

Tensorlines: Advection-Diffusion based Propagation through Diffusion Tensor Fields

David Weinstein, Gordon Kindlmann, Eric Lundberg

Center for Scientific Computing and Imaging
Department of Computer Science
University of Utah
{dmw|gk|lundberg}@cs.utah.edu

Abstract

Tracking linear features through tensor field datasets is an open research problem with widespread utility in medical and engineering disciplines. Existing tracking methods, which consider only the preferred local diffusion direction as they propagate, fail to accurately follow features as they enter regions of local complexity. This shortcoming is a result of *partial voluming*; that is, voxels in these regions often contain contributions from multiple features. These combined contributions result in ambiguities when deciding local primary feature orientation based solely on the preferred diffusion direction. In this paper, we introduce a novel feature extraction method, which we term *tensorline* propagation. Our method resolves the above ambiguity by incorporating information about the nearby orientation of the feature, as well as the anisotropic classification of the local tensor. The nearby orientation information is added in the spirit of an advection term in a standard diffusion-based propagation technique, and has the effect of stabilizing the tracking. To demonstrate the efficacy of tensorlines, we apply this method to the neuroscience problem of tracking white-matter bundles within the brain.

1 Introduction

Diffusion tensor magnetic resonance imaging (DT-MRI) is becoming a more common and useful medical research tool as magnetic resonance scanners become capable of delivering ever higher resolution scans. As with increased exposure of any new type of data, new applications arise which result in a need for new visualization methods. As a particular example, neuroscientists are interested in visualizing diffusion tensor data within the cranial volume in order to examine the connectivity between different regions of the brain and ultimately to better understand the brain's functional organization [9].

DT-MRI data describes the way water diffuses through a volume [1]. Since white matter tends to diffuse water along the fiber directions, diffusion tensor data should be a strong indicator of white matter orientation, and therefore of cognitive functional organization.

A method akin to vector field streamline advection should be a strong candidate for such neuro-connectivity studies. By seeding a streamline in a particular region of the brain, a neuroscientist might be able to see what other regions of the brain are anatomically attached through gray and white matter paths. If we had very high resolution data, capable of imaging anatomy at the scale of individual fibers, then such a visualization might be straightforward.

Unfortunately, clinical neural diffusion tensor data are still relatively low resolution and are also rather noisy. When a region of white matter is measured at the resolution of today's MR scanners, they are sampling a volume composed of many fibers, which can have a range of orientations. As we shall discuss, it is precisely this

partial voluming that makes tracking fibers through diffusion tensor fields unstable.

2 Background

2.1 Diffusion Tensors

A diffusion tensor can be represented mathematically as a 3×3 symmetric, semi-positive definite matrix. By examining the eigenvectors and eigenvalues of this matrix, we can gain intuition about the tensor.

We will refer to the eigenvalues as λ_1 , λ_2 and λ_3 , and their corresponding eigenvectors as \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 , where $\lambda_1 \leq \lambda_2 \leq \lambda_3$. Scaling \mathbf{e}_1 by λ_1 , we obtain the *major* eigenvector, \mathbf{v}_1 . Similarly, if we scale \mathbf{e}_2 and \mathbf{e}_3 by λ_2 and λ_3 , we obtain the *medium* and *minor* eigenvectors, \mathbf{v}_2 and \mathbf{v}_3 , respectively. If we imagine an ellipsoid, with axes corresponding to \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 , we have a geometric representation which uniquely defines the diffusion tensor and intuitively represents the probable shape a water droplet would take as it was diffused by the tensor.

A diffusion tensor that has only one large eigenvalue is said to be *linearly isotropic* [16], and has a cigar-shaped geometric representation. This classification is common within white matter regions of the brain, as the myelin sheaths cause water to diffuse preferentially along their axonal lengths. When a tensor has two relatively large eigenvalues, it preferentially diffuses water in the plane spanned by \mathbf{e}_1 and \mathbf{e}_2 . This case is referred to as *planar anisotropy* and is geometrically represented by a pancake shape. The final classification, *spherical anisotropy*, refers to tensors for which all three eigenvalues are approximately equal. In this case, the tensor diffuses relatively uniformly in all directions.

It is worth noting that planar anisotropy can result from either material which diffuses in sheets or from partial voluming effects as two linear features cross near each other [11]. Similarly, very complex regions with many different fiber orientations can result in an isotropic tensor.

In real data, there are not just three unique tensor values (linear, planar, spherical), but rather there is a continuum of classifications within these three extremes. To quantify this continuum, Westin [16] introduced three definitions, corresponding to linear (c_l), planar (c_p) and spherical (c_s) anisotropy coefficients:

$$c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3}, \quad (1)$$

$$c_p = \frac{2(\lambda_2 - \lambda_3)}{\lambda_1 + \lambda_2 + \lambda_3}, \quad (2)$$

$$c_s = \frac{3\lambda_3}{\lambda_1 + \lambda_2 + \lambda_3}, \quad (3)$$

By design, these three coefficients sum to unity. This makes

them amenable as barycentric coordinates [7], as shown in Figure 1. In this figure, we see a histogram of the c_l , c_p , and c_s coordinates for all tensors in the field. Most interestingly, there is no clustering within the histogram. This suggests that there is substantial partial voluming taking place throughout the measured volume.




Figure 1: Barycentric histogram of a diffusion tensor MRI dataset. The coordinates correspond to the amount of linear, planar and spherical anisotropy in the tensor. The lack of clustering suggests considerable partial voluming is taking place within the volume.

2.2 Previous Work

Traditional methods for visualizing tensor data include brush strokes [8], glyphs [2, 5, 10], ellipsoids [6, 13, 14], stream-polygons [12] and hyperstreamlines [3, 4]. Using a 2D approach, Laidlaw [8] uses brush strokes to represent the diffusion tensor data through various stroke shape, color and texture cues. This method works well in 2D, but it is not clear how it can be effectively extended to 3D. A more 3D-oriented approach is the use of glyphs. Glyphs are geometric icons which are useful for depicting the tensors at particular locations, or uniformly distributed through the domain. Ellipsoids are a specific type of glyph, which geometrically represents the diffusion tensor by rendering the deformation a sphere would undergo when acted upon by the tensor matrix. As with all glyphs, though, a 3D field full of ellipsoids rapidly becomes cluttered and uninformative. In contrast to these other methods, which show global information about the volume, hyperstreamlines and stream-polygons show local information about the field along a particular path through the volume. In the case of neural connectivity, it is actually the path itself which is of primary interest. Stream-polygons string together polygons depicting local field properties. Similarly, hyperstreamlines depict flow paths by propagating particles along the major eigenvector direction (the direction of most likely diffusion) for each tensor, and stretching a cylinder about that core. Other properties of the field can then be encoded onto the cylinder; for example, the transverse shape of the cylinder can be warped to conform to the ellipsoid spanned by the medium and minor eigenvectors. In this way, hyperstreamlines can completely represent the six degrees of freedom encoded in the diffusion tensor.

While hyperstreamlines provide a powerful method for tensor visualization, the propagation path itself is not always ideal. Whereas streamlines produced the path indicated in light gray in Figure 2, which can “get lost” in isotropic regions, our method produced the




Figure 2: Visualization of hyperstreamline and tensorline propagation through non-isotropic (left, right and bottom) and isotropic regions (top middle) in synthetic data. The directions of the major eigenvectors are indicated with arrows. The hyperstreamline core is shown with the light gray line, and the tensorline is shown in dark gray. Note that the tensorline continues along its present course, despite encountering a region of isotropic diffusion.

path of the dark gray *tensorline* from that same image - a path that continued along its present course when it encountered a region of isotropic diffusion.

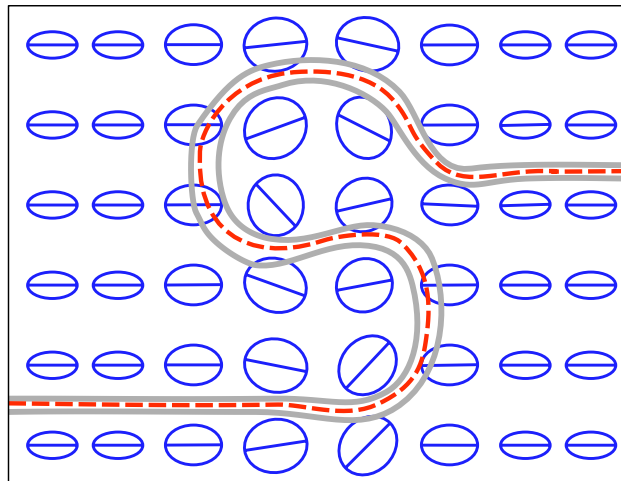


Figure 3: Hyperstreamline advection through a nearly isotropic region. The S-shaped path is an artifact of the noise in the region and conveys misleading information about the direction of flow through the field.

In Figure 3, we see one example of how propagating a streamline according to the \mathbf{v}_1 direction of the tensors can be misleading. For explanatory simplicity, we have displayed a two-dimensional slice of data with two-dimensional diffusion tensors indicated by the ellipses. In this figure we see a field which varies from being somewhat linearly anisotropic on the left and right, to nearly isotropic in the middle. We have indicated the \mathbf{v}_1 directed streamline with a dotted line and the borders of a hyperstreamline with gray. What is not clear from either of these visualizations is that diffusion near the middle of the frame is nearly isotropic. The S-shaped path through the middle of the field is not indicative of complex structure in the

data, but might be an artifact of measurement noise or partial voluming. However, it is in no way obvious to the viewer that this is the case.

Extending our visualization to three dimensions, we can have confusion in isotropic regions as well as in planar anisotropic regions. In the case of planar anisotropy, the confusion is a result of the field having nearly equivalent first and second eigenvalues and, therefore, major and medium eigenvectors that are only meaningful insofar as they span a particular plane, but not in their particular orientations within that plane. Similarly, for isotropic regions, none of the eigenvector directions are individually meaningful.

In this paper, we present an alternative propagation method which does not suffer from the misleading effects described above. We term this visualization method *tensorlines*. Tensorlines are described in detail in Section 3. In Section 4, we show the results of applying tensorlines to neural diffusion tensor data. We conclude in Section 5 with a summary and a discussion of possible extensions.

3 Methods

Physically, diffusion is a probabilistic phenomenon. A diffusion tensor specifies the probability density function (pdf) of where a particle's Brownian motion will move it over time [1]. To visualize this pdf, we cover the unit sphere with dots, as is shown for the unit circle in Figure 4. All of the points on the unit sphere (circle) are then transformed by the diffusion tensor matrix, resulting in an ellipsoid (ellipse). Furthermore, the resulting distribution of dots on that ellipsoid (ellipse) corresponds to the probable distribution of particles as they diffuse from the origin. As we can see in Figure 4, the dots have a higher density in the \mathbf{v}_1 direction, and a lower density in \mathbf{v}_2 . This corresponds to the higher likelihood that a particle will be diffused in the \mathbf{v}_1 direction, and a lower likelihood that it will be diffused in some other direction.




Figure 4: Redistribution of uniform sample resulting from anisotropic diffusion tensor. Particles have a higher probability of being diffused in the \mathbf{v}_1 (major eigenvector) direction of the ellipsoid.




Figure 5: Redistribution of uniform sample resulting from nearly isotropic diffusion tensor. Probability of particle diffusion is approximately the same in every direction.

However, if we look at the case in Figure 5, the dots on the diffusion ellipsoid (ellipse) are nearly uniformly distributed. This case




Figure 6: Remapping of unit vector through four different tensor matrices. Each row gives the geometric interpretation of applying a different tensor matrix (left matrix and left ellipsoid) to the same example vector (middle vector and middle circle) and the resultant transformed vector (right vector). The gray axis of the ellipses is the direction the major eigen-axis and is the direction in which a hyperstreamline would be propagated; in contrast, the dark arrow on the right is the diffusion modulated direction. Our method uses a combination of these two terms to produce more stable propagation paths through isotropic regions.

is representative of any nearly isotropic diffusion tensor. In contrast to the previous example, here there is approximately the same probability for particle diffusion in *any* direction. In such a case, choosing the major eigenvector as the diffusion direction is very much an arbitrary decision. This instability is also depicted in Figure 6, where we are examining the effects of applying various 2D tensor matrices to a unit vector. In each row of the image, we apply a different tensor. The matrix is given on the left, then the corresponding ellipse is shown (with the major eigenvector axis indicated in gray), followed by the vector upon which the matrix will operate, and finally the transformed vector is shown on the right. We also indicate the effects of the transformation on the x and y unit vectors, depicted in dashed lines. We note that for the last row, the diffusion tensor is nearly isotropic (as evidenced by the nearly circular ellipse). The vector we are transforming by this matrix is almost orthogonal to the first eigenvector. However, since the tensor is nearly isotropic, the output vector on the right is only slightly rotated from its initial position shown in the middle.

3.1 Propagation

If we follow an individual particle's path as it moves through the volume (being probabilistically diffused as it travels), we get a streamline traced through the field. Delmarcelle's method propagates these hyperstreamlines by always choosing to diffuse the particle in the direction of the major eigenvector of the tensor through which it is traveling. Delmarcelle has thus reduced the problem of advecting through a tensor field to the problem of advecting through the \mathbf{v}_1 vector field. While the streamline paths generated with this method are, in fact, the most likely pure diffusion paths, they can at times be misleading, as we saw in Figure 3.

Delmarcelle's hyperstreamlines assume a pure diffusion model;

however, in regions of the data with planar and spherical anisotropy, the first principal component is a rather arbitrary direction. These ambiguities result in unstable propagation. Our method stabilizes the propagating by incorporating two additional to propagate through ambiguous regions. Because these terms supplement the diffusion motion with a velocity term, we refer to them as *advection* vectors, in the spirit of advection-diffusion particle physics.

3.2 Implementation

The “advection” vector used to stabilize propagation is combination of two vectors, with relative weightings chosen by the user. These vectors correspond to the *incoming* direction \mathbf{v}_{in} (the direction of the previous propagation step), and the *outgoing* direction \mathbf{v}_{out} (the incoming vector, transformed by the tensor matrix). Specifically, we compute:

$$\mathbf{v}_{out} = \mathbf{D}\mathbf{v}_{in}, \quad (4)$$

where \mathbf{D} is the diffusion tensor matrix. We note that as a preprocess, we scale our diffusion matrix \mathbf{D} by $\frac{2}{e_{max}}$, where e_{max} is the largest eigenvector in our field. This scaling has the effect of normalizing the diffusion term to be to scale with the advection terms.

As mentioned above, our propagation direction is a combination of \mathbf{v}_1 , \mathbf{v}_{in} and \mathbf{v}_{out} . The way in which these vectors are combined to determine the next propagation step vector, \mathbf{v}_{prop} , depends on the shape of the local tensor:

$$\mathbf{v}_{prop} = c_l \mathbf{v}_1 + (1 - c_l)((1 - w_{punct})\mathbf{v}_{in} + w_{punct}\mathbf{v}_{out}), \quad (5)$$

where w_{punct} is a user-controlled parameter (described below) and c_l is the linear anisotropy coefficient of the local tensor.

The equation above was chosen because it satisfies the following desirable conditions:

Anisotropy	Direction In	Desired Out
Linear	Any	e_1
Planar	Tangential to disk	\mathbf{v}_{in} or \mathbf{v}_{out}
Planar	Normal to disk plane	\mathbf{v}_{out}
Spherical	Any	\mathbf{v}_{in} or \mathbf{v}_{out}

The first and last rows of this table are straightforward. In contrast, the second and third rows describe what should happen in regions of planar anisotropy. That is, these two cases cannot be disambiguated based on anisotropy type, and so we resort to a user-controlled parameter w_{punct} . This coefficient can take on values from 0 to 1, and affects how much the propagation should be encouraged to “puncture” through planar tensors oriented normal to its path, versus turning into the plane. This property depends largely on the type of data being investigated, which is why it has been left as a user-definable coefficient. For example, when identifying white matter association tracts, a puncture coefficient of 0.20 worked well in practice, as shown in Figure 2. For more rigid datasets, the appropriate coefficient choice would likely be somewhat higher.

Discretely propagating along the \mathbf{v}_{prop} tensorlines, we generate different paths than we did by advecting through the \mathbf{v}_1 velocity field. This is illustrated in Figure 7, where we revisit the case shown in Figure 3. Now we have added solid tensorlines, and we note that where the hyperstreamlines wandered through the isotropic region in the middle, the tensorlines continued straight through, with only minor fluctuations.

¹Since \mathbf{v}_1 and $-\mathbf{v}_1$ are both valid eigenvectors, we can avoid “doubling back” on ourselves as we propagate by always checking to see if $(\mathbf{v}_{in} \cdot \mathbf{v}_1) < 0$, and if it is, we negate \mathbf{v}_1 .




Figure 7: Comparison of tensorlines (solid) and hyperstreamline (core is dashed, borders are gray). Note the tensorlines continue with only minor fluctuations through the isotropic region in the center of the figure, whereas the hyperstreamline is diverted into an S-shaped path.

4 Results




Figure 8: The six diffusion tensor matrix coefficients for a single slice of neural diffusion tensor MRI data.

In the previous sections, we have shown simple images to illustrate indicate the theoretic efficacy of tensorline propagation. In this section we show a tensorline visualization of actual diffusion tensor MRI data. The dataset is a $128 \times 128 \times 60$ volumetric dataset, with a diffusion tensor matrix at each voxel. An image of the six components for one slice of the data is shown in Figure 8. The 60 slices extend from the tops of the eyes at the bottom, to the top of the cortical surface at the top.




Figure 9: Visualization of tensorlines (in yellow) and hyperstreamlines (cyan) as they spread out from the corpus callosum through the coronal radiate. Cutting planes of the tensor volume are shown for reference, along with an isosurface indicating the corpus callosum and cingulum bundles.

In Figure 9, as well as the color plate images, we see bundles of the coronal radiata and sagittal stratum emanating from the internal capsule. Tensorlines are shown in yellow, and hyperstreamlines are shown in cyan. Cutting planes of the linear anisotropy of the tensor field are shown in gray-scale for reference, along with a gray isosurface of that volume, indicating the corpus callosum and cingulum bundle. Because they are not as susceptible to planar anisotropy, the tensorlines do a better job of tracking the white matter fibers as they rise from the deeper brain structures up toward the cortical surface.

5 Conclusions and Future Work

In this paper, we have introduced a novel propagation method, called tensorlines, for visualizing diffusion tensor fields. Tensorlines extend the traditional propagation methods in order to stabilize propagation through regions with non-linear preferential diffusion.

In the future, we could like to generate surfaces around tensorlines in a way similar to Delmarcelle's method for surfaces about hyperstreamlines [3]. We could encode the deviation from \mathbf{v}_1 as a texture on the surface, providing the user with a visual clue as to the difference between the various vectors being weighted in Equation 5.

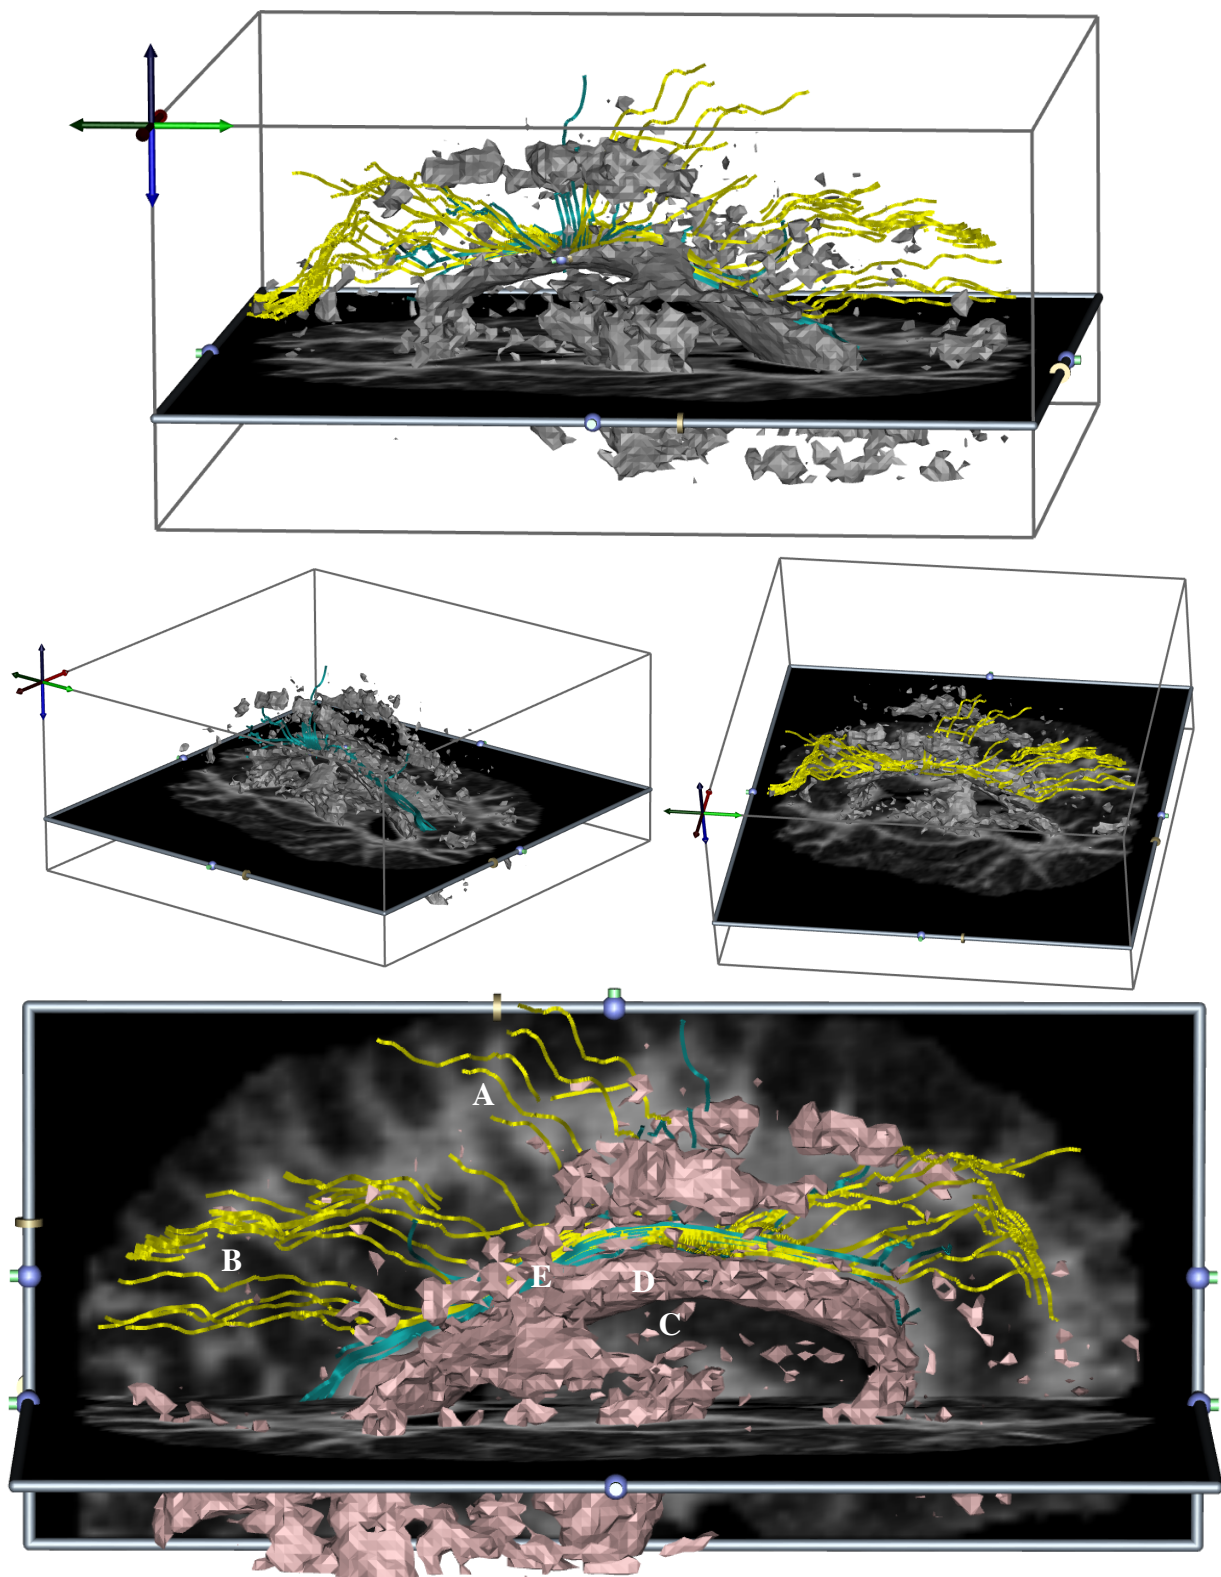
We are also interested in stabilizing advection by propagating groups of streamlines together as a cohesive bundle [15]. Our earlier work in this area focused on bundles which advected through flow fields. In the future, we would like to extend these ideas to also apply to tensor fields.

6 Acknowledgments

This work was supported in part by awards from the Department of Energy and the National Science Foundation. The authors would like to thank Chris Johnson, Chuck Hansen and Matthew Bane for their valuable comments and suggestions, as well as Helen Hu and Cameron Christensen for their help with the figures. We also gratefully acknowledge Andy Alexander at the University of Utah's Center for Advanced Medical Technologies for providing the neural diffusion tensor dataset.

References

- [1] P.J. Basser, J. Mattiello, and D. Le Bihan. Estimation of the effective self-diffusion tensor from the NMR spin-echo. *Magnetic Resonance*, pages 247–254, 1994.
- [2] W.C. de Leeuw and J.J. van Wijk. A probe for local flow field visualization. In *IEEE Visualization 93 Proceedings*, pages 39–45, 1993.
- [3] T. Delmarcelle and L. Hesselink. Visualization of second order tensor fields and matrix data. In *IEEE Visualization 92 Proceedings*, pages 316–323, 1992.
- [4] T. Delmarcelle and L. Hesselink. Visualizing second-order tensor fields with hyper streamlines. *IEEE Computer Graphics and Applications*, pages 25–33, 1993.
- [5] R.B. Haber. Visualization techniques for engineering mechanics. *Computing Systems in Engineering*, pages 37–50, 1990.
- [6] G.D. Kerlick. Moving iconic objects in scientific visualization. In *IEEE Visualization 90 Proceedings*, pages 124–130, 1990.
- [7] G.L. Kindlmann and D.M. Weinstein. Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization 99 Proceedings*, 1999.
- [8] 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 *IEEE Visualization 1998 Proceedings*, pages 127–134, 1998.
- [9] N. Makris, A.J. Worth, G. Sorensen, G.M. Papadimitriou, O. Wu, T.G. Reese, V.J. Wedeen, T.L. Davis, J.W. Stakes, V.S. Caviness, E. Kaplan, B.R. Rosen, D.N. Pandya, and D.N. Kennedy. Morphometry of *in vivo* human white matter association pathways with diffusion weighted MRI. *Annals of Neurology*, 42(6):951–962, 1997.
- [10] F.J. Post, T. van Walsum, Post F.H., and D. Silver. Iconic techniques for feature visualization. In *IEEE Visualization 95 Proceedings*, pages 288–295, 1995.
- [11] C. Poupon, J.F. Mangin, V. Frouin, J. Régis, F. Poupon, M. Pachot-Clouard, D. Le Bihan, and I. Bloch. Regularization of MR diffusion tensor maps for tracking brain white matter bundles. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 489–498, 1998.
- [12] W.J. Schroeder, C.R. Volpe, and W.E. Lorensen. The stream polygon: A technique for 3D vector field visualization. In *IEEE Visualization 91 Proceedings*, pages 126–132, 1991.
- [13] D. Silver, N. Zabusky, V. Fernandez, and M. Gao. Ellipsoidal quantification of evolving phenomena. *Scientific Visualization of Natural Phenomena*, pages 573–588, 1991.
- [14] C. Upson, R. Wolff, R. Weinberg, and D. Kerlich. Two and three dimensional visualization workshop. In *Course Number 19, SIGGRAPH 89*, 1989.
- [15] D.M. Weinstein. Stream bundles - cohesive advection through flow fields. Technical Report UUCS-99-005, University of Utah, Department of Computer Science, Salt Lake City, UT, 1999.
- [16] C-F. Westin, S. Peled, H. Gubjartsson, R. Kikinis, and F.A. Jolesz. Geometrical diffusion measures for MRI from tensor basis analysis. In *Proceedings of ISMRM*, 1997.



Bundles of the coronal radiata (a) and sagittal stratum (b) emanating from the internal capsule (c). Tensorlines are shown in yellow, and hyperstreamlines are shown in cyan. Cutting planes of the linear anisotropy of the tensor field are shown in gray-scale for reference, along with a gray isosurface of that volume, indicating the corpus callosum (d) and cingulum bundle (e). Because they are not as susceptible to planar anisotropy, the tensorlines do a better job of tracking the white matter fibers as they rise from the deeper brain structures up toward the cortical surface.