432.1057441>
20. Schirski, M., Kuhlen, T., Hopp, M., Adomeit, P., Pischinger, S., Bischof, C.: Virtual tubelets—efficiently visualizing large amounts of particle trajectories. *Comput. Graph.* **29**(1), 17–27 (2005). DOI <http://dx.doi.org/10.1016/j.cag.2004.11.004>
21. Schultz, T., Seidel, H.P.: Estimating crossing fibers: A tensor decomposition approach. *IEEE Transactions on Visualization and Computer Graphics* **14**(6), 1635–1642 (2008). DOI <http://doi.ieeecomputersociety.org/10.1109/TVCG.2008.128>
22. Tricoche, X.: Vector and tensor field topology simplification, tracking, and visualization. Ph.D. thesis, University of Kaiserslautern, Germany (2002)
23. Tricoche, X., Scheuermann, G., Hagen, H.: Tensor topology tracking: A visualization method for time-dependent 2D symmetric tensor fields. In: Eurographics 2001 Proceedings, Computer Graphics Forum 20(3), pp. 461–470. The Eurographics Association, Saarbrücken, Germany (2001). DOI <http://dx.doi.org/10.1111/1467-8659.00539>
24. Turing, A.: The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London* **237**(641), 37 – 72 (1952)
25. Turk, G.: Generating textures on arbitrary surfaces using reaction-diffusion. In: SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques, pp. 289–298. ACM, New York, NY, USA (1991). DOI <http://doi.acm.org/10.1145/122718.122749>
26. Weiskopf, D., Ertl, T.: A hybrid physical/device-space approach for spatio-temporally coherent interactive texture advection on curved surfaces. In: GI '04: Proceedings of Graphics Interface 2004, pp. 263–270. Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada (2004)
27. van Wijk, J.J.: Image based flow visualization. In: SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques, pp. 745–754. ACM, New York, NY, USA (2002). DOI <http://doi.acm.org/10.1145/566570.566646>
28. van Wijk, J.J.: Image based flow visualization for curved surfaces. In: VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03), p. 17. IEEE Computer Society, Washington, DC, USA (2003). DOI <http://dx.doi.org/10.1109/VISUAL.2003.1250363>
29. Zhang, S., ?agatay Demiralp, Laidlaw, D.H.: Visualizing diffusion tensor mr images using streamtubes and streamsurfaces. *IEEE Transactions on Visualization and Computer Graphics* **9**(4), 454–462 (2003). DOI <http://doi.ieeecomputersociety.org/10.1109/TVCG.2003.1260740>
30. Zheng, X., Pang, A.: Hyperlic. In: VIS '03: Proceedings of the 14th IEEE Visualization 2003 (VIS'03), p. 33. IEEE Computer Society, Washington, DC, USA (2003). DOI <http://dx.doi.org/10.1109/VISUAL.2003.1250379>