

Diffusion-Based Tractography: Visualizing Dense White Matter Connectivity from 3D Tensor Fields

S. Muraki¹, I. Fujishiro², Y. Suzuki², Y. Takeshima²

¹National Institute of Advanced Industrial Science and Technology, Japan

²Institute of Fluid Science, Tohoku University, Japan

Abstract

In this paper, we present a novel method, called diffusion-based tractography (DBT), for visualizing diffusion tensor magnetic resonance imaging datasets. The DBT method generates 3D textures similar to the line integral convolution (LIC) by smearing 3D random dot textures. In contrast to the LIC method, which only traces a single direction, the DBT method takes into account both linear and planar diffusion components, and suppresses excessive blur by an analysis of three decomposed components. We will demonstrate that the DBT method is effective for visualizing dense white matter connectivity from 3D diffusion tensor fields and that it is suitable for hardware acceleration using commodity graphics processors.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Line and Curve Generation I.3.8 [Computer Graphics]: Applications I.6.8 [Simulation and Modeling]: Visual

1. Introduction

Diffusion tensor magnetic resonance imaging (DT-MRI) is a technique used to measure the anisotropic diffusion properties of biological tissues. Since the anisotropic diffusion of water within the white matter of the brain follows the direction of a nerve axon [BMPL94], visualizing DT-MRI datasets is quite useful for diagnosing various brain diseases.

Figure 1 presents examples of DT-MRI datasets, which are the 20th slices of 30 transversal slices. The diffusion tensor can be computed from several DT-MRI datasets (S_i) obtained by applying different motion-probing-gradient (MPG) directions and one null dataset (S_0) obtained without applying MPG. A single diffusion tensor T is represented by a symmetric 3×3 second-order matrix:

$$T = \begin{pmatrix} T_{xx} & T_{xy} & T_{xz} \\ T_{xy} & T_{yy} & T_{yz} \\ T_{xz} & T_{yz} & T_{zz} \end{pmatrix}.$$

Since there are six unknowns in T , we need DT-MRI datasets of at least six different MPG directions.

The signal intensity $S_i = S_i(x, y, z)$ under an MPG $G_i = (G_x, G_y, G_z)^T$ is represented by

$$S_i = S_0 \exp(-bG_i^T T G_i),$$

where b is a strength parameter called the b -value. By solving the simultaneous equation

$$\ln \frac{S_i}{S_0} = -bG_i^T T G_i$$

of several different G_i , we can determine a diffusion tensor for each voxel. However, the visualization of tensor volume data, whose voxel has a symmetric 3×3 second-order matrix, needs a more sophisticated technique than those used in the conventional vector volume data visualizations.

In this paper, we propose the *diffusion-based tractography* (DBT) method as a new technique for visualizing dense white matter connectivity from DT-MRI datasets. A simple local operation of the DBT method is suitable for parallel processing using PC clusters or hardware acceleration using graphics processors (GPUs).

The remainder of this paper is organized as follows. The next section will give an overview of related work. In Section 3, we will describe the basic principle and computation scheme of the DBT method. In section 4, we will demonstrate the expressiveness of the DBT method by using cross-section displays and volume renderings. The last section will conclude the paper with a few remarks on future work.




Figure 1: Slices of seven DT-MRI datasets. Data size: $256 \times 256 \times 30$. Voxel size: $0.8984375\text{mm} \times 0.8984375\text{mm} \times 3.3\text{mm}$.

2. Related Work

Diffusion tensor T is a real symmetric matrix, and has three orthogonal eigenvectors $\{v_1, v_2, v_3\}$ and three corresponding positive eigenvalues $\{\lambda_1, \lambda_2, \lambda_3\}$. By using the orthogonal matrix

$$V = (v_1, v_2, v_3),$$

we can decompose T into the diagonal matrix:

$$\begin{aligned} T &= VDV^T \\ &= V \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} V^T \end{aligned} \quad (1)$$

In this case, v_1 represents the direction of the fastest diffusion and λ_1 the speed of the diffusion, where we assume that $\lambda_1 \geq \lambda_2 \geq \lambda_3$.

The simplest tensor field visualization method is to generate a scalar field, which represents the anisotropy of the diffusion for each voxel, and to illustrate the scalar values by a 2D cross-section image. Fractional anisotropy (FA) and relative anisotropy (RA) are often used in the field of medicine as metrics of anisotropy [PB00]. In the field of visualization research, three different kinds of anisotropy metrics, c_l , c_p , and c_s (linear, planar, and spherical), are more popular [WPG*97]:

$$\begin{aligned} c_l &= \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3} \\ c_p &= \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3} \\ c_s &= \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}. \end{aligned}$$

Instead of using scalar fields, there are known methods assigning a glyph, which represents the diffusion property, to each pixel location. The most popular glyph is an ellipsoid, whose axes represent three eigenvectors and their corresponding eigenvalues [PB00]. Laidraw et al. [LAKR98] used 2D multilayer brush strokes to represent several parameters of the tensor field. Although these methods could compress a lot of information into a single image, it was difficult to demonstrate the continuity of the tensor field. It was also unsuitable for visualizing the 3D tensor field since many glyphs or brush strokes obstruct the 3D structure of the tensor fields.

Tractography tracks streamlines in the major eigenvector direction and the resulting streamline corresponds well with the nerve axon tract [XZC*99]. However, the voxel of DT-MRI is much bigger than the diameter of a nerve axon, and a single voxel of DT-MRI can include many axon fibers of different directions. Because of this partial volume effect, the planar diffusion components are observed in regions where fibers cross or fan out, and special care is needed to track the actual fiber direction [WKL99, ZB02]. Zhang et al. [ZDL03] visualized both linear and planar diffusion components by using primitives called *streamtubes*, which are similar to the *hyperstreamline* [DH93] and *streamsurfaces*.

Line integral convolution (LIC) [CL93], which visualizes the vector field by smearing a random dot texture along the local streamlines, is currently one of the most popular vector visualization methods. Some strategies for using LIC as a 3D flow visualization tool have been investigated (e.g., [IG97, SFCN02]). Hsu attempted to apply the LIC method, but only in the major or medium eigenvector direction of the diffusion tensor field [Hsu01]. There have also been a few attempts to extend LIC for tensor fields by introducing a multidimensional convolution kernel [SEHW02, ZP03]. We assume that the motivation of using convolutions for diffusion tensor visualization comes from the analogy with the mathematical equivalence between the isotropic diffusion and the Gaussian filter convolution. However, anisotropic diffusion cannot be equivalent to the multidimensional convolution in general, and there is no mathematical proof that the multidimensional convolution of a random dot texture can always visualize the tensor field structure.

Kindlmann et al. [KWH00] proposed several DT-MRI visualization techniques based on volume rendering. They used reaction-diffusion simulation to generate 3D patterns similar to ellipsoidal glyphs that enhance the diffusion tensor property. O'Donnell et al. [OHW02] also used diffusion simulation to quantify the white matter connectivity. However, these papers did not mention the use of their diffusion simulations for fibrous structure visualization. We assume that this is because there is a strong blurring effect caused by the spherical diffusion components of the diffusion tensor, which we will discuss in the next section.




Figure 2: Diffusion of random dot pattern by numerical simulation. (a) Initial value $u_0(x, y, z)$. The random dot patterns are generated in the white matter candidate area $\{(x, y, z) | S_0(x, y, z) \geq 100.0, c_l \geq 0.1\}$. (b)-(d) Diffused random dots. Only the 32nd slice images are shown though this simulation is performed in 3D. Data size: $256 \times 256 \times 107$. Voxel size: $0.8984375\text{mm} \times 0.8984375\text{mm} \times 0.8984375\text{mm}$. On-line update ($\Delta t = 0.333$).

3. Diffusion-Based Tractography

In this section, we propose the diffusion-based tractography (DBT) method for visualizing dense white matter connectivity from diffusion tensor volume data. The DBT method smears 3D random dots, which are generated in the white matter candidate area, by numerically solving the following diffusion equation:

$$\frac{\partial u_T(x, y, z)}{\partial t} = \nabla \cdot \mathbf{T}(x, y, z) \nabla u_T(x, y, z). \quad (2)$$

The first-order explicit Euler integration is adopted as the numerical solution scheme for Equation (2), which computes approximate solutions from an initial solution $u_0(x, y, z)$ by performing the following iteration:

$$u_{T+1}(x, y, z) = u_T(x, y, z) + \frac{\partial u_T(x, y, z)}{\partial t} \Delta t \quad (3)$$

for certain steps, where Δt denotes a time step. We can approximate the right side of Equation (2) with a certain finite difference scheme.

There are two ways of computing Equation (3), i.e. on-line update and batch update. The on-line update method consecutively applies Equations (2) and (3) for each voxel whereas the batch update method applies Equation (3) after computing Equation (2) for all voxels. In general, the batch update tends to diverge and requires setting the time step smaller than the on-line update. However, the batch update method is useful when the diffusion simulation is performed on massively parallel processors or GPUs.

As the initial solution $u_0(x, y, z)$, we used a 3D random dot texture similar to that of LIC. However, we do not generate dots in the background of the brain or in areas where the isotropic metric is sufficiently high. Here, we can easily distinguish the brain area and the background by thresholding the null dataset S_0 . Accordingly, we can avoid generating patterns in the isotropic area (e.g. ventricles and the gray matter), which obstructs the view of the fibrous structure of the white matter.

A method similar to LIC using a diffusion equation to generate the 3D texture has been proposed in [DPR00]. However, their method seeks to visualize a vector field, and to generate diffusion tensors from the vector field by using tensor products. In contrast, our method uses genuine diffusion tensors to generate a 3D texture similar to LIC for visualizing the 3D tensor field structure.

Figure 2 depicts the result of the numerical simulation of the random-dot diffusion, where we used the diffusion tensor field obtained from the seven DT-MRI datasets shown in Figure 1. Each DT-MRI dataset ($256 \times 256 \times 30$ voxels) was super-sampled by tri-linear interpolation into $256 \times 256 \times 107$ voxels so that the slice interval becomes the same as the pixel interval (0.8984375 mm). All experiments throughout this paper were performed on a Xeon 3.0 GHz Linux (Fedora Core 1) PC with 2 GB RAM. For the initial solution, we generated random dots in the white matter candidate area specified with $\{(x, y, z) | S_0(x, y, z) \geq 100.0, c_l \geq 0.1\}$. We used the on-line update method using $\Delta t = 0.333$ since the iteration did not converge when $\Delta t \geq 0.5$. Contrary to our expectation, however, the random dots became blurred images as the iterations continued, and we could not obtain images that visualize the fibrous structure of white matter. We investigated the distribution of anisotropy metrics by generating a barycentric histogram [WKL99], as shown in Figure 3, which clearly shows that the proportion of spherical diffusion (higher c_s) is dominant in this dataset and that these spherical diffusions seem to be working as smoothing filters to excessively blur the random dots.




Figure 3: Barycentric histogram (log scale intensity).




Figure 4: Results of DBT. First row: Linear diffusion $T = VD_lV^T$. Second row: Planar diffusion $T = VD_pV^T$. Third row: Mixed-diffusion. Only the 32nd slice images are shown though these simulations are performed in 3D. Data size: $256 \times 256 \times 107$. Voxel size: $0.8984375\text{mm} \times 0.8984375\text{mm} \times 0.8984375\text{mm}$. On-line update ($\Delta t = 1.0$).

We then decomposed the diagonal matrix D of Equation (1) as

$$D = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} = D_l + D_p + D_s,$$

where

$$D_l = \begin{pmatrix} \lambda_1 - \lambda_2 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$D_p = \begin{pmatrix} \lambda_2 - \lambda_3 & 0 & 0 \\ 0 & \lambda_2 - \lambda_3 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$D_s = \begin{pmatrix} \lambda_3 & 0 & 0 \\ 0 & \lambda_3 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix},$$

and defined a new diffusion tensor T' by combining these diagonal matrices to suppress the excessive spherical diffusion effects.

The first row of Figure 4 illustrates the results of using the linear diffusion tensor

$$T' = VD_lV^T$$

by numerically solving the following partial differential equation (PDE):

$$\frac{\partial u_T(x, y, z)}{\partial t} = \nabla \cdot T'(x, y, z) \nabla u_T(x, y, z).$$

The initial solution was binary random dots, the same as in Figure 2. However, we could perform on-line updates with $\Delta t = 1.0$ since the strong spherical diffusion component was not included in T' . In contrast to Figure 2, fibrous patterns similar to those in LIC appeared in the regions like the corpus callosum as the iteration continued. Random dot patterns were also observed in regions like the corona radiata, where the nerve axons ran perpendicular to the image, since the diffusion parallel to the image was not strong in these areas.

The second row of Figure 4 depicts the results of using the planar diffusion tensor

$$T' = VD_pV^T.$$

The initial solution was the binary random dots generated

in the regions of $\{(x, y, z) | S_0 \geq 100.0, c_p \geq 0.4\}$. As the iteration continued, the random dots were smoothed out and became surfaces. These surfaces were similar to streamsurfaces [ZDL03], but we could generate them with relative ease.

As pointed out in many diffusion tensor visualization studies [WKL99, ZB02], tracking only major eigenvectors fails to extract the actual neural pathway. Furthermore, using planar diffusion in regions where linear diffusion is sufficiently high excessively blurs the DBT texture. We then define the *mixed-diffusion model* that empirically changes T' for four different areas. For the region of $\{(x, y, z) | S_0 \geq 100.0, c_l \geq 0.1, c_p < 0.4\}$, it uses the linear diffusion tensor as T' . For the region of $\{(x, y, z) | S_0 \geq 100.0, c_l < 0.1, c_p \geq 0.4\}$, it uses the planar diffusion tensor as T' . For the region of $\{(x, y, z) | S_0 \geq 100.0, c_l \geq 0.1, c_p \geq 0.4\}$, it uses the linear-planar diffusion tensor:

$$\begin{aligned} T' &= \mathbf{V}(\mathbf{D}_l + \mathbf{D}_p)\mathbf{V}^T \\ &= \mathbf{V} \begin{pmatrix} \lambda_1 - \lambda_3 & 0 & 0 \\ 0 & \lambda_2 - \lambda_3 & 0 \\ 0 & 0 & 0 \end{pmatrix} \mathbf{V}^T \quad (4) \\ &= \mathbf{V}\mathbf{D}_a\mathbf{V}^T \end{aligned}$$

as T' to diffuse a random dot texture to consider both the linear and planar diffusion components while excluding the spherical diffusion component.

The third row of Figure 4 illustrates the result of this mixed-diffusion model. The fibrous patterns of the first row of Figure 4 are smoothed out in the regions where the planar diffusion is strong, while the fibrous pattern in the strong linear diffusion area, e.g. corpus callosum, is maintained. Figure 5 depicts the relationship between the anisotropy metrics and the diagonal matrixes of the mixed-diffusion model. According to Figure 5, the linear-planar diffusion matrix is used in the area of $\{(x, y, z) | S_0 \geq 100.0, c_l < 0.1, c_p < 0.4\}$. Since there is no structure to be visualized, it is natural to consider that no diffusion should be performed in this area. However, unpleasant discontinuities appear around the edge of white matter if we omit the diffusion, as shown in Figure 6(a). Figure 6(b) is the pseudo-color display of Figure 6(a) emphasizing the discontinuities, and Figure 6(c) is the same pseudo-color display of Figure 4(k). Figure 6(c) clearly demonstrates that moderate diffusion around the white matter edge (red area) functions as an antialias filter to reduce the discontinuities. In the area of $\{(x, y, z) | S_0 \geq 100.0, c_l < 0.1, c_p < 0.4\}$, the spherical diffusion is dominant and the diffusion caused by the linear-planar diffusion tensor is sufficiently small to maintain the details.

We also applied the batch update. Although the iteration diverged with $\Delta t = 1.0$, we could obtain the result similar to Figure 4 (k) after 200 iterations with $\Delta t = 0.2$.




Figure 5: Relation between the anisotropy metrics and the diagonal matrixes of the mixed-diffusion model.

4. Discussions

One of the merits of the DBT method is that it provides a dense 3D texture that visualizes the tensor field structure. Once this kind of volume data is obtained, we can extract arbitrary cross-section images of the 3D DBT texture. Since medical doctors are familiar with 2D slice images, demonstrating the DBT slice image and the normal MR image of the same slice location will be quite useful for understanding white matter connectivity that cannot be seen in the latter modality.

Figure 7 illustrates slice images of DBT texture generated in the same way as Figure 4 (k). Figure 7(a) is a sagittal slice at $x = 126$ revealing the structure of the corpus callosum. Since the corpus callosum is a bundle of nerve fibers connecting two hemispheres, it is seen as a random dot pattern in this sagittal slice. Figure 7(a) also depicts part of the vertically running corticospinal tracts and the pons. Figure 7(b) is a coronal slice at $y = 128$ revealing the intersection of pontocerebellar fibers and corticospinal tracts. It also illustrates the fibrous structure of the corpus callosum, and the white matter connectivity of the frontal and the temporal gyri. Figure 7(c) is a transversal slice at $z = 91$ that depicts the fibrous structure of the pontocerebellar fibers. The random dot patterns are observed in the cross sections of the corticospinal tracts. Figure 7(d) is a transversal slice at $z = 21$. A vertically running pair of cingulum bundles is observed.

The fibrous structures of Figure 7 are too thick to see the structure of human white matter because the voxel interval of the original DT-MRI (0.8984375 mm) is too large compared to the size of the human nerve fiber. However, we can improve the resolution of the DBT texture by super-sampling the DT-MRI datasets. Figure 8 illustrates DBT texture generation from super-sampled DT-MRI datasets. The seven DT-MRI datasets were super-sampled to $255 \times 300 \times 160$ voxels by using tri-linear interpolation so that the voxel interval became 0.6 mm. We chose 0.6 mm to fit the data to the memory size of our computer (2 GB). In this iteration, we used the batch update method, and it took 150 iterations to obtain a satisfactory result with $\Delta T = 0.2$. However, we need to notice that this kind of resolution change tends to increase




Figure 6: Exceptional treatment of $\{(x,y,z)|S_0 \geq 100.0, c_l < 0.1, c_p < 0.4\}$ area. (a) No diffusion. (b) Pseud-color display of (a). (c) Pseud-color display of Figure 4(k).




Figure 7: 2D cross-section images of a 3D DBT texture. Data size: $256 \times 256 \times 107$. Voxel size: $0.8984375\text{mm} \times 0.8984375\text{mm} \times 0.8984375\text{mm}$. X-axis: right to left. Y-axis: posterior to anterior. Z-axis: superior to inferior.




Figure 8: The DBT texture generated from the super-sampled DT-MRI datasets. Data size: $255 \times 300 \times 160$. Voxel size: $0.6\text{mm} \times 0.6\text{mm} \times 0.6\text{mm}$. Batch update($\Delta T = 0.333$)

the memory consumption and the computation time. Table 1 summarizes the timing results of our experiments.

Figure 9 depicts the cross-section images of the high-resolution DBT texture sliced similar to Figure 7. Fibrous structures are much clearer than those in Figure 7. Figure 10 is the cross-section images of the FA metric, which is often used in the field of medicine, for comparison. One can understand how useful the directional information of the DBT cross-section image is for visualizing the nerve fiber structure. We presented only flat cross-section images of the DBT texture. However, it will be interesting to consider a curved cross-section surface along a certain path such as the visual pathway.

Table 1: Timing results. V : voxel interval (millimeter). U : update method (O : on-line, B : batch), E : eigenvalue decomposition time (seconds), N : number of iterations, DBT : iteration time (seconds). System: Xeon 3.0 GHz; Linux (Fedora Core 1) PC with 2GB RAM.

Datasize	V	U	E	ΔT	N	DBT
$256 \times 256 \times 107$	0.8984375	O	40	1.0	30	13
$256 \times 256 \times 107$	0.8984375	B	40	0.2	200	123
$255 \times 300 \times 160$	0.6	O	77	0.333	75	81
$255 \times 300 \times 160$	0.6	B	78	0.2	150	208



Figure 9: 2D cross-section images of the high-resolution DBT texture. Data size: $255 \times 300 \times 160$. Voxel size: $0.6\text{mm} \times 0.6\text{mm} \times 0.6\text{mm}$. X-axis: right to left. Y-axis: posterior to anterior. Z-axis: superior to inferior.




Figure 10: 2D cross-section images of the FA metric. Data size: $255 \times 300 \times 160$. Voxel size: $0.6\text{mm} \times 0.6\text{mm} \times 0.6\text{mm}$. X-axis: right to left. Y-axis: posterior to anterior. Z-axis: superior to inferior.

Figure 11 illustrates the volume rendered images of the high resolution DBT texture. We used a transfer function that returns the opacity values for the isosurface rendering of the DBT texture (threshold 0.5) and color values from the anisotropy metrics, according to Figure 5. Figure 11(a) is the front view of the complete white matter. The intersection between the corpus callosum and corona radiate and the intersection between pontocerebellar fibers and corticospinal tracts look green because of their high planar diffusion. This corresponds well with the report of Wiegell et al. [WLW00]. Figure 11(b) is the top view, where the fibrous structure of the corpus callosum and a pair of cingulum bundles are easily recognized. Figure 11(c) is the right view of the left hemisphere. The reddish cross-section reveals the strong linear diffusion inside the corpus callosum.

5. Conclusions

In this paper, we proposed a texture-based method, called DBT, for visualizing the white matter connectivity from DT-MRI datasets. The mixed-diffusion model was introduced by decomposing the diffusion tensor into three components and mixing them together to suppress excessive spherical diffusion components while accounting for both linear and planar

diffusion components. The visualization results of the DBT method are consistent with the neuroanatomy.

Since the DBT method can generate a 3D texture by numerically solving a simple PDE, it is suitable for parallel processing and hardware acceleration using GPUs. However, the DBT method needs to keep six tensor elements for each voxel during the iteration, and uses more graphics memory than the GPU implementation of isotropic diffusion, such as reaction-diffusion simulation [HCSL02]. Furthermore, simulations on a GPU use batch updates and need more iterations than on-line updates. In our future work, we are going to implement the DBT method on a graphics PC cluster [MOK*01], and to study how to generate a much finer DBT texture and a better tracking method of much longer neural pathway.

DBT can qualitatively visualize the local connectivity of human white matter. However, medical doctors need quantitative properties of the white matter connectivity for their diagnosis. We hope to prove that the tensor decomposition technique we used in the DBT method is useful also for quantifying the white matter connectivity by using diffusion simulation [OHW02].




Figure 11: Volume rendered images of the high-resolution DBT texture. Data size: $255 \times 300 \times 160$. Voxel size: $0.6\text{mm} \times 0.6\text{mm} \times 0.6\text{mm}$.

Acknowledgements

Part of this work has been financially supported by ACT-JST, and JSPS under Grant-in-Aid for Scientific Research (B) 18300026. We have benefited from a pilot DBT code implemented by Shoko Ando in an early stage of the work.

References

- [BMPL94] BASSER P., MATTIELLO J., PIERPAORI C., LEBIHAN D.: MR diffusion tensor spectroscopy and imaging. *Bio-physical Journal* 66 (1994), 259–267.
- [CL93] CABRAL B., LEEDOM L.: Imaging vector fields using line integral convolution. In *Proc. SIGGRAPH 93* (1993), pp. 263–270.
- [DH93] DELMARCELLE T., HESSELINK L.: Visualizing second-order tensor fields with hyperstreamlines. *IEEE Computer Graphics and Applications* 13, 4 (1993), 25–33.
- [DPR00] DIEWALD U., PREUER T., RUMPF M.: Anisotropic diffusion in vector field visualization on euclidean domains and surfaces. *IEEE Trans. Visualization and Computer Graphics* 6, 2 (2000), 139–149.
- [HCSL02] HARRIS M. J., COOMBE G., SCHEUERMANN T., LASTRA A.: Physically based visual simulation on graphics hardware. In *Proc. SIGGRAPH/EUROGRAPHICS Workshop on Graphics Hardware* (2002), pp. 109–118.
- [Hsu01] HSU E.: Generalized line integral convolution rendering of diffusion tensor fields. In *Proc. International Society for Magnetic Resonance in Medicine, 9th Scientific Meeting and Exhibition* (2001).
- [IG97] INTERRANTE V., GROSCH C.: Strategies for effectively visualizing 3D flow with volume LIC. In *Proc. Visualization '97* (1997), pp. 421–424.
- [KWH00] KINDLMANN G. L., WEINSTEIN D., HART D.: Strategies for direct volume rendering of diffusion tensor fields. *IEEE Trans. Visualization and Computer Graphics* 6, 2 (2000), 124–138.
- [LAKR98] LAIDRAW D. H., AHRENS E. T., KREMERS D., READHEAD C.: Visualizing diffusion tensor imaging of the mouse spinal cord. In *Proc. Visualization '98* (1998), pp. 127–134.
- [MOK*01] MURAKI S., OGATA M., KAJIHARA K., MA K.-L., KOSHIZUKA K., LIU X., NAGANO Y., SHIMOKAWA K.: Next-generation visual supercomputing using PC clusters with volume graphics hardware devices. In *Proc. SC2001* (2001).
- [OHW02] O'DONNELL L., HAKER S., WESTIN C.-F.: New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor warped space. In *Proc. MICCAI 2002* (2002).
- [PB00] PIERPAORI C., BASSER P.: Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine* 36 (2000), 893–906.
- [SEHW02] SIGFRIDSSON A., EBBERS T., HEIBERG E., WIGSTROM L.: Tensor field visualization using adaptive filtering of noise fields combined with glyph rendering. In *Proc. Visualization 2002* (2002), pp. 371–378.
- [SFCN02] SUZUKI Y., FUJISHIRO I., CHEN L., NAKAMURA H.: Hardware-accelerated selective volume rendering of 3D LIC textures. In *Proc. Visualization 2002* (2002), pp. 458–488.
- [WKL99] WEINSTEIN D., KINDLMANN G., LUNDBURG E.: Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *Proc. Visualization '99* (1999), pp. 249–253.
- [WLW00] WIEGELL M. R., LARSSON H. B. W., WEDEEN V. J.: Fiber crossing in human brain depicted with diffusion tensor MR imaging. *Radiology* 217, 3 (2000), 897–903.
- [WPG*97] WESTIN C.-F., PELED S., GUBJARTSSON H., JOLETSZ F., KIKINIS R.: Geometrical diffusion measures for MRI from tensor basis analysis. In *Proc. ISMRM'97* (1997).
- [XZC*99] XUE R., ZHIL P. C. M. V., CRAIN B. J., SOLAIYAPAN M., MORI S.: In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magnetic Resonance in Medicine* 42 (1999), 1123–1127.
- [ZB02] ZHUKOV L., BARR A. H.: Oriented tensor reconstruction: Tracing neural pathways from diffusion tensor MRI. In *Proc. Visualization 2002* (2002), pp. 387–394.
- [ZDL03] ZHANG S., DEMIRALP C., LAIDLAW D. H.: Visualizing diffusion tensor MR images using streamtubes and stream-surfaces. *IEEE Trans. Visualization and Computer Graphics* 9, 4 (2003), 454–462.
- [ZP03] ZHENG X., PANG A.: HyperLIC. In *Proc. Visualization 2003* (2003), pp. 249–256.