

Abstract of “Revealing White Matter Fiber Structure with Diffusion Imaging” by Song Zhang, Ph.D., Brown University, August 2006.

The complexity of the brain has made the study of neuroanatomy a notoriously demanding task. Diffusion magnetic resonance imaging probes fiber structures *in vivo* by measuring the Brownian motion of water molecules. Tractography generates curves from the diffusion imaging data that correlate with the neural fiber pathways. The resulting high-dimensional data sets present a two-fold challenge in both visualization and data analysis. This dissertation presents several diffusion MRI visualization and analysis methods for revealing the white matter structures in the brain. I present a model that describes linear and planar features of the brain anatomy respectively. I discuss strategies in seeding, model placement, display environments and interaction. To further explore the structures in the model, I emulate the way neuroscientists conceptualize neural fibers by grouping closely-related fibers into fiber bundles and studying the consistency, anatomical meaning and other properties of these fiber bundles. Apart from qualitative exploration, I also present quantitative methods such as metrics on tractography models or forward model for the tractography models and employ these methods to study white matter integrity and increase the accuracy of the tractography models. I applied these visualization and analysis methods in biomedical applications such as the studies on multiple sclerosis, brain tumor, and lemur brain. The work in this dissertation advances both qualitative visualization and quantitative analysis of the 3D diffusion imaging models.

Revealing White Matter Fiber Structure with Diffusion Imaging

by

Song Zhang

B. S., Nankai University, 1996

Sc. M, Brown University, 2000

Ph.D. Thesis

Providence, Rhode Island

August 2006

This thesis by Song Zhang is accepted in its present form by
the Department of Computer Science as satisfying the dissertation requirement
for the degree of Doctor of Philosophy.

Date _____
_____ David H. Laidlaw, Director

Recommended to the Graduate Council

Date _____
_____ John Hughes, Reader

Date _____
_____ Pascal Van Hentenryck, Reader

Approved by the Graduate Council

Date _____
_____ Karen Newman
Dean of the Graduate School

Vita

Education

- Ph.D., Computer Science, Brown University, expected 2006.
Revealing White Matter Fiber Structure with Diffusion Imaging
- M.Sc., Computer Science, Brown University, 2000.
Visualizing Diffusion Tensor MRI Using Streamtubes and Streamsurfaces
- B.Sc., Computer Science, Nankai University, China, 1996.

Experience

- Research Assistant, Brown University Computer Science Department
August 1998 – August 2006
- Teaching Assistant, Brown University Computer Science Department
January – May 2006, Software System Design
- Teaching Assistant, Brown University Computer Science Department
January – May 2005, Software System Design
- Teaching Assistant, Brown University Computer Science Department
January – May 2002, Computer Networks
- Internship, Microsoft (Summer 2001), Automatic Map Label Layout

Honors and Awards

- Burroughs Welcome Fund Fellowship: 2004-2005
- Nankai University Scholarship: 1994-1995

Publications

Book Chapters

- A. Vilanova, Song Zhang, G. Kindlmann, and David H. Laidlaw. **An introduction to visualization of diffusion tensor imaging and its applications.** In *Visualization and Image Processing of Tensor Fields*. Springer-Verlag, 2005.
- Song Zhang, Gordon Kindlmann, and David H. Laidlaw. **Diffusion tensor MRI visualization.** In *Visualization Handbook*. Academic Press, June 2004.

Refereed Journal Articles

- Robert H. Paul, David H. Laidlaw, David F. Tate, Stephanie Lee, Karin Hoth, John Gunstad, Song Zhang, Ronald A. Cohen, and Tim Flanigan. **Neuropsychological and neuroimaging outcome of HIV-associated progressive multifocal leukoencephalopathy in the era of ART: a case report.** *HIV Medicine*, 2006. In press.
- Jack H. Simon, Song Zhang, David H. Laidlaw, David E. Miller, Mark Brown, John Corboy, and Jeffrey Bennett. **Identification of fibers at risk for degeneration by diffusion tractography in patients at high risk for MS after a clinically isolated syndrome.** *Journal of Magnetic Resonance Imaging*, 2005. In press.

- Cagatay Demiralp, Cullen Jackson, David Karelitz, Song Zhang, and David H. Laidlaw. **A qualitative and quantitative comparison of cave and fish-tank virtual reality displays.** *IEEE Transaction on Visualization and Computer Graphics*, 2005. In press.
- Jason A. Kaufman, Eric T. Ahrens, David H. Laidlaw, Song Zhang, and John M. Allman. **Anatomical analysis of an aye-aye brain (*daubentonia madagascariensis*, primates: Prosimii) combining histology, structural magnetic resonance imaging, and diffusion-tensor imaging.** *Anatomical Record*, 287A(1):1026–1037, November 2005.
- Song Zhang, Mark Bastin, David H. Laidlaw, Saurabh Sinha, Paul A. Armitage, and Thomas S. Deisboeck. **Visualization and analysis of white matter structural asymmetry in diffusion tensor MR imaging data.** *Magnetic Resonance in Medicine*, 51(1):140–147, 2004.
- Andreas Wenger, Daniel Keefe, Song Zhang, and David H. Laidlaw. **Interactive volume rendering of thin thread structures within multivalued scientific datasets.** *IEEE Transactions on Visualization and Computer Graphics*, 10(6):664–672, November/December 2004.
- Song Zhang, Cagatay Demiralp, and David H. Laidlaw. **Visualizing diffusion tensor MR images using streamtubes and streamsurfaces.** *IEEE Transactions on Visualization and Computer Graphics*, 9(4):454–462, October 2003.

Reviewed Conference Papers and Abstracts

- Stephanie Y. Lee, Stephen Correia, David F. Tate, Robert H. Paul, Song Zhang, Stephen P. Salloway, Paul F. Malloy, and David H. Laidlaw. **Quantitative**

tract-of-interest metrics for white matter integrity based on diffusion tensor MRI data. *ISMRM*, 2006.

- Song Zhang, Stephen Correia, David F. Tate, and David H. Laidlaw. **Correlating DTI fiber clusters with white matter anatomy.** *ISMRM*, 2006. Under Review.
- Stephen Correia, S. Lee, Paul Malloy, N. Mehta, Song Zhang, Steven Salloway, and David H. Laidlaw. **Diffusion-tensor MRI tractography methods for assessing white matter health and its relationship to cognitive functioning.** In *Proceedings of the International Neurophysiological Society*, Boston, February 2006.
- Stephen Correia, Thea Brennan-Krohn, Song Zhang, David H. Laidlaw, Paul Malloy, and Steven Salloway. **Diffusion-tensor imaging and executive function in subcortical ischemic vascular disease and mild cognitive impairment.** In *International Neuropsychological Society Conference*, February 2005.
- Stephen Correia, Thea Brennan-Krohn, Erin Schlicting, Song Zhang, David H. Laidlaw, Paul Malloy, and Steven Salloway. **Diffusion-tensor imaging in vascular cognitive impairment and mild cognitive impairment: relationship with executive functioning.** In *Proceedings of the 2nd Congress of the International Society for Vascular, Cognitive and Behavioural Disorders*, Florence, Italy, June 2005.
- Jack H. Simon, Song Zhang, David H. Laidlaw, D. E. Miller, M. Brown, J. Corboy, D. Singel, and J. Bennett. **Strategy for detecting neuronal fibers at risk for neurodegeneration in earliest MS by streamtube tractography at 3T.** In *Proceedings of ISMRM*, Miami, FL, May 2005.

- Jack H. Simon, John Corboy, Song Zhang, and David H. Laidlaw. **Visualization of fibers at risk for neuronal tract degeneration in early MS by streamtube diffusion tractography.** Platform Presentation at the Academy of Neurology, Miami, FL, April 2005.
- David Tate, J. Gunstad, Robert Paul, Song Zhang, David H. Laidlaw, and Ronald Cohen. **Fractional anisotropy in subcortical white matter regions of interest among cardiovascular patients.** In *Proceedings of ISMRM*, Miami, FL, May 2005.
- Song Zhang and David H. Laidlaw. **DTI fiber clustering and cross-subject cluster analysis.** In *Proceedings of ISMRM*, Miami, FL, May 2005.
- Daniel Acevedo, Song Zhang, David H. Laidlaw, and Chris Bull. **Color rapid prototyping for diffusion tensor MRI visualization.** MICCAI 2004 Short Papers, September 2004.
- Eric T. Ahrens, J. M. Allman, E. Bush, David H. Laidlaw, and Song Zhang. **Comparative 3d anatomy of the prosimian brain: DTI and histological studies.** In *Proceedings of the Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*, 2004.
- Thea Brennan-Krohn, Stephen Correia, Song Zhang, David H. Laidlaw, Paul Malloy, and Steven Salloway. **Diffusion-tensor imaging and executive function in subcortical vascular disease.** In *Workshop on aging connections: Advanced MRI of age-related white matter changes in the brain*. ISMRM, 2004.
- Stephen Correia, Thea Brennan-Krohn, Song Zhang, David H. Laidlaw, Paul Malloy, and Steven Salloway. **Diffusion-tensor imaging and executive function in subcortical ischemic vascular disease and mild cognitive**

impairment. In *Proceedings of the International Neurophysiological Society*, 2004.

- Steve Correia, Song Zhang, David H. Laidlaw, Paul Malloy, and Steven Salloway. **Diffusion-tensor imaging: linear, planar, and spherical diffusion in CADASIL.** In *9th International Conference on Alzheimer's Disease and Related Disorders*, 2004.
- David H. Laidlaw, Song Zhang, Mark Bastin, S. Correia, Steven Salloway, and Paul Malloy. **Ramifications of isotropic sampling and acquisition orientation on DTI analyses.** In *Proceedings of the Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*, 2004.
- Vadim Slavin, David H. Laidlaw, Song Zhang, Rovert Pelcovits, George Lorient, and Anderw Callan-Jones. **Visualization of topological defects in nematic liquid crystals using streamtubes, streamsurfaces and ellipsoids.** In *IEEE Visualization 2004 Poster Compendium*, October 2004.
- Song Zhang, Jack H. Simon, David H. Laidlaw, Mark Brown, and David M. Miller. **Visualization of the interaction of multiple sclerosis lesions with adjacent white matter fibers using streamtubes and streamsurfaces.** In *IEEE Visualization Poster Compendium*, October 2004.
- Song Zhang and David H. Laidlaw. **DTI fiber clustering in the whole brain.** In *IEEE Visualization 2004 Poster Compendium*, October 2004.
- Song Zhang and David H. Laidlaw. **A model for some subcortical DTI planar and linear anisotropy.** MICCAI 2004, September 2004.
- Cagatay Demiralp, David H. Laidlaw, Cullen Jackson, Daniel Keefe, and Song Zhang. **Subjective usefulness of CAVE and fish tank VR display systems for a scientific visualization application.** In *IEEE Visualization*

Poster Compendium, Seattle, WA, 2003.

- Marco DaSilva, Song Zhang, Cagatay Demiralp, and David H. Laidlaw. **Visualizing the differences between diffusion tensor volume images.** In *Proceedings ISMRM Workshop in Diffusion MRI: Biophysical Issues*, pages 237–238, March 2002.
- Song Zhang, David H. Laidlaw, Mark Bastin, Saurabh Sinha, and Thomas S. Deisboeck. **Application of DTI visualization and analysis on a data set from a brain tumor patient.** In *IEEE Visualization Poster Compendium*, October 2002.
- Marco DaSilva, Song Zhang, Cagatay Demiralp, and David H. Laidlaw. **Visualizing diffusion tensor volume differences.** In *Visualization '01 Work in Progress Proceedings*, pages 16–17, October 2001.
- Song Zhang and David H. Laidlaw. **Elucidating neural structure in diffusion tensor MRI volumes using streamtubes and streamsurfaces.** In *Proc. 9th International Society of MR in Medicine*, April 2001.
- Song Zhang, Cagatay Demiralp, Daniel Keefe, Marco DaSilva, Benjamin D. Greenberg, Peter J. Basser, Carlo Pierpaoli, E. A. Chiocca, T. S. Deisboeck, and David H. Laidlaw. **An immersive virtual environment for DT-MRI volume visualization applications: a case study.** In *Proceedings of IEEE Visualization 2001*, pages 437–440, October 2001.
- Song Zhang, Cagatay Demiralp, Marco DaSilva, Daniel Keefe, David H. Laidlaw, Benjamin D. Greenberg, Peter J. Basser, Carlo Pierpaoli, E.A. Chiocca, and T. S. Diesboeck. **Toward application of virtual reality to visualization of DT-MRI volumes.** In *Proceedings MICCAI*, October 2001.

Acknowledgments

I would like to thank my advisor David H. Laidlaw for all these years of support and guidance. I would like to thank my committee members John F. Hughes and Pascal Van Hentenryck for many inspiring discussions.

I would like to thank my collaborators. Steve Correia, David Tate, Benjamin D. Greenberg, Robert Paul, Paul Malloy and Steve Salloway from the Brown Medical School; Mark Bastin, Saurabh Sinha and Paul A. Armitage from the University of Edinburgh; Jason A. Kaufman and John M. Allman from Caltech; Jack H. Simon, John Corboy, David E. Miller and Mark Brown from the University of Colorado — Denver; Anna Vilanova from Eindhoven University of Technology; Gordon Kindlmann and Thomas S. Deisboeck from Harvard Medical School; Peter J. Basser and Carlo Pierpaoli from the National Institutes of Health.

I would like to thank my friends and fellow graduate students at Brown CS for keeping my spirit up during all these years, particularly Aris Anagnostopoulos, Çağatay Demiralp, Elisabeta Marai, Ye Sun, Dan Keefe, Dan Acevedo, Ying Xing, Srikanta Tirthapura, Lijuan Cai and Wenjin Zhou.

Thanks to Shiyin Wang for the editing work on this dissertation.

Thanks to NIH, NSF and the Brown Brain Science Program for their fundings.

Finally I would like to thank my parents Zongdi Zhang and Jingzhen Wu for everything.

Contents

List of Tables	xv
List of Figures	xvii
1 Introduction	1
1.1 Overview	1
1.2 Contributions	4
2 An Introduction to Diffusion Imaging	6
3 Streamtubes and Streamsurfaces Model	10
3.1 Related Work	11
3.2 Streamtubes and Streamsurfaces	14
3.2.1 Anisotropy Classification	15
3.2.2 Streamtubes for Linear Anisotropy	16
3.2.3 Streamsurfaces for Planar Anisotropy	21
3.2.4 Anatomical Landmarks	26
3.2.5 Results and Discussion	27
4 Visualization Schemes	32
4.1 Virtual Reality	33
4.2 Streamtubes Visualization with Volume Rendering	35
4.2.1 Related Work	39

4.2.2	A Layered Volume-Rendering Framework	42
4.2.3	Threads and Halos	43
4.2.4	Layering Volumes	46
4.2.5	Exploratory Culling and Transfer Function Manipulation . . .	47
4.2.6	Results and Discussion	49
4.3	Rapid Prototyping	51
4.4	Conclusions	53
5	Fiber Bundle Model	55
5.1	Introduction	56
5.2	Related Work	59
5.3	Data Acquisition and Preparation	60
5.4	Generating DTI Integral Curves	60
5.5	Single Linkage Algorithm	66
5.6	Exploration, Evaluation and Automatic Labeling	69
5.6.1	Interactive Exploration of Clustering Models	69
5.6.2	Evaluation	70
5.6.3	Automatic Fiber Bundle Labeling	73
5.7	Discussion and Conclusion	74
6	Quantitative Study of the Streamtubes Model	77
6.1	Quantitative Tract-of-Interest Metrics for White Matter Integrity . .	78
6.1.1	Tract of Interest (TOI) Selection	80
6.1.2	Proposed Metrics	81
6.1.3	Reproducibility Study	84
6.1.4	Whole Brain Metrics	85
6.1.5	Interhemispheric Fibers	88
6.1.6	Conclusion	88
6.2	Inverse Modeling	89

6.2.1	Model	91
6.2.2	Forward Modeling	92
6.2.3	Optimization	94
7	Applications	99
7.1	Multiple Sclerosis	100
7.2	Tumor Analysis	103
7.2.1	Method	107
7.2.2	Results	109
7.3	The Aye-aye Brain	114
7.4	Conclusions	117
8	Summary and Conclusion	120

List of Tables

3.1	Parameters used to select streamtubes for Figures 3.2, 3.4, 3.9 and 3.10.	20
3.2	Parameters used to select streamsurfaces selection for Figures 3.2, 3.7, 3.9, and 3.10.	26
4.1	Advantages reported for the CAVE and fishtank VR.	36
5.1	Confidence ratings for the 12 white matter features in the clustering models of four data sets, S1-S4 . WMF – white matter fiber tract; ff – frontal forceps; slf – superior longitudinal fasciculus; ilf – inferior longitudinal fasciculus; af – arcuate fasciculus; uf – uncinata fasciculus; ac – anterior commissure; cb – cingulum bundle; scp – superior cerebellar peduncle; mcp – middle cerebellar peduncle; icp – inferior cerebellar peduncle; cst – cortico–spinal tract; v indicating very likely located, l indicating likely located, p indicating possibly located, and u indicating unlocated. A rating pair shows the ratings for the same tract in each cerebral hemisphere.	72
6.1	Results of the reproducibility study (mean \pm SD) and coefficient of variance (COV; means and SD are rounded)	85
6.2	Group comparison, whole brain corrected for age	86
7.1	Comparison of Diffusion Metric Data Obtained From Right and Left Hemispheres for Four Normal Healthy Volunteers	110

7.2	Comparison of Diffusion Metric Data Obtained From Right and Left Hemispheres in a 33-Year-Old Female Volunteer Scanned on Three Separate Occasions	111
7.3	Comparison of Diffusion Metric Data Obtained From Different Sized Regions-of-Interest Covering Tumor/Edema (Ipsilateral) and Contralateral Volumes in a 72-Year-Old Male Patient With a Right Parietal Metastasis of a Lung Adenocarcinoma	112

List of Figures

1.1	A depiction of the neuron taken from (Enchanted Learning, 2006).	2
3.1	(a) shows the hyperstreamlines from (Delmarcelle and Hesselink, 1993). (b) shows the streamlines tracked in the diffusion tensor field from (Xue et al., 1999). (c) shows the volume rendering technique from (Kindlmann et al., 1999) which assigns color and opacity with barycentric transfer function.	12
3.2	Visualization of a human brain using our method. Geometric models include the red streamtubes, the green streamsurfaces, the blue ventricle surface, and the wireframed skull surface. Large neural structures such as the corpus callosum and the corona radiata are represented by the red streamtubes in this view.	14
3.3	Barycentric map for the definition of anisotropy regions. Here, linear anisotropy regions are defined as $c_l > 0.3$ (Vertically hatched area); planar anisotropy regions are defined as $c_p > 0.3$ (horizontally hatched area). The small crosshatched triangle has both linear anisotropy and planar anisotropy by this definition.	16
3.4	Streamtubes representing regions of linear anisotropy in a human brain. The long U-shaped tubes pass through the corpus callosum.	17

3.5	In order to emphasize important differences between a pair of trajectories, we average the distance between them only over the region where they are at least T_t apart; smaller differences are assumed to be insignificant.	19
3.6	Generating a streamtube from a trajectory and cross-section. The second and third eigenvalues of the diffusion tensor, λ_2 and λ_3 , define the relative shape of the cross-section. Its orientation is defined by the corresponding eigenvectors.	21
3.7	Streamsurfaces extend along major and medium eigenvectors in regions of planar anisotropy. Part of the surfaces matches the area where two tracts of neural fibers cross each other.	22
3.8	Expanding a streamsurface. Begin at point labeled Seed, add the red edges; repeat the expansion at points 2, 3, 4, etc., as many times as possible.	25
3.9	Front view of the human brain data set visualization. The culling parameters are shown in Tables 3.1 and 3.2.	29
3.10	Side view of the human brain data set visualization. The culling parameters are the same as those used in Figure 3.9.	30
3.11	Human brain data set visualization with average linear anisotropy threshold set to 0.20 instead of 0.30. The number of streamtubes increases from 472 to 4,538.	31
4.1	A model of the CAVE viewed from behind the back wall. The user stands inside the three rear-projected walls and sees a virtual world on the head-tracked stereo displays.	33
4.2	The virtual reality applications run in the CAVE (left) and fishtank VR display (right).	34

4.3	Interactive renderings of a human brain dataset. The renderings (top) show collections of threads consistent with major white-matter structures: IC=internal capsule, CR=corona radiata, CB=cingulum bundle, CC=corpus callosum, diagrammed on the bottom. Components of the tensor-valued data control thread direction, color, and density as described in the text. Direct volume rendering simultaneously shows the ventricles (labeled V) in blue for anatomical context.	37
4.4	Left: a direct-volume-rendered layer showing regions with different diffusion anisotropy. Right: a thread layer showing both the anisotropy and diffusion direction. The two layers are combined in the center image which shows significantly more information and requires little additional visual bandwidth.	39
4.5	Three thread volumes showing brain connectivity information (the front of the head points right). Shown without halos on the left, with halos in the center, and after slightly shifting the halos away from the viewer to brighten the threads on the right.	39
4.6	Black arrows illustrate the path of virtual viewing rays through a cross section of a volumetric thread and halo. When the rays, coming from the viewer's perspective, pass through part of the dark halo volume before reaching the red thread, the resulting color is a darker red than when the halo volume is shifted away from the viewer (the bottom case).	43
4.7	Filtering a thread into a volume (2D view). For each voxel within a radius of two voxels from a thread, we use the shortest distance to the thread as the input to the filter (at right). The grid on the left and the horizontal axis on the right both show single-voxel spacing.	45

4.8	Filtering a halo into a volume (2D view); red depicts the thread and black the halo around it. The red curve is the filter for the thread and the black curve is the filter with which the halo is generated. The grid on the left and the horizontal axis on the right both show single-voxel spacing.	45
4.9	A sequence of renderings of a thread density volume with increasing length threshold from left to right. The rightmost image shows only long threads.	46
4.10	The interactive exploration tool. Clockwise from upper left are a 2D barycentric widget, a 1D widget, a 2D Cartesian widget, and a 2D Cartesian culling widget.	47
4.11	A data-flow diagram of the rendering setup for Figure 4.10.	50
4.12	(a,b) A plaster model showing areas of linear and planar water self-diffusion obtained from DT-MR images. (c) Detail of support structures (dark gray surfaces around tubes) for the streamtubes; these surfaces are created using the second and third eigenvectors of the tensors that produce the tubes.	52
5.1	(a) shows the shortcoming of defining the proximity between two curves as the distance between two selected points. The Hausdorff and Fréchet distances of A and B are close to those of B and C, even if B is considered more similar to A than C in certain contexts. (b) shows the motivation behind the mean thresholded closest distances. A and B are considered different if they branch for a portion of their lengths. Without the threshold, the mean closest distances between A and B is low if they stay close for a large part of their lengths.	64

5.2	Integral curve models generated from a $128 \times 128 \times 90$ data set with a voxel size of $1.7mm \times 1.7mm \times 1.7mm$. (a) Shows 438,042 integral curves generated from a jittered regular grid of $256 \times 256 \times 180$, with minimum length set to $13mm$; minimum linear anisotropy set to 0.15 and minimum average linear anisotropy set to 0.1. (b) Shows the model generated from the same data set and the same jittered regular grid sampling, and with the added culling and gray matter projection constraints. The culling threshold on d_{St} is set to $1.0mm$ with $t = 0.5mm$. The model in (b) has 6,113 integral curves. Color is mapped to the linear anisotropy value. Redder means higher linear anisotropy. . . .	65
5.3	The clustering result of the model in 5.2(b) using single linkage algorithm and d_{Lt} as proximity measure with a threshold of $3.5mm$. The top left corner of the interactive interface contains three sliders for adjusting the proximity threshold, minimum mean fractional anisotropy values and minimum number of curves in a cluster.	68
5.4	Snapshots of exploring the clustering models shown in Fig. 5.3. a,b,c,d and e show the clustering results with proximity threshold set to $1.5mm$, $2.5mm$, $3.5mm$, $4.5mm$ and $5.5mm$ respectively. f,g,h,i and j show the same model/threshold with minimum cluster size threshold set to 10.	69
5.5	The interactive interface can be used to quickly evaluate a new clustering algorithm or proximity measure. a,b,c,d show the same model used in Fig. 5.4 with the same single linkage algorithm, but different proximity measure (d_{St} instead of d_{Lt}). a, b, c, d and e show the clustering results with the proximity threshold set to $0.5mm$, $1.4mm$, $1.5mm$, $2.1mm$ and $2.5mm$ respectively. The abrupt change of cluster structures around proximity threshold $1.5mm$ suggests that the d_{St} is more prone to this kind of jump than d_{Lt}	70

5.6	Four of the fiber bundles in the template. (a) left cingulum bundle, (b) right cingulum bundle, (c) left uncinate, (d) forceps minor. Blue surface represents the ventricles and a slice of $T2$ image is shown as context. Colors are randomly generated for each bundle of each subject.	71
5.7	The sum of matching scores for all the fiber bundles in the template. (a) shows the result of S4. Proximity threshold is $3.6mm$ for the optimal match, while the expert-selected threshold is $3.9mm$. (b) shows the result of S5. Proximity threshold is $3.2mm$ for the optimal match, while the expert-selected threshold is $3.5mm$.	75
5.8	The matched bundles for subject 4 and 5. Color is fixed on each label. Blue surface represents the ventricles. Note the correspondence between the same bundles across the two subjects.	75
6.1	The 3D views of whole brain streamtube models for one healthy control (a) and one VCI patient (b).	87
6.2	The IHF selection in one healthy control (a) and one VCI patient (b).	87
6.3	This figure shows the results of the forward imaging model on the central portion of a mid-coronal section through a human brain. The section includes parts of the cingulum bundles (green in (a)) and the corpus callosum (red in (a)). (a) and (b) show DTIs generated from the original DWIs and simulated DWIs, with the diagonal tensor elements mapped to red, green, and blue; (d) and (e) show a section in one of the original DWIs and the simulated DWIs; (c) shows the fiber model used in this forward modeling.	94
6.4	This plot shows the fibers optimization processes for the greedy algorithm and the simulated annealing. For both algorithms the χ^2 value converges to around 77,600. The greedy algorithm converges faster, and it does not appear to get trapped in a local minimum.	96

6.5	<p>a), b) ,c) and d) show the snapshots of the value $\chi_v^2 = \sum_{q=1}^m \left(\frac{I_{v,q} - \hat{I}_{v,q}}{\sigma} \right)^2$ for each voxel v in the data volume during different stages of the simulated annealing. S is the number of accepted moves. e), f), g) and h) show the fiber model generated at these stages. N is the number of fibers. Note that g) and h) both have 600 fibers, but the optimization of the configuration lowered the χ^2 value from g) to h) significantly.</p>	97
7.1	Fast-spin-echo images through the region of corpus callosum.	100
7.2	<p>Left: Streamtube map culled to fibers intersecting lesions. Right: Streamtube map with additional culling to fibers passing through corpus callosum and intersecting lesion.</p>	101
7.3	<p>Left: T2W image showing few focal T2-hyperintensities. Right: Sagittal fiber map showing in red the location of fibers in the midline corpus callosum that connect to the focal MS lesions.</p>	101
7.4	<p>The process of creating 2D barycentric histograms from DTI data. a: For a given slice, the brain is divided in half with a straight line following the falx cerebri on the T_2-weighted EP image. b: A barycentric histogram is generated for each side of the brain. In this example, the histogram comes from the tumor-bearing section shown in a. Note the large number of voxels near the $c_s = 1$ vertex, reflecting the cystic component of the tumor. c: The difference histogram obtained by subtracting normalized barycentric histograms calculated from tumor-bearing and contralateral sections. In this difference histogram, zero maps to medium gray because the difference is signed. Note that the most striking difference occurs near the $c_s = 1$ vertex. d: The definition of the diffusion anisotropy regions in the barycentric histogram. These regions are chosen to capture the differences between regions A and B in part c.</p>	104

7.5 Geometric models representing the diffusion metric data obtained from a 33-year-old female volunteer viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. Linear diffusion is represented by red streamtubes, planar diffusion by green streamsurface fragments, and CSF by the blue surface. Note the symmetry between two sides of the brain in this normal subject. 105

7.6 Maps of (a) T_2 -weighted signal intensity, (b) $\langle D \rangle$, (c) FA, and (d) color composite diffusion metric data obtained from a 72-year-old male patient with a right parietal metastasis of a lung adenocarcinoma. In d, maps of linear (red channel), planar (green), and spherical (blue) diffusion are overlaid so that the color in each voxel represents the relative magnitude of the three diffusion metrics. Note the high $\langle D \rangle$, low FA, and spherical (isotropic) diffusion of the tumor/edema volume. Also note that the tumor/edema volume is surrounded by a large region of planar diffusion. 106

7.7 Geometric models representing the diffusion metric data obtained from the 72-year-old male patient viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. Linear diffusion is represented by red streamtubes, planar diffusion by green streamsurface fragments, and CSF by the blue surface. The yellow surface, which is rendered from the signal abnormality on the T_2 -weighted EP images, depicts the shape of the tumor/edema volume. Note that the tumor/edema volume is partly surrounded by regions of planar diffusion. 107

- 7.8 3D DTI in the intact fixed brain of the mouse lemur. Shown on the left are hemi-segments of para-sagittal, coronal, and horizontal slices taken from the 3D volume. A composite RGB image was formed from the diagonal tensor elements, where red=rostral-caudal, green=dorsal-ventral, and blue=lateral. A dominant color in a region indicates the directionality of a track. Diffusion weighted volumes (28) were acquired using a 3D PGSE sequence. The diffusion gradients were applied along 7 different directions, and four values of b-matrix were used along each direction (maximum $b \sim 2000 \text{ mm}^2/\text{s}$). The data were acquired with $512 \times 256 \times 256$ points and a $\sim 70 \text{ um}$ isotropic resolution. On the right a streamtube representation of the dataset. For anatomical context, a T2-weighted grayscale para-sagittal slice from the 3D volume is shown in the yellow plane. 115
- 7.9 Comparison of different imaging modalities in the brain of the Aye-Aye lemur. 3D DTI data were acquired in a fixed, intact, brain at $\sim 100 \text{ um}$ resolution. Shown is the caudal portion of a single para-sagittal slice from the 3D dataset. The panels from left-to-right represent the T2-weighted anatomical image, an RGB composite image of the diagonal elements of the diffusion tensor (see above), and the corresponding fiber-stained histology section. Note that the anatomical image and histology essentially provide scalar information about the location of the fiber tracts, whereas the DTI provides fiber location as well as directionality of the fibers. Scale bar=1cm. 116

7.10	<p>Left: DTI slice from the Aye-Aye data set showing targeted fiber tracts for histological comparisons. DTI slice is a RGB composite image of the diagonal elements of the diffusion tensor (see above). Center: ‘Best-match’ histology slice to the DTI image shown on the left. Micrographs at several magnifications are shown; the highest magnification (right panel) shows the boundary between the cingulum bundle and the splenium of the corpus callosum. Histology was performed on the same brain as the MRI by embedding in celloidin, serially sectioning at 40 um using a sliding microtome, and staining with the Gallyas technique. Right: Diffusion ellipsoid rendering showing the relative diffusion anisotropy on a voxel-by-voxel basis. The observed orientation of the diffusion ellipsoids is agrees qualitative with the histology.</p>	118
7.11	<p>Fiber tracings from a sequence of sites in the cingulum bundle in the Aye-Aye brain. (A) depicts polar histograms centered on each sample site depicting fiber orientation and spread; the histograms encodes the mean polar angle (θ) of the fiber orientation and the standard deviation (SD). Data from (A) were extracted directly from microscopic images of the histology slice shown in (B). (C) shows the corresponding DTI slice of the cingulum bundle rendered in an RGB format (red=anterior-posterior, green=superior-inferior, blue=medial-lateral).</p>	119

Chapter 1

Introduction

1.1 Overview

The complexity of the brain makes the study of neuroanatomy a demanding task. The building block of this complex structure is the neuron, the functional cellular unit of the nervous system. The neuron consists of dendrites, a cell body, an axon, and axon terminals (see Fig. 1.1). Axons form wire-like structures in the brain's white matter which transfers information between the dendrites, cell bodies, and axon terminals. There are about 100 billion neurons in the human brain, and for some types of neurons the axon can be as long as a meter. Neuroscientists strive to understand these neurons and their interconnections.

Traditional methods for examining the brain, including post-mortem visual inspection and chemical tracer methods are invasive, time-consuming, and limited in application. To address these problems, a new method for studying neuroanatomy was pioneered in the 1980's. Known as diffusion magnetic resonance imaging (MRI), this technique measures the average diffusion of water molecules during a specified time interval over a small region called a voxel. In a glass of water, water molecules change course only by colliding with other water molecules engaging in the same constant, seemingly random motion known as Brownian motion. In the brain's white

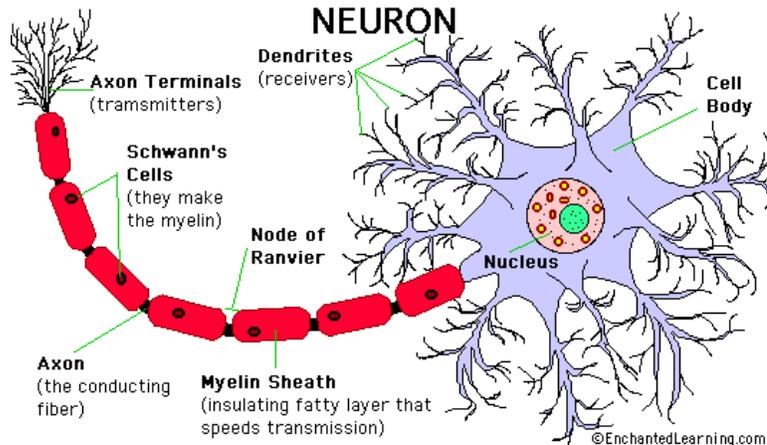


Figure 1.1: A depiction of the neuron taken from (Enchanted Learning, 2006).

matter, water-molecule diffusion is hindered by various components of the tissue such as the cell membrane, the axon sheath and the collagen. Different compositions of materials lead to different diffusion patterns inside (intracellular) and outside (extracellular) the neurons. It is generally believed that the diffusion measured by diffusion MRI contains a mixture of intercellular and extracellular diffusion and the exchange between the two. Chapter 2 gives some background information in diffusion imaging. Although the exact mechanism of the water diffusion in the white matter is still unknown, experiments show that in coherently oriented fiber structures, the preferred direction of diffusion measured by diffusion MRI coincides with the fiber orientation. Armed with this powerful new technique, we can hope to model the diffusion process in white matter and infer neural fiber structures using diffusion MRI data.

The thesis of this dissertation is that the diffusion imaging data can reveal white-matter structures with visualization and data analysis techniques. This dissertation presents our work in developing the diffusion-imaging visualization and analysis techniques. I present comparison between different visualization schemes, validation of the diffusion imaging methods and applications of the visualization and analysis techniques in defense of the thesis. It is important to point out that diffusion-imaging

visualization is an inherently interdisciplinary area, and many parts of this dissertation are results of collaborations.

We first introduce the components of our models, known as streamtubes and streamsurfaces, that encode diffusion information through their shapes, colors, sizes and spatial placements. We then design several display and interaction schemes that enable users to view these complex geometric models in a clear and informative way. To further assist users in identifying anatomically related fiber bundle structures, we develop fiber-clustering techniques. Apart from qualitative visual exploration, we also perform quantitative studies on the streamtubes model: we design a number of metrics on the streamtubes model and test them on both normal and diseased brains. We also implement a forward model that optimizes our streamtubes model by simulating the diffusion MRI signals from that model. Finally we enumerate several applications of our techniques to biology and medicine.

In Chapter 3, we propose a set of 3D models, built from the components we call streamtubes and streamsurfaces, that convey the connectivity, coherency and other information about the white matter structures from diffusion imaging. Streamtubes extend the tractography models, which consist of curves from the diffusion imaging data that correlate with neural fiber pathways. The trajectories of the streamtubes follow the fastest direction of diffusion, while the color along the streamtubes represent the coherency of the neural fibers. Streamsurfaces follow planar diffusion, which arises from crossing fibers or material boundaries. The seeding scheme is designed to cover the whole data volume while avoiding artificial clusters.

In Chapter 4, we design and experiment with different display environments for the purpose of brain visualization. We employ two virtual-reality display systems, the CAVE and the fishtank VR system, and compare the benefits of the two. We also develop a thread-and-halo technique suitable for interactive volume rendering volume rendering to render the fiber structures together with a number of other components for exploring the brain models. Finally, we explore color rapid prototyping plaster

models as visualization tools for the brain models.

In Chapter 5, we present a method for automatically grouping white matter into coherent anatomical structures, using a proximity-based algorithm for sampling, culling, and clustering diffusion tensor imaging integral curves. We then demonstrate the accuracy of the method in forming bundles of curves that are anatomically consistent when guided and evaluated by an expert rater. We then develop an automated labeling algorithm that models the expert rater’s decisions and we evaluate the results.

In Chapter 6, we present two quantitative methods to evaluate the streamtubes model. We first try to assess white matter integrity by calculating a number of metrics on the DTI integral curves. Then we try to increase the accuracy of the streamtubes model by building a forward model from the streamtubes to diffusion MRI signals and applying this forward model to optimization of the streamtubes model.

A model and visualization scheme for diffusion MRI can only be tested through a variety of applications in neuroanatomy, medicine and biology. Indeed, without applications to guide the development of computational and visualization tools, these tools are far less likely to be useful. In Chapter 7, we present several applications that apply the qualitative and quantitative modeling and visualization techniques discussed through the previous chapters. In turn, these practical applications validate and provide heuristics to our modeling approach.

1.2 Contributions

The contributions of this dissertation include

- Streamtubes and streamsurfaces (Zhang et al., 2003) are the first models to incorporate more than just major eigenvector field information in tractography methods.

- A clustering algorithm for streamtubes (Zhang and Laidlaw, 2004) and a matching algorithm for comparing fiber bundles across subjects (Zhang and Laidlaw, 2005).
- A forward model for diffusion MRI and a optimization scheme for tractography models (Zhang and Laidlaw, 2006).
- Application of a barycentric mapping for anisotropy histogram in detecting white matter changes in the presence of the brain tumor (Zhang et al., 2004).
- Comparison between CAVE and Fishtank VR system for brain visualization (Demiralp et al., 2006).
- Application of diffusion MRI methods in studying white matter diseases (Simon et al., 2005; Zhang et al., 2001a) and lemur brain (Ahrens et al., 2004).

Chapter 2

An Introduction to Diffusion Imaging

This chapter introduces some background knowledge about diffusion, diffusion imaging and diffusion tensor imaging which the following chapters use extensively.

Water molecules in human tissue constantly collide randomly with one another and with other molecules, a phenomenon called Brownian motion; one consequence of this is that water that starts in one area spreads to nearby areas over time, a process called diffusion. Magnetic resonance imaging can be programmed to generate signals relating to this diffusion process. These signals can then be used to fit a model to describe the water diffusion in the brain. Since this water diffusion is restricted by the tissue structures, the diffusion model can then be employed in visualizing and analyzing brain anatomy. This chapter briefly reviews some background knowledge in diffusion imaging. For a more detailed introduction, see (Buxton, 2002; Tuch, 2002).

In pure water, the seemingly random Brownian motion results in a dynamically expanding Gaussian distribution of water molecules released from one point (Einstein, 1905). In human tissues, however, cell membranes and large protein molecules limit the motion of water molecules. The geometrical and physical properties of the tissue

determine the rate and direction of diffusion. We can thus infer the microstructure of human tissue by measuring the diffusion of the water molecules, just as we feel out the shape of a mug by measuring the deformation of our finger tips.

The discovery of the nuclear magnetic resonance (NMR) effect (Purcell et al., 1946; Bloc, 1946) in 1946 was the beginning of work that has led to the current form of diffusion magnetic resonance imaging (MRI). Two important landmarks were the discovery of the spin echo (Hahn, 1950), whose signal is perturbed by the water molecule diffusion, and MR imaging (Lauterbur, 1973), which determines exactly where the NMR signal originates within the sample. Diffusion imaging was the first imaging modality that could measure the diffusion of water in human tissues in vivo. Although the exact mechanism of the generation of diffusion MRI signals in biological tissues is not fully understood, it is generally believed that the quantity measured by diffusion MRI is a mixture of intracellular diffusion, intercellular diffusion, and the exchange between the two sides of the cell membrane (Stanisz and Henkelman, 2002).

Inferring tissue structure from the diffusion process requires exploring the orientation dependence of the diffusion. This dependence can be described by the diffusion propagator $P(r, r', \tau)$, which is the probability of a water molecule traveling from position r' to r in diffusion time τ (Callaghan, 1993). In practice, the number of diffusion directions we can measure in a clinical scan is limited by scanning time, making it impossible to sample complicated diffusion propagator with enough resolution.

A diffusion tensor (Basser et al., 1994) is a simplified diffusion propagator. It fully describes the diffusion propagator for free diffusion in a uniform anisotropic medium (Gaussian diffusion). For example, a diffusion tensor is a good model for diffusion in uniformly oriented white matter structures such as the corpus callosum, but is insufficient in areas where different tracts cross or merge. The coefficients of the diffusion tensor, \mathbf{D} , are related to the diffusion-weighted MRI(DWI) signals by (Bihan, 1991): $\tilde{\mathbf{I}} = \mathbf{I}_0 \exp(\mathbf{b} : \mathbf{D})$, where \mathbf{I}_0 is the 0-weighted diffusion image, the tensor \mathbf{b} characterizes the diffusion-encoding gradient pulses used in the MRI

sequence, and $\mathbf{b} : \mathbf{D} = \sum_{i=1}^3 \sum_{j=i}^3 \mathbf{b}_{ij} \mathbf{D}_{ij}$ is the tensor dot product.

A 3D diffusion tensor is a 3×3 positive symmetric matrix:

$$\mathbf{D} = \begin{bmatrix} \mathbf{D}_{xx} & \mathbf{D}_{xy} & \mathbf{D}_{xz} \\ \mathbf{D}_{xy} & \mathbf{D}_{yy} & \mathbf{D}_{yz} \\ \mathbf{D}_{xz} & \mathbf{D}_{yz} & \mathbf{D}_{zz} \end{bmatrix}$$

Diagonalizing \mathbf{D} , we get three positive eigenvalues λ_1, λ_2 and λ_3 (in decreasing order) and their corresponding eigenvectors \vec{e}_1, \vec{e}_2 and \vec{e}_3 . Many scalar indices and visualization methods are based on the eigenvalues and eigenvectors of diffusion tensor imaging (DTI) measurements.

One geometric representation of Gaussian diffusion is a diffusion ellipsoid. These ellipsoids represent the surface of constant mean-squared displacement of diffusing water molecules at some time τ after they are released from one point. The shape of a diffusion ellipsoid is inherently related to the eigenvalues and eigenvectors of the diffusion tensor: the three principal radii are proportional to the eigenvalues and the axes of the ellipsoid aligned with the three orthogonal eigenvectors of the diffusion tensor. The disparity of diffusion along different directions, or diffusion anisotropy, is often defined on the three eigenvalues of the diffusion tensor (Pierpaoli and Basser, 1996; Westin et al., 1997a).

DTI measurements have been validated within acceptable error on the fibrous muscle tissue of the heart (Scollan et al., 1998; Hsu et al., 1998). However, in a voxel containing nonuniformly oriented neural fibers DTI measures an average signal from all the fibers within the voxel, which usually results in an apparent reduction of anisotropy and increase in uncertainty (Jones, 2003). To resolve the uncertainty in these areas, q -space spectral imaging (Assaf and Cohen, 2000) and other high angular resolution diffusion imaging methods (Tuch et al., 2002; Frank, 2002) have been explored.

Image acquisition for diffusion MRI is a very active research area. Progress is frequently reported on resolution improvement and reductions in imaging time, noise,

and distortion.

Chapter 3

Streamtubes and Streamsurfaces

Model

In this chapter, we propose a set of 3D models that convey the connectivity, coherency and other information about the white matter structures from the diffusion imaging. We distinguish between coherent fibrous structures and incoherent mesh-like planar structures and represent values in the two areas using streamtubes and streamsurfaces, respectively. Streamtubes represent structures with primarily linear diffusion, typically fiber tracts; streamtube direction correlates with tract orientation. The cross-sectional shape and color of each streamtube represent additional information from the diffusion tensor at each point. Streamsurfaces represent structures in which diffusion is primarily planar.

Our algorithm chooses a quantitatively representative subset of the streamtubes and streamsurfaces for display. We describe the set of metrics used for the culling process, which reduces visual clutter and improves interactivity. We also generate anatomical landmarks to identify the locations of such structures as the eyes, skull surface, and ventricles. The final models are complex surface geometries that can be imported into many interactive graphics software environments.

The work in this chapter is published in (Zhang et al., 2003).

3.1 Related Work

Researchers have designed visualization methods that represent an entire data set in a 2D slice of a 3D diffusion tensor field. Several attempts to visualize 3D DT-MRI data have also been made recently. Pierpaoli *et al.* use arrays of ellipsoids to represent a two-dimensional diffusion tensor field (Pierpaoli and Basser, 1996). Since the diffusion tensor matrix is symmetric and has positive eigenvalues, an ellipsoid is its natural geometric representation. Each axis of the ellipsoid represents one eigenvector and its corresponding eigenvalue. Laidlaw *et al.* normalize the size of the ellipsoid to give a more continuous appearance and, in a second method, borrow concepts from oil painting to display diffusion tensor images (Laidlaw et al., 1998). Two-dimensional brush strokes built up in several layers represent different features of the diffusion tensors.

The two-dimensional methods in (Pierpaoli and Basser, 1996) and (Laidlaw et al., 1998) visualize a diffusion tensor field by completely visualizing the tensors of discrete sample points. When applied to a 3D data set, such methods have two limitations:

- Visualizing every sample point in the 3D data set produces so many ellipsoids or brushstrokes that internal structures are difficult to see.
- The continuity inherent in biological tissues is not properly represented in the final image. For example, neural fibers in the brain are difficult to locate within an array of ellipsoids.

Several methods for 3D diffusion tensor field visualization have been developed to address these problems, each making different choices of the subset of the tensor information to represent and how to represent it.

Dickinson, in his design of visualization software for tensor-field datasets, tracks tensor field lines that are everywhere parallel to an eigenvector of a 3D tensor field (Dickinson, 1989). Delmarcelle *et al.* built on this in proposing hyperstreamlines, visual

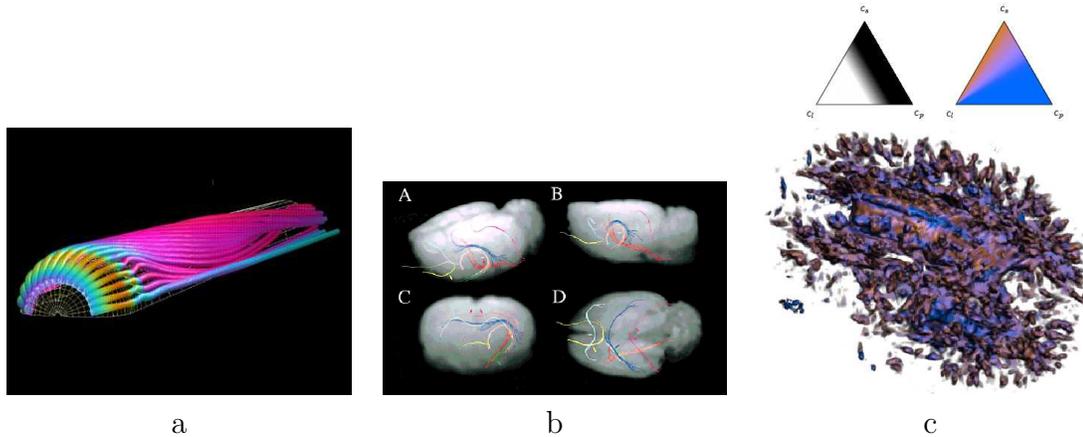


Figure 3.1: (a) shows the hyperstreamlines from (Delmarcelle and Hesselink, 1993). (b) shows the streamlines tracked in the diffusion tensor field from (Xue et al., 1999). (c) shows the volume rendering technique from (Kindlmann et al., 1999) which assigns color and opacity with barycentric transfer function.

icons for representing tensor information along its trajectory (Delmarcelle and Hesselink, 1993) (see Fig. 3.1(a)). The idea behind this method, which is analogous to using ellipsoids to represent one diffusion tensor, is to visualize *all* the information but only at *some* locations. The trajectory of the hyperstreamline is generated by the vector field defined by the major eigenvector of the tensor. The cross-sectional shape and color along the trajectory encode information about the other two eigenvectors and the magnitude of the major eigenvector. The hyperstreamline method has been applied to both stress tensor fields and momentum flux density tensor fields, but not to diffusion tensor fields.

Xue *et al.* track streamlines in the major eigenvector field of the diffusion tensor field to visualize DT-MRI data sets (Xue et al., 1999) (see Fig. 3.1(b)). For DT-MRI data sets, the major eigenvector of the tensor matrix is in the direction of fastest diffusion. Xue *et al.* use a fiber-tracking method dubbed FACT (Mori et al., 1999) to track the linear features in biological tissues following the major eigenvector of each diffusion tensor. There are other methods for tracking linear features in diffusion tensor

data sets: Weinstein *et al.* stabilize the tracking in regions with nonlinear preferential diffusion using their tensorlines method (Weinstein et al., 1999), and Basser *et al.* use numerical methods to solve a Frenet equation describing the evolution of a fiber tract (Basser et al., 2000). These methods are, however, constrained to visualizing only one vector field, and neither paper discusses the sampling and placement of the tracts in the image.

Kindlmann *et al.* take a volume-rendering approach to the problem (Kindlmann et al., 1999): their method displays only some of the information but displays that information densely within a volume. A hue-ball and a barycentric map assign color and opacity to each point based on the diffusion measurements (see Fig. 3.1(c)). The result shows the data set at high resolution. However, the composition into a single pixel of many data points along a ray makes it difficult to discern any given point, and the paths of fibrous structures are hard to see.

Our goal is to visualize connectivity and tissue microstructure in MR diffusion tensor images of biological tissue. The streamtubes we use here build on hyperstreamlines but attempt to overcome their limitations in this context:

- The cross-section of a hyperstreamline can grow quite large, limiting the density of hyperstreamlines in a scene and thus the level of detail we can visualize.
- The images in hyperstreamline papers generally contain only a few hyperstreamlines. To visualize microstructures in biological tissues, we need more hyperstreamlines in the scene.
- While a hyperstreamline is an intuitive visual representation for linear structures in biological tissue, it is not very effective for representing planar diffusion.

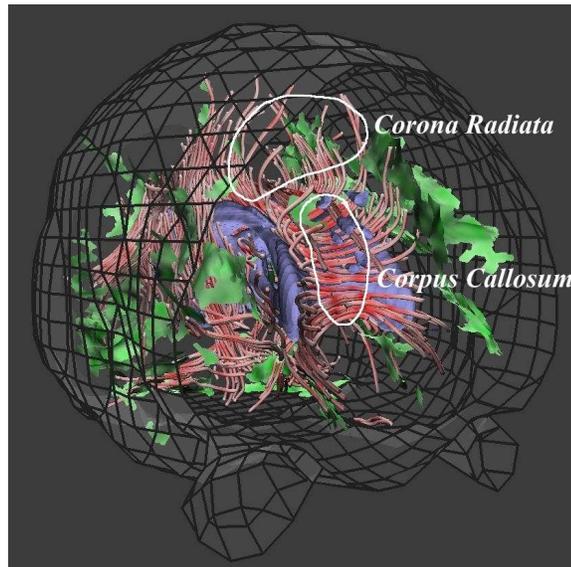


Figure 3.2: Visualization of a human brain using our method. Geometric models include the red streamtubes, the green streamsurfaces, the blue ventricle surface, and the wireframed skull surface. Large neural structures such as the corpus callosum and the corona radiata are represented by the red streamtubes in this view.

3.2 Streamtubes and Streamsurfaces

We distinguish between regions of linear and planar anisotropy and employ streamtubes and streamsurfaces respectively to visualize these two types of regions (Zhang and Laidlaw, 2001). Section 3.2.1 describes classifying different kinds of anisotropies. In section 3.2.2 we discuss issues related to streamtubes, including their definition, the extension of the trajectories they follow, how to sample the seed points for an initial trajectory set so that all the data are covered, and how to select a representative set from the whole trajectory set. Section 3.2.3 explains the strategies for generating streamsurfaces in regions of planar anisotropy. We also generate geometric representations of anatomical landmarks to provide context, as discussed in section 3.2.4.

3.2.1 Anisotropy Classification

We classify diffusion anisotropy into linear diffusion anisotropy, planar diffusion anisotropy and diffusion isotropy with Westin’s metrics. Westin *et al.* (Westin et al., 1997a) define three characteristic qualities of a diffusion tensor value based on the three eigenvalues, $\lambda_1 \geq \lambda_2 \geq \lambda_3$:

$$c_l = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2 + \lambda_3},$$

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

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

where c_l represents linear anisotropy, c_p represents planar anisotropy, and c_s represents isotropy. Note that these values are complementary, $c_l + c_p + c_s = 1$.

A diffusion tensor with one eigenvalue much larger than the other two has large c_l , corresponding to linear anisotropy; white-matter tracts tend to produce tensors with linear anisotropy (Cormans et al., 1994). A diffusion tensor with two large and one small eigenvalues has large c_p , corresponding to planar anisotropy; sheet-like structures and crossings of fiber tracts in biological tissues are likely to yield planar anisotropy. Diffusion tensors whose three eigenvalues are roughly the same imply an underlying structure with no preferred diffusion direction; these tensors are isotropic and have a large c_s . Gray matter in the brain tends to produce isotropic diffusion tensors (Basser et al., 1994).

There are other anisotropy metrics, such as relative anisotropy (RA) or fractional anisotropy (FA) (Pierpaoli and Basser, 1996), that measure the variance of the eigenvalues and provide a directionless scale value. Such metrics measure differences among anisotropies but do not distinguish between linear and planar anisotropies and isotropies. Westin’s metrics, on the other hand, measure these kinds of anisotropies separately.

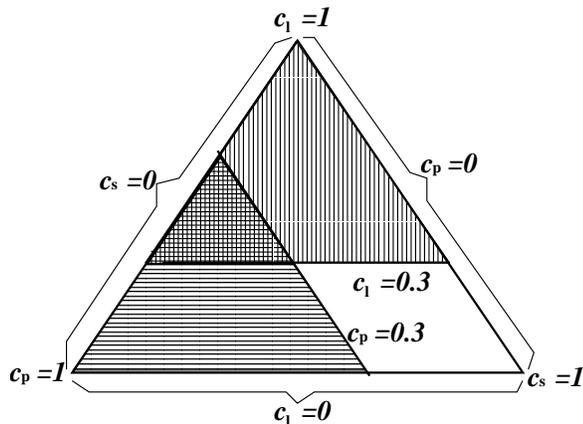


Figure 3.3: Barycentric map for the definition of anisotropy regions. Here, linear anisotropy regions are defined as $c_l > 0.3$ (Vertically hatched area); planar anisotropy regions are defined as $c_p > 0.3$ (horizontally hatched area). The small crosshatched triangle has both linear anisotropy and planar anisotropy by this definition.

We define linear anisotropy regions and planar anisotropy regions by setting thresholds on c_l and c_p . We can use a barycentric map (Kindlmann and Weinstein, 1999) to look at the defined anisotropy regions (see Figure 3.3). Note that this definition allows diffusion tensors to have both linear and planar anisotropy, so that streamtubes and streamsurfaces can be generated in the same region. From our observations, those regions are usually where linear structures transition to planar ones, often where neural fibers fan out or approach intersecting tracts.

3.2.2 Streamtubes for Linear Anisotropy

We chose streamtubes as the geometric primitive to represent linear anisotropy because it can naturally represent the underlying linear structures, can carry additional information provided by the diffusion tensors, and has the potential to reduce visual clutter.

The visual mapping of the streamtube is similar to that of the hyperstreamline. The trajectory sweeps along the major eigenvector field, and the cross-sectional shape

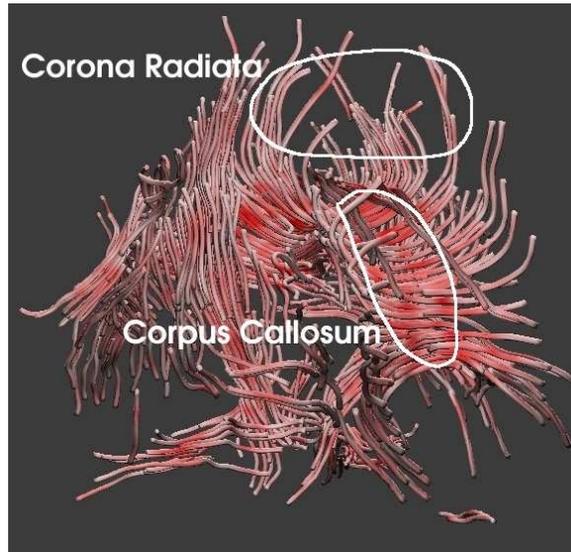


Figure 3.4: Streamtubes representing regions of linear anisotropy in a human brain. The long U-shaped tubes pass through the corpus callosum.

is an ellipse representing the ratio of the other two eigenvectors. The radius corresponding to the medium eigenvector is set to a constant to keep the streamtube slim, while the ratio between the medium and the minor eigenvectors is preserved. The color of the streamtube shows the linear anisotropy value: greater anisotropy is redder.

The shape of a streamtube is determined primarily by its trajectory. In other words, the major eigenvector field is the primary information visualized by streamtubes. We first fill the regions of linear anisotropy with a dense set of streamtubes and then select a representative subset from them. The four steps that generate a representative set of streamtubes are: (1) picking the seed points for a dense set of trajectories covering the regions of linear anisotropy, (2) calculating trajectories in the major eigenvector field, (3) selecting a sparser representative set of trajectories from the dense set, and (4) visualizing additional information such as the direction of the medium and minor eigenvectors. Figure 3.4 shows the streamtubes generated from a human brain data set.

Seed Points

Streamtube seed points are chosen so that streamtubes pass through all regions of high linear anisotropy. Recall that we reduce the number of trajectories later, so it is acceptable to generate many trajectories initially.

The seed points are not restricted to the sample points of the volume image. At any point within the data volume, we use tricubic B-spline functions to interpolate the tensor field (Basser et al., 2000), so there is no limit on how many seed points we use or where we put them. We generate a seed point from every sample point and jitter it within a voxel (Dippé and Wold, 1985).

Trajectories

Each streamline begins from a seed point and follows the major eigenvector field both forward and backward. We find the integral curve passing through the seed point using a second-order Runge-Kutta integration method (Press et al., 1992; Basser et al., 2000). (We also experimented with fourth-order Runge-Kutta, with the results being very similar to those generated by the second-order method.) Because streamtubes represent regions of high linear anisotropy, their trajectories are restricted to these regions. Also, the streamtubes are clipped to the data volume and to regions with sufficiently high signal-to-noise ratios.

Culling

Our sampling method starts with 1.2 million seed points and produces more than 150,000 trajectories on a $256 \times 256 \times 144$ human brain data set. Visualizing all the streamtubes is not only expensive but also undesirable, as including too many streamtubes in the scene would block the inner structures. A selection algorithm culls most of the trajectories and keeps only a representative set. We use three metrics for the culling process: the length of a trajectory, the average linear anisotropy along a trajectory, and the similarity between a trajectory and the group of trajectories

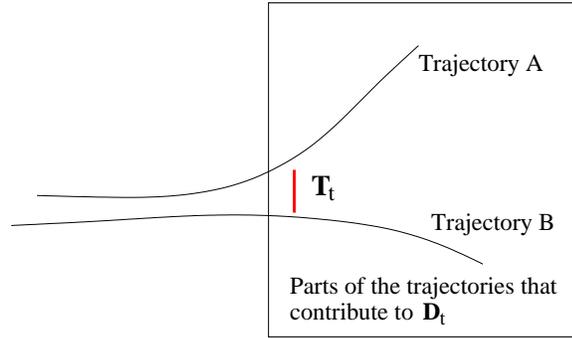


Figure 3.5: In order to emphasize important differences between a pair of trajectories, we average the distance between them only over the region where they are at least T_t apart; smaller differences are assumed to be insignificant.

already selected. Trajectories are kept or discarded according to their metrics. For example, a trajectory that is too similar to a selected trajectory is discarded.

We define the similarity between two trajectories using a distance measure:

$$D_t = \frac{\int_{s_0}^{s_1} \max(\text{dist}(s) - T_t, 0) ds}{\int_{s_0}^{s_1} \max\left(\frac{\text{dist}(s) - T_t}{|\text{dist}(s) - T_t|}, 0\right) ds} \quad (3.1)$$

where s parameterizes the arc length of the shorter trajectory, s_0 and s_1 are the starting and end points of s , and $\text{dist}(s)$ is the shortest distance from location s on the shorter trajectory to the longer trajectory. T_t ensures that we label two trajectories as different if they differ significantly over any portion of the arc length. For example, if $T_t = 0$, the pair of trajectories shown in Figure 3.5 would have a small D_t value, because they run close together over much of their length, making the denominator large. Setting T_t as in Figure 3.5 makes D_t larger because it reduces the denominator.

Note that other features such as the curvature of the trajectories can also be used to measure similarity. Since DTI data are averaged over the size of the voxel, the trajectories tracked are usually smooth. Thus the distance between the trajectories makes a decent measurement of the similarity.

For each of the three criteria, we set a threshold to limit the streamtubes we draw. Table 3.1 shows the thresholds for this culling process for Figures 3.2, 3.4, 3.9 and 3.10.

Table 3.1: Parameters used to select streamtubes for Figures 3.2, 3.4, 3.9 and 3.10.

Streamtube length	> 18.0 mm
Average linear anisotropy	> 0.30
T_t used in D_t for streamtubes	0.89 mm
Distance between lines, D_t	> 4.5 mm

Our effort in integral curve seeding and culling can be interpreted as Poisson Disk sampling in the high dimensional space of the integral curves. The Poisson Disk distribution can be defined as the limit of a uniform sampling process with a minimum-distance rejection criterion (Cook, 1986). Successive points are independently drawn from the uniform distribution. If a point is at least distance R from all points in the set of accepted points, it is added to that set. Otherwise, it is rejected. The choice of R controls the minimum allowable distance between points and, indirectly, the density of the Poisson Disk pattern.

Shape and Color

The trajectory visually represents the major eigenvector field of the diffusion tensor field. Once we have generated a representative set of streamtube trajectories, we map additional information to them and construct the streamtubes. The cross-section of a streamtube at a given point is an ellipse representing the other two eigenvectors and eigenvalues, as shown in Figure 3.6. We set the radius along the medium eigenvector to a constant, but preserve the aspect ratio of the second and third eigenvalues. Adjacent cross-sections are connected to form the streamtube. The color of the streamtube is based on the linear anisotropy value, with saturated red for the maximum and white for the minimum.

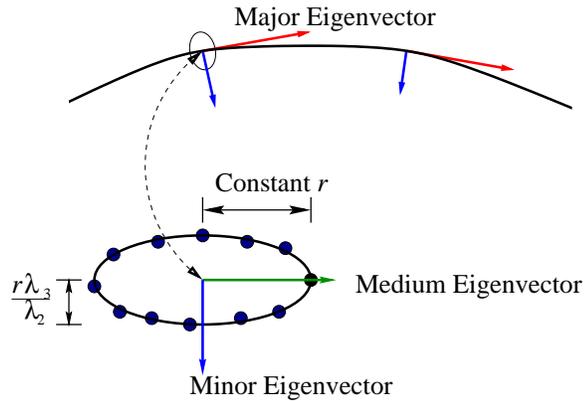


Figure 3.6: Generating a streamtube from a trajectory and cross-section. The second and third eigenvalues of the diffusion tensor, λ_2 and λ_3 , define the relative shape of the cross-section. Its orientation is defined by the corresponding eigenvectors.

3.2.3 Streamsurfaces for Planar Anisotropy

Diffusion tensors with planar anisotropy could result from a surface structure, a boundary between different materials, or a crossing of multiple linear features. Visualizing regions of planar anisotropy may help to illustrate these kinds of structures.

The streamsurface is the approximation of the surface that extends along both the major eigenvector and medium eigenvector. At any point on a streamsurface, the major and medium eigenvectors lie in the tangent plane to the surface. Given a starting point in the volume, we expand a streamsurface by following these two vector fields. We first generate many streamsurfaces and then select a subset of them for display. Colors are mapped to the surfaces to represent the planar anisotropy, c_s , at each point.

Seed Points

We choose the seed points to generate streamsurfaces similarly to those for streamtubes: we generate seed points by jittering every sample point in the data set.

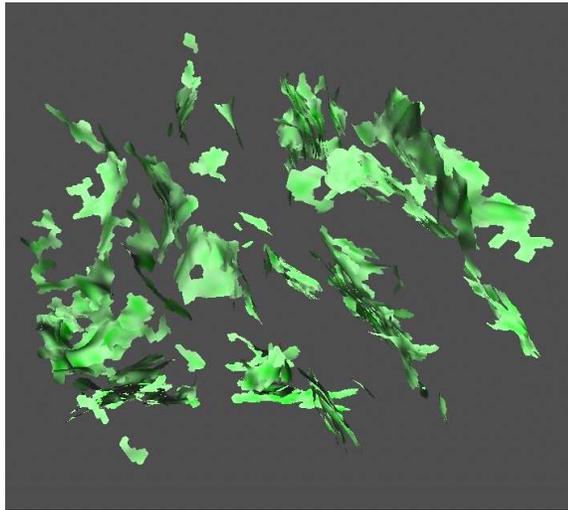


Figure 3.7: Streamsurfaces extend along major and medium eigenvectors in regions of planar anisotropy. Part of the surfaces matches the area where two tracts of neural fibers cross each other.

Streamsurface Integrability

Unlike streamtubes in a C^∞ vector field, streamsurfaces in two C^∞ vector fields do not always exist. In this section we give the necessary and sufficient condition for the existence of the streamsurface.

We use the notations and theorems from (Bishop and Crittenden, 1964). $[X, Y]$ is the Lie bracket of vector fields X and Y .

A p -dimensional distribution on a manifold M ($p \leq \dim(M)$) is a function θ defined on M which assigns to each $m \in M$ a p -dimensional linear subspace $\theta(m)$ of M_m . A p -dimensional distribution θ on M is of class C^∞ if there are C^∞ vector fields X_1, \dots, X_p defined in a neighborhood U of m and such that for every $n \in U$, $X_1(n), \dots, X_p(n)$ span $\theta(n)$. An integral manifold N of θ is a submanifold of M such that $di(N_n) = \theta(i(n))$ for every $n \in N$. We say that a vector field X belongs to the distribution θ and write $X \in \theta$, if for every m in the domain of X , $X(m) \in \theta(m)$. A distribution θ is involutive if for all C^∞ vector fields X, Y which belong to

θ , we have $[X, Y] \in \theta$. A distribution θ is integrable if for every $m \in M$ there is an integral manifold of θ containing m .

— Geometry of Manifolds (Bishop and Crittenden, 1964)

If we define e_1, e_2 as the first and second eigenvector fields of the diffusion tensor field, a streamsurface as we defined earlier is integrated on a 2D distribution in Euclidean space \mathbb{R}^3 that assigns to each $m \in \mathbb{R}^3$ the 2D linear subspace spanned by e_1 and e_2 .

Theorem 3.2.1 *Given C^∞ orthogonal vector fields e_1 and e_2 on \mathbb{R}^3 , the 2D distribution η on \mathbb{R}^3 that assigns to each $m \in \mathbb{R}^3$ the 2D linear subspace spanned by $e_1(m)$ and $e_2(m)$ is completely integrable if and only if $[e_1, e_2] \in \eta$.*

Proof From Frobenius' Theorem, we know that a C^∞ distribution is completely integrable if and only if it is involutive. Thus we only need to prove that if $[e_1, e_2] \in \eta$, then for all C^∞ vector fields U, V which belong to η , $[U, V] \in \eta$.

Since $e_1(m)$ and $e_2(m)$ are orthogonal for all $m \in \mathbb{R}^3$, we can write

$$U(m) = a(m)e_1(m) + b(m)e_2(m)$$

$$V(m) = c(m)e_1(m) + d(m)e_2(m)$$

Thus,

$$\begin{aligned} [U, V] &= [a \cdot e_1 + b \cdot e_2, c \cdot e_1 + d \cdot e_2] \\ &= [a \cdot e_1, c \cdot e_1] + [a \cdot e_1, d \cdot e_2] + [b \cdot e_2, c \cdot e_1] + [b \cdot e_2, d \cdot e_2] \\ &= [a \cdot e_1, d \cdot e_2] - [c \cdot e_1, b \cdot e_2] \\ &= a \cdot d[e_1, e_2] + a(e_1\{d\})e_2 - d(e_2\{a\})e_1 - c \cdot b[e_1, e_2] - c(e_1\{b\})e_2 + b(e_2\{c\})e_1 \\ &\in \eta \end{aligned}$$

Conversely, if η is completely integrable, then for all $U, V \in \eta$, $[U, V] \in \eta$. Let $U = e_1, V = e_2$, we have $[e_1, e_2] \in \eta$.

Surface Generation

We design the surface generation algorithm to try to follow the planar structure suggested by the planar anisotropy. We define $T(v)$ as the diffusion tensor at point v and $P(v)$ as the plane that contains the major and medium eigenvectors of $T(v)$ as well as v . $P(v)$ is tangent to the streamsurface.

1. Starting from a seed point, v , we set the initial directions radially along six evenly distributed directions within $P(v)$, and extend curved edges to follow the shape of the surface (see step 3). We then generate a triangle for every pair of the neighboring edges (two edges are neighbors if their projections on $P(u)$ are neighbors).
2. For every new vertex u extended from step 1, we project to $P(u)$ the existing triangles that are attached to u , and then extend curved edges (see step 3) from the initial directions on $P(u)$ that are not covered by the existing triangles. Then we generate a triangle for every pair of neighboring edges. This step is repeated for every newly generated vertex until the terminating conditions in step 4 are met.
3. We extend the edges from a vertex u by integrating in the 2D vector field V , which is defined on the plane P_1 that is both perpendicular to $P(u)$ and contains the initial direction of extension in $P(u)$. V is defined at every point x in P_1 as the the linear combination of the normalized major and medium eigenvector of $T(x)$ that lies within P_1 . We ensure the consistency of the integration directions by swapping $V(x)$ if the dot product of $V(x)$ and the initial direction is negative. We use second order Runge-Kutta for the integration.
4. The extension stops when it goes outside the data boundary, hits a low planar anisotropy region, or enters a region of low signal-to-noise ratio.

Figure 3.8 illustrates the order of the first few extensions starting from a seed point.

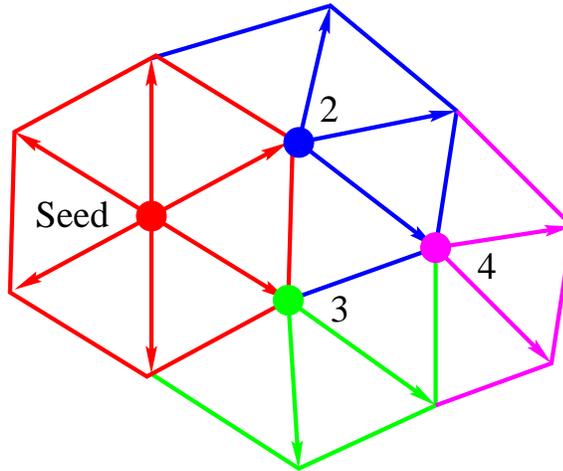


Figure 3.8: Expanding a streamsurface. Begin at point labeled Seed, add the red edges; repeat the expansion at points 2, 3, 4, etc., as many times as possible.

To verify the algorithm for generating the streamsurfaces, one can construct a phantom diffusion tensor field with known geometry. For example, if we design the eigenvectors and eigenvalues as

$$e_1(x, y, z) = (1, 0, 0)$$

$$e_2(x, y, z) = (0, \cos(x), \sin(x))$$

$$e_3(x, y, z) = (0, -\cos(x), \sin(x))$$

$$\lambda_1(x, y, z) = 1$$

$$\lambda_2(x, y, z) = 0.8$$

$$\lambda_3(x, y, z) = 0.1,$$

and construct the diffusion tensor field from these values, then for $-10 < x, y, z < 10$, the streamsurface should rotate along the X axis.

Table 3.2: Parameters used to select streamsurfaces selection for Figures 3.2, 3.7, 3.9, and 3.10.

Surface size	$> 7.00 \text{ mm}^2$
Average planar anisotropy	> 0.30
T_s used in D_s for streamsurfaces	1.34 mm
Distance between surfaces, D_s	$> 8.9 \text{ mm}$

Culling

As in streamtube selection, we use the following three criteria in surface selection: the area of the surface, the average planar anisotropy of the surface, and the distance between this surface and other selected surfaces. We eliminate a surface if it is too small, too low in average planar anisotropy, or too close to other selected surfaces.

The distance between two surfaces is defined as

$$D_s = \frac{\iint_{\Omega} \max(\text{dist}(\omega) - T_s, 0) d\omega}{\iint_{\Omega} \max\left(\frac{\text{dist}(\omega) - T_s}{|\text{dist}(\omega) - T_s|}, 0\right) d\omega} \quad (3.2)$$

Like the definition of distance between curves, this metric gives the average above-threshold distance between two surfaces.

Table 3.2 shows the thresholds for culling streamsurfaces for Figures 3.2, 3.7, 3.9, and 3.10.

3.2.4 Anatomical Landmarks

Feedback from preliminary results showed that biologists can explore images more effectively if they can see some familiar anatomical structures. Thus, we generate isosurfaces of a few anatomy features from T_2 -weighted images using AVS (Upson et al., 1989).

Water has a higher T_2 -weighted intensity than white and gray matter; white and gray matter, in turn, have higher T_2 -weighted intensities than air. We generate an approximate boundary between fluid and other tissue that echoes the shape of the

ventricles by generating isosurfaces at a certain level between the intensity of water and that of other tissue. While these shapes are not precise, their presence makes the images much easier to interpret. Surfaces of other anatomical structures, such as blood vessels or tumors, could also be generated for certain applications.

3.2.5 Results and Discussion

We applied our method to a human brain data set with $256 \times 256 \times 40$ voxels and a resolution of $0.89 \text{ mm} \times 0.89 \text{ mm} \times 3.2 \text{ mm}$. Each slice of the data was acquired at a resolution of 128×128 and zero-filled. Figures 3.9 and 3.10 shows a visualization of the human brain on a desktop display (the indications of skull and eyes help orient the image). Through the semi-transparent skull surface, we can see the ventricle (colored blue) in the middle, and the streamtubes both around it and throughout the space.

The image shows that the streamtubes correlate well with major neural structures. The corpus callosum, depicted by the streamtubes running across the top of the ventricle, contains almost all the neurons that cross the brain from one hemisphere to the other. The internal capsule, clearly shown from the side view in Figure 3.10, is the second most obvious white matter structure visible on a myelin stain and is the major fiber tract that carries information between the cortex and the brainstem. The cerebral peduncles are the continuation of the internal capsule as it runs down into the midbrain.

All our users were eager for more detail in the visualization, saying that more anatomical detail and streamtubes in the regions of interest would enhance their understanding. To this end, Figure 3.11 shows the visualization of the same human brain data set used in Figure 3.9 with the average linear anisotropy threshold set to 0.20: the number of streamtubes increases from 472 in Figure 3.9 to 4,538 in Figure 3.11. Setting more seed points will also increase the number of streamtubes. However, the level of detail we can visualize is fundamentally limited by the voxel size of the data set. Setting many seed points within one voxel will yield redundant

streamtubes that will be discarded by the culling algorithm.

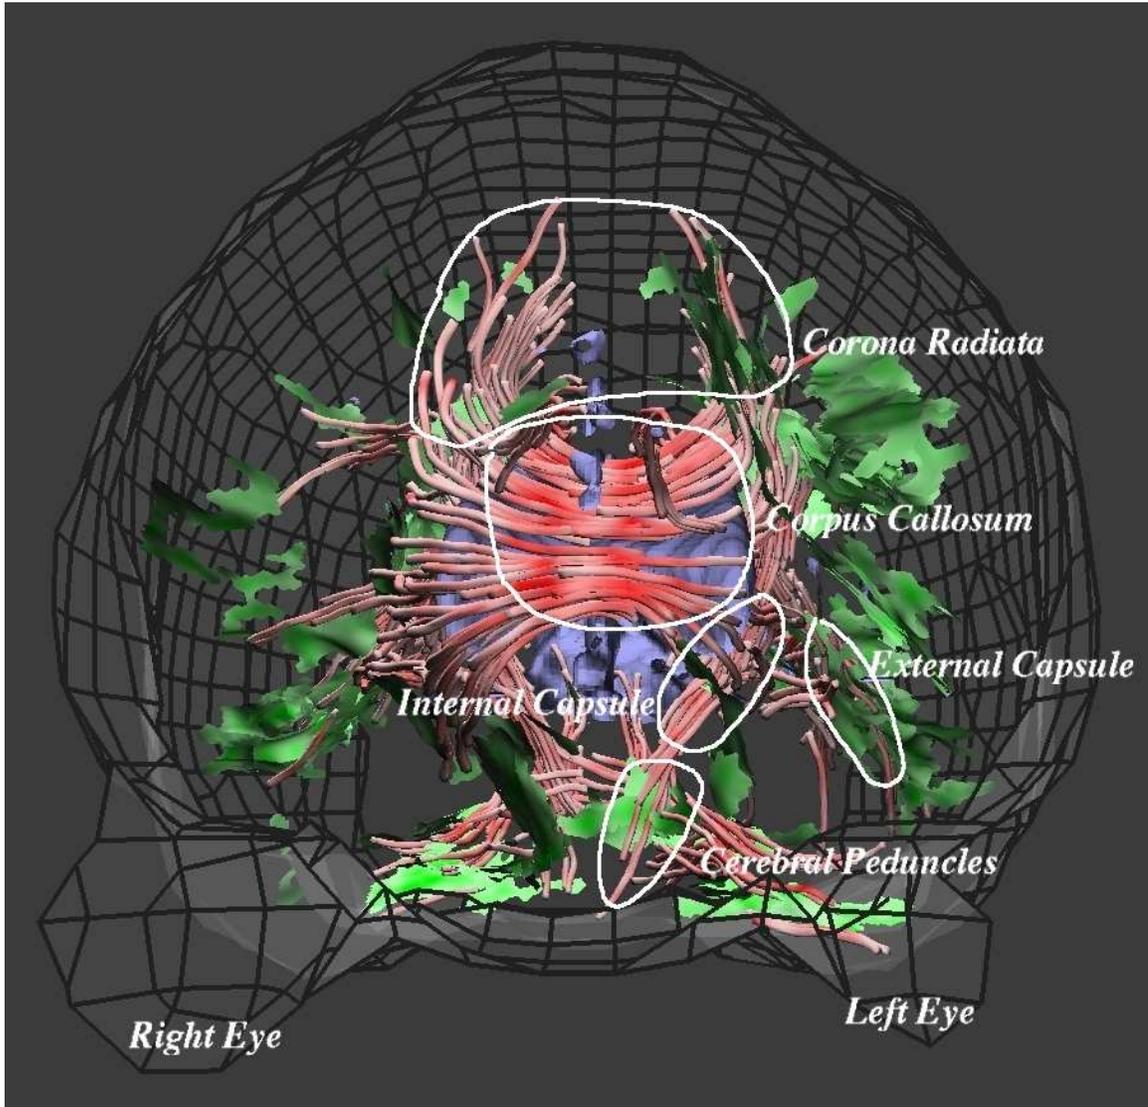


Figure 3.9: Front view of the human brain data set visualization. The culling parameters are shown in Tables 3.1 and 3.2.

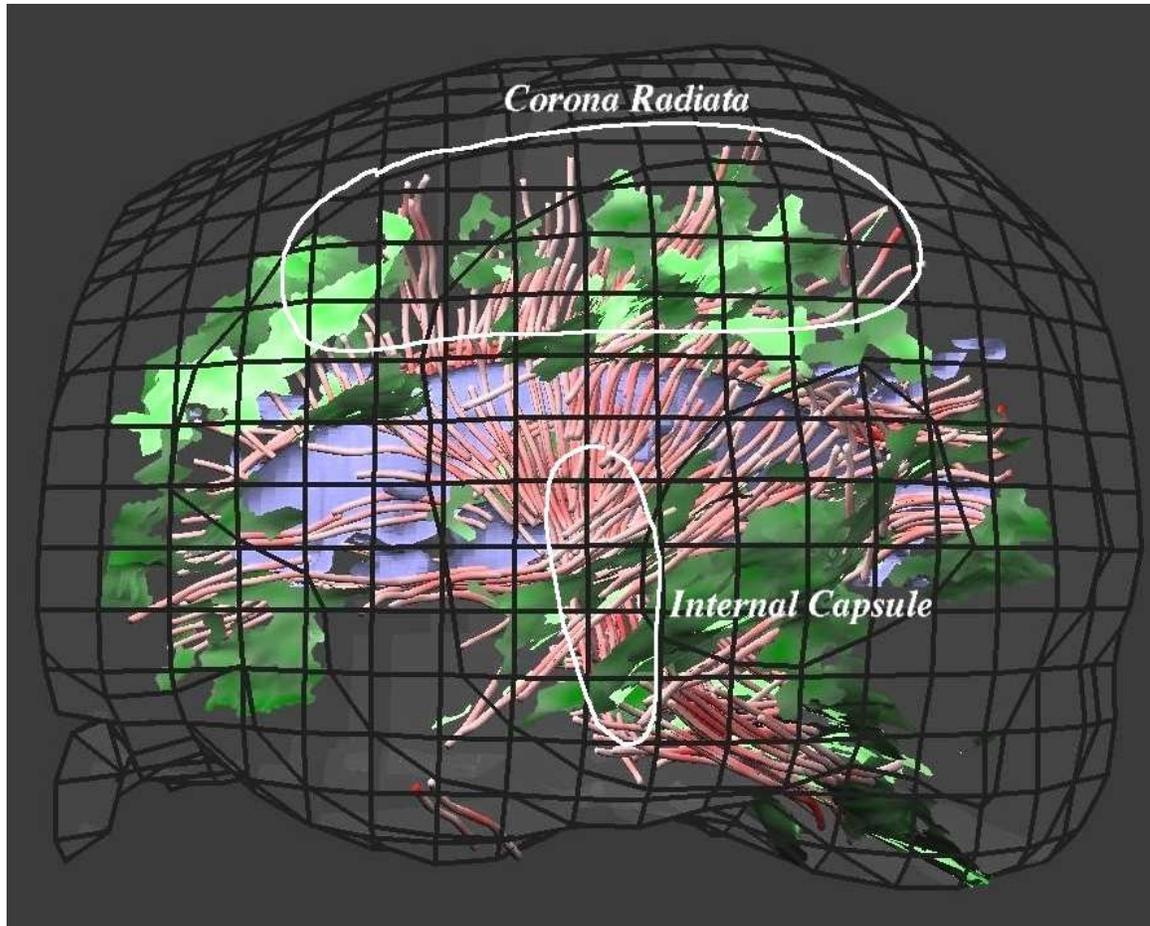


Figure 3.10: Side view of the human brain data set visualization. The culling parameters are the same as those used in Figure 3.9.

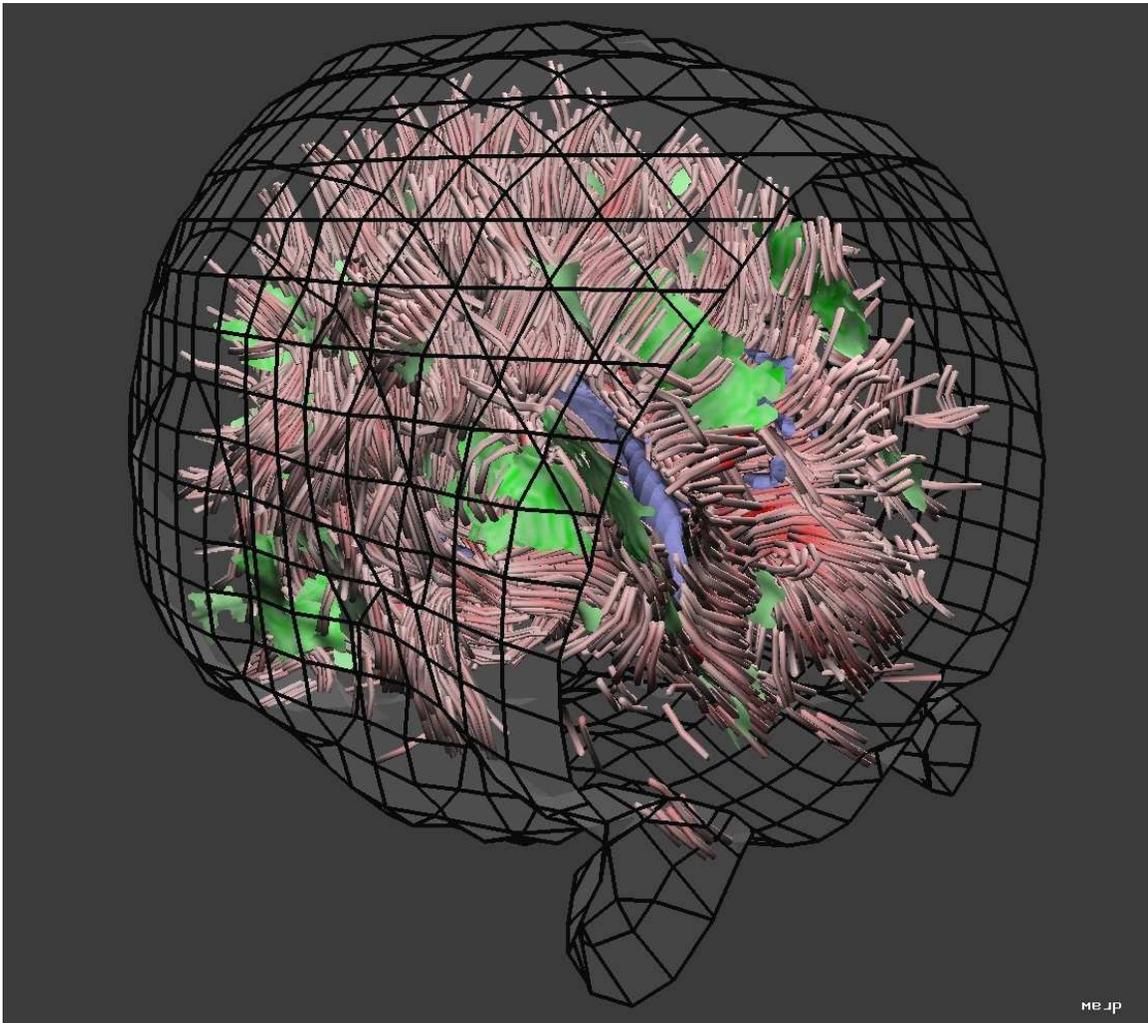


Figure 3.11: Human brain data set visualization with average linear anisotropy threshold set to 0.20 instead of 0.30. The number of streamtubes increases from 472 to 4,538.

Chapter 4

Visualization Schemes

The streamtubes, streamsurfaces and anatomical landmarks are complex 3D shapes that are difficult to grasp in a still picture. In this Chapter, I present three different display and interaction schemes for exploring these complicated brain models.

In Section 4.1, I will discuss the use of two virtual environment systems: the CAVE and the fishtank VR system. These environments employ stereo and head-tracking for the display of our surface models for the streamtubes, streamsurfaces and the ventricles, and each includes its own interaction scheme. A qualitative study shows that the stereo vision helps users understand the model, and the users prefer the small but crisp display from the fishtank VR.

In Section 4.2, we describe a thread and halo technique suitable for interactive volume rendering of thin linear structures, together with a number of components for exploring the brain models and data. Volume rendering has the defining property of mapping from the tensor field attributes to a rendered image, without introducing geometry. However, combining the volume scene with fiber tracts together creates a more informative image.

In Section 4.3, we present work toward creating color rapid prototyping plaster models as visualization tools for the brain models. The physical representations of these virtual models allows the users to review the data with a very robust, natural

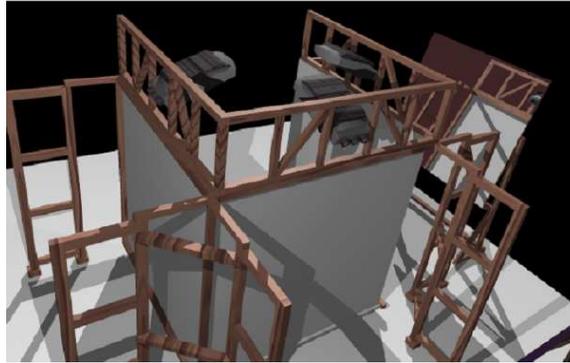


Figure 4.1: A model of the CAVE viewed from behind the back wall. The user stands inside the three rear-projected walls and sees a virtual world on the head-tracked stereo displays.

and fast haptic interface: their hands.

4.1 Virtual Reality

Our CAVE, shown in Figure 4.1, is a virtual reality environment in the shape of a cube with rear-projected front and side walls and a front-projected floor. Each of the four displays has a resolution of 1024×768 . The user wears LCD shutter glasses with an attached tracker and uses the wand, a mouse-like input device that is also tracked, to interact with the virtual world.

Compared to still pictures and desktop display, an immersive virtual environment such as the CAVE (Cruz-Neira et al., 1993) offers the advantages of a head-tracked stereo display, a large display surface, and interactivity. We take advantage of these features by showing our visualization in the CAVE.

We put geometric models including streamtubes, streamsurfaces, and anatomical landmarks at approximately the center of the CAVE, and orient them so that the front of the brain always faces the opening of the CAVE. We then display a 2D section that carries a slice of the $T2$ -weighted MRI data, the position and orientation of which can be chosen interactively. We can also use the 2D section to carry any other data set



Figure 4.2: The virtual reality applications run in the CAVE (left) and fishtank VR display (right).

that is registered with the 3D models. The user clicks a wand button to change the orientation of the slice among axial, coronal and sagittal, and uses two other buttons to move it back and forth. We also draw a yellow line akin to a laser pointer from the wand to where the wand is pointing (since physically pointing out an object is ineffective in the CAVE).

For the environment, we texture-map the virtual walls and the floor in the same positions as the projection walls and floor, and set up a table beneath the geometric models.

Figure 4.2 left shows two users interacting with the brain visualization in the CAVE.

We also set up a fishtank VR display consisting of a 22" (20" viewable) Mitsubishi Diamond Pro 2070-SB desktop monitor. Users wear a pair of LCD shutter glasses that support active stereo viewing; the glasses have a tracker attached to relay their position and orientation to the computer. We used a Polhemus 3Space Fastrak. The fishtank VR display is set to the same resolution as the CAVE walls: 1024×768 pixels.

Figure 4.2 right shows two users studying the brain model on the fishtank VR display.

We did a subjective usefulness study comparing the two virtual reality systems (Demiralp et al., 2003). We had five domain-expert users. They were asked to use the brain application both in CAVE and Fish Tank VR platforms. While using the application, they were asked to compare the advantages of each platform for their purposes. They did so by talking to us while using the application. Most often we questioned their arguments by bringing counter-arguments, which helped to expose the reasoning behind the users' observations. The users were then asked to give an overall preference for one of the two VR platforms.

All five users used the application in both platforms. Three of the users started with the fish tank version of the application and the other two with the CAVE version.

Results are shown in Table 4.1. Overall, one user preferred CAVE and four of them preferred Fish Tank VR.

We also conducted quantitative comparison of CAVE and Fish Tank VR. The results are published in (Demiralp et al., 2006).

4.2 Streamtubes Visualization with Volume Rendering

In (Wenger et al., 2004), we describes a thread and halo technique suitable for interactive volume rendering of thin linear structures together with a number of components that make it useful for visualization of multivalued 3D scientific data. That work is excerpted here. Two scientific applications drive our volume-rendering work: understanding brain anatomy and pathology, and understanding blood flow in coronary arteries. These driving applications have provided the problems, and, as Brooks suggests, the extent to which our application facilitates the solution of these problems helps to evaluate and guide our algorithm and tool development (Brooks, 1996).

Creating comprehensive and accurate visualizations for exploring 3D multivalued data is challenging. The first challenge is to create visualizations in which the data

Table 4.1: Advantages reported for the CAVE and fishtank VR.

	User				
	1	2	3	4	5
<u>Advantages reported for the CAVE:</u>					
Has bigger models, one can see more.		✓	✓	✓	✓
Has larger field of view.		✓			
More suitable for gestural expression and natural interaction.		✓			
Possible to walk around.			✓		
<u>Advantages reported for the fishtank VR:</u>					
Has sharper and crisper images.	✓		✓	✓	✓
Displays the data more compactly, spatial relationships between the structures are easier to see.	✓				
Feels more comfortable, non-claustrophobic and sitting is better than standing.	✓	✓			✓
Works better for collaboration, especially with two people.	✓				
Pointing to objects on the screen is easier.		✓		✓	
More time efficient to use; doctors prefer to work-and-go.					✓
Would work better for telemedicine-like collaboration.			✓		
More intuitive for surgery planning because doctors are used to working with brain models at full scale or smaller.			✓		
<u>Overall preferred display:</u>					
The CAVE.		✓			
The fishtank VR.	✓		✓	✓	✓

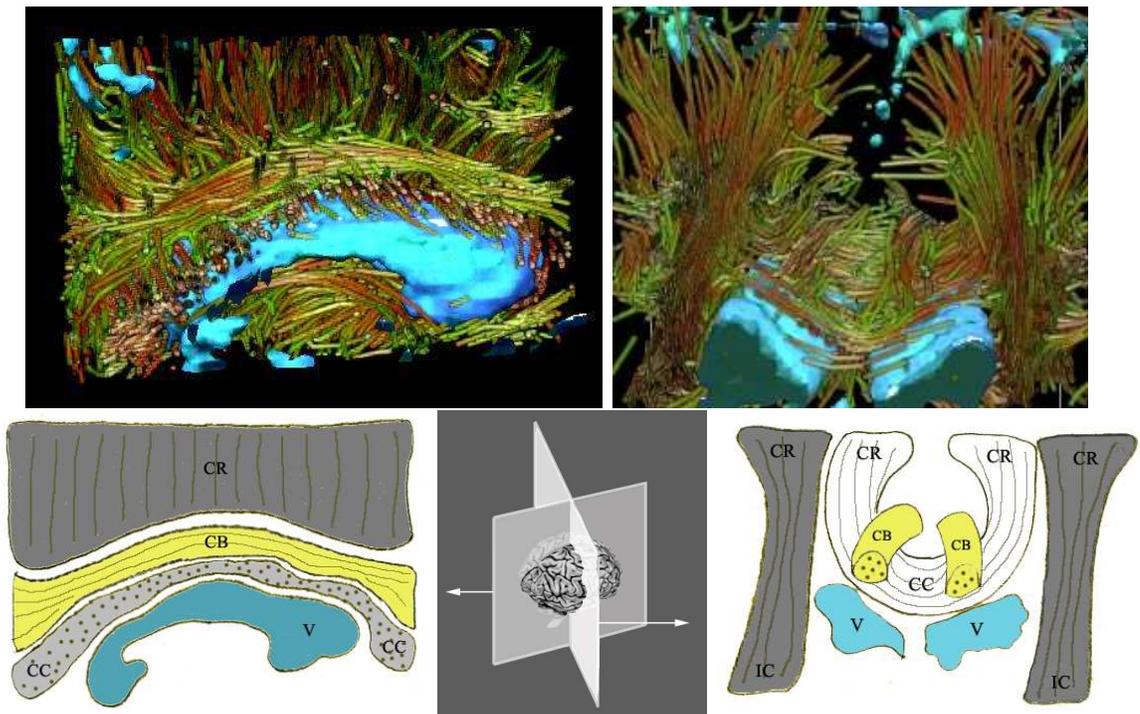


Figure 4.3: Interactive renderings of a human brain dataset. The renderings (top) show collections of threads consistent with major white-matter structures: IC=internal capsule, CR=corona radiata, CB=cingulum bundle, CC=corpus callosum, diagrammed on the bottom. Components of the tensor-valued data control thread direction, color, and density as described in the text. Direct volume rendering simultaneously shows the ventricles (labeled V) in blue for anatomical context.

nearer to the viewer does not excessively obscure data farther away. The second challenge is to represent many values and their interrelationships at each spatial location.

Interactive volume rendering with user controlled transfer functions can provide a promising approach for overcoming much of the obscuration problem. By using transparency effectively, transfer functions can be designed that show important features in a dataset throughout a volume. Interactive control allows a scientist to weight the relative importance (usually tied to an opacity level) of data values or features while exploring the dataset.

In our work, we rely heavily on volume rendering techniques. In fact, we use a multi-layer volume rendering approach, similar to (Hauser et al., 2001) to enable us to fully represent multivalued datasets. We also make heavy use of transfer functions and provide interactive controls that are tailored to our application domains.

The key contribution of our work is a clear volumetric vector-field representation that can be rendered interactively. Datasets that can benefit from this representation are common in fluid flow research and medical imaging. Our thread and halo representation, shown in Figure 4.3, together with direct volume rendering, provides clear visual indications of complex linear forms, depth relationships among multiple densely packed threads, and changing data values along the length of the thread. In this paper, we demonstrate that our threads and halos technique can be incorporated into a multi-layer volume rendering scheme and displayed at interactive frame rates on modern consumer graphics cards. We also describe the benefits of such an implementation for our scientific collaborators.

In the next section we discuss related work. We then describe our layered volume rendering framework, threads and halos, and our interactive controls. Results for our two driving applications are then presented and discussed along with some conclusions from this work.

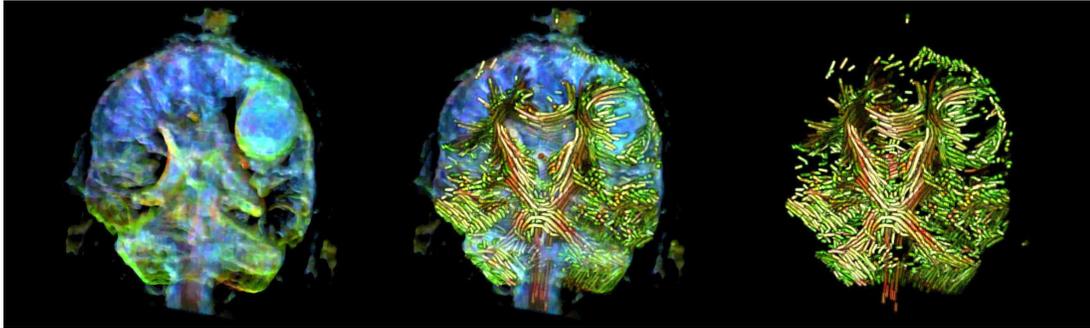


Figure 4.4: Left: a direct-volume-rendered layer showing regions with different diffusion anisotropy. Right: a thread layer showing both the anisotropy and diffusion direction. The two layers are combined in the center image which shows significantly more information and requires little additional visual bandwidth.

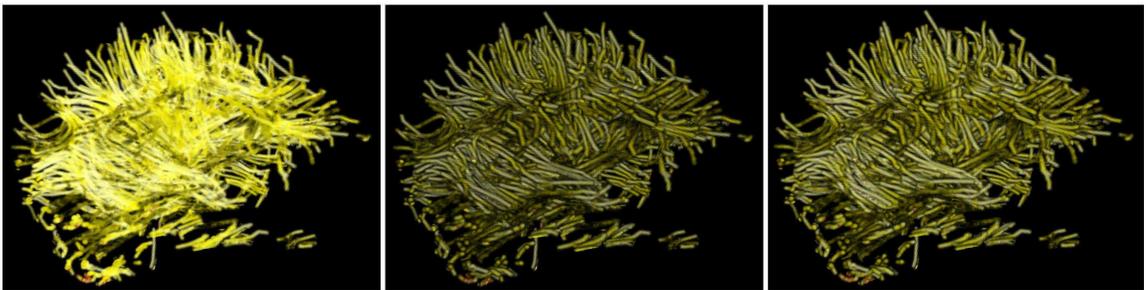


Figure 4.5: Three thread volumes showing brain connectivity information (the front of the head points right). Shown without halos on the left, with halos in the center, and after slightly shifting the halos away from the viewer to brighten the threads on the right.

4.2.1 Related Work

Below we survey relevant work in vector field visualization, and hardware-accelerated volume rendering.

Visualization of Vector Fields

Of the extensive work on creating effective vector field visualizations, the following two papers are most closely related to our work. Interrante and Grosch (Interrante and Grosch, 1997) visualized 3D flow with volume line integral convolution (LIC). As they

demonstrated with offline rendering, their visibility-impeding halos improve depth perception and help make complex 3D structures easier to analyze. Our technique builds on this work to produce a similar effect interactively.

Zöckler *et al.* (Zöckler et al., 1996) introduced illuminated field lines to visualize 3D vector fields. Our illuminated thread representation is similar, but our volumetric rendering approach renders at a rate independent of the tube or line complexity and combines with our other volumetric layers to create visualizations that convey more information.

Li *et al.* (Li et al., 2003) presented a volumetric texture based method to visualize 3D vector fields. They scan convert properties of the streamlines into a volumetric texture and then use these properties to look up color and opacity information in an appearance texture. Our approach differs in that our threads, while individually less visually complex, are much thinner and more densely packed; we also represent data more complicated than vector fields.

Hardware-Accelerated Volume Rendering

Cabral *et al.* (Cabral et al., 1994) introduced a 3D texture approach for volume rendering using view-aligned slices that exploit trilinear interpolation. In addition, we make significant use of hardware texture compression to reduce texture memory consumption.

There are also several volume-rendering implementations that make use of dedicated hardware (Wu et al., 2003) or distributed hardware (Parker et al., 1999), that are capable of visualizing multi-valued volumetric datasets.

Our multi-layer volume rendering is closely related to the two-level volume rendering presented by Hauser *et al.* (Hauser et al., 2001). In this scheme, the two levels are an object level and a global level. Different rendering styles, such as direct volume rendering, maximum intensity projection, or value integration, are used for each level.

Our system is based on the same concept of rendering multiple volumes of information into the same visualization space. However, we use separate volumetric datasets for our layers of information, rather than classifying a single volume of data as either focus or context. In some cases our halos and threads could be conceptualized as together forming a *focus* level for the visualization with any additional direct volume rendered layers forming the *context* level.

Kniss *et al.* (Kniss et al., 2002a) use interactive transfer functions operating on directional derivative values to select boundaries in scalar-valued datasets. We use this technique to visualize our scalar-valued datasets, although with less sophisticated interactive manipulation widgets. In (Kniss et al., 2002b), multi-dimensional transfer functions and dual-domain interaction are applied to multivariate meteorological data. They found, as we did, that multi-dimensional transfer functions provide a powerful tool to explore multivariate datasets.

Lum and Ma (Lum and Ma, 2002) implemented a hardware-accelerated parallel nonphotorealistic volume renderer that uses multi-pass rendering on consumer-level graphics cards. Their system emphasizes edges or depth ordering using artistically motivated techniques. Like Lum and Ma, we utilize multiple rendering passes to enhance visual cues, but our rendering is targeted to exploratory visualization of multi-valued data, which has significant implications for the interface, implementation, and results.

Stompel *et al.* (Stompel et al., 2002) use nonphotorealistic (NPR) volume rendering to more effectively visualize multivariate volumetric data. They use stroke rendering to display a vector field simultaneously with a scalar field and produce several NPR effects, including silhouettes. Their silhouettes help to emphasize depth discontinuities just like our halos but will not work for features as small as our threads because a reliable gradient cannot be calculated.

Hair, Fur, and Thread Rendering

Several hair and fur rendering algorithms inspired our work. Kajiya and Kay (Kajiya and Kay, 1989) introduced texels to render realistic-looking fur. Kajiya and Kay also developed a Phong-like lighting model for fur; our approach is similar but targets free-floating threads. Instead of providing parameters for lighting, we store derived values from the multivalued datasets along with tangent and density values throughout the volume.

Lengyel (Lengyel, 2000) uses a volumetric texture approach to render short threads in real time. Unlike his short threads, our data-defined threads remain individually distinguishable. Like Lengyel, we use Banks’ (Banks, 1994) hair-lighting model but with a different implementation appropriate for volume rendering.

4.2.2 A Layered Volume-Rendering Framework

Our visualization framework has four steps. We begin with primary multivalued volumetric data.

Calculate Derived Datasets

Since the primary data is often difficult to interpret directly, our first step is to calculate derived volumes of data with more intuitive interpretations. For example, DTI datasets are second-order tensor fields. It is often useful to decompose these into several scalar and vector fields to reduce the problem of rendering a tensor field into one of rendering several simpler fields.

Define Visual Abstractions

In the abstraction step, we group the data volumes into layers of visual representations that are filtered into and stored in separate volumes. The simplest abstraction is a straight mapping of a scalar value, such as speed, to a volume. Threads are another

abstraction that are good at illustrating vector fields. Any conversion of the derived data to a clearer visual representation fits into this step of the framework.

Map Data with Interactive Transfer Functions

The mapping step defines transfer functions that assign color and opacity to the volume layers and shader programs that control lighting parameters.

Visualize and Explore

In the final step of the framework, we render the multiple volumes together in the same visualization space and use interaction widgets to control the appearance of each layer.

4.2.3 Threads and Halos

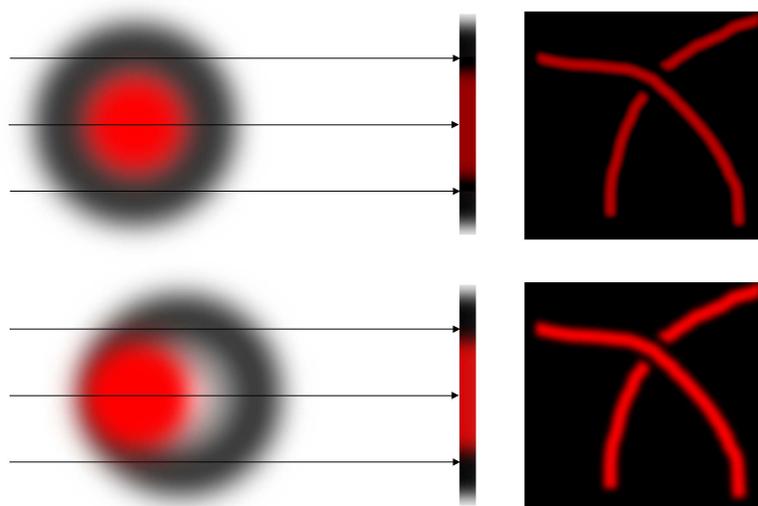


Figure 4.6: Black arrows illustrate the path of virtual viewing rays through a cross section of a volumetric thread and halo. When the rays, coming from the viewer's perspective, pass through part of the dark halo volume before reaching the red thread, the resulting color is a darker red than when the halo volume is shifted away from the viewer (the bottom case).

We represent continuous directional information using threads. We densely populate the volume with threads so as to represent as much of the underlying data as possible. To clarify individual threads, we add a “halo,” modeled after those presented by Interrante and Grosch (Interrante and Grosch, 1997). Compare the images in Figure 4.5 to see the depth clarifying effect of using halos with the threads.

The threads and halos are each stored in a volume texture that is precomputed. There is exactly one halo in the halo volume for each thread in the thread volume. Each halo follows the same path as its thread but has a slightly larger radius. Figure 4.6 demonstrates how the halos extend beyond the sides of the threads when rendered to obscure any geometry behind them.

Unfortunately, the halo also slightly obscures the thread from a frontal view. This has the effect of darkening the thread rather than completely hiding it since the halo is semi-transparent and there is only a small amount of it in front of the thread. This is seen in the darker, middle image in Figure 4.5. We compensate for the darkening effect to produce images like the one on the right of the figure by shifting the entire halo volume so that there is less of the halo showing between the thread and the viewer.

As illustrated in Figure 4.6, the virtual viewing rays that pass through only the halo on the silhouette of each thread will produce a black color. If a portion of the halo exists in front of the thread with respect to the viewer, as in the top of the figure, then the viewing ray passing through this area will produce a color partly saturated with red and partly with black, resulting in dark red. If the halo is offset away from the viewer at a distance of one voxel, then the viewing rays pass through far less, or none in the case of the middle ray, of the halo volume before reaching the red color of the thread. The result is a brighter red color for the thread.

The threads and halos are filtered into volumes using a cubic B-spline filter that has a support of four voxels for the thread (see Figure 4.7) and six voxels for the halos (see Figure 4.8). Since the threads and halos are represented in volumes, the

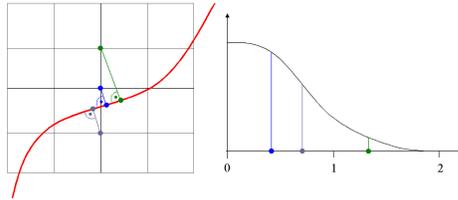


Figure 4.7: Filtering a thread into a volume (2D view). For each voxel within a radius of two voxels from a thread, we use the shortest distance to the thread as the input to the filter (at right). The grid on the left and the horizontal axis on the right both show single-voxel spacing.

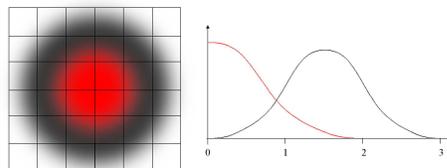


Figure 4.8: Filtering a halo into a volume (2D view); red depicts the thread and black the halo around it. The red curve is the filter for the thread and the black curve is the filter with which the halo is generated. The grid on the left and the horizontal axis on the right both show single-voxel spacing.

rendering time is independent of the number of threads displayed. However, the diameter of the threads is limited by the resolution of the volume.

Lighting for the threads is calculated using a restricted version of the lighting model in (Banks, 1994), which defines intensity I as

$$I = k_d I_t \left(\sqrt{1 - (T \cdot L)^2} \right)^p + k_s \left(\sqrt{1 - (T \cdot H)^2} \right)^n \quad (4.1)$$

Here I_t is the combined output color from the transfer functions, T the tangent, L the light vector, H the vector halfway between the light vector and the view vector, n the specular exponent, p the excess-brightness diffuse exponent, and k_d and k_s the diffuse and specular contributions respectively. Our restricted lighting model assumes directional lighting and a constant halfway vector for the entire volume.

We have implemented several thread lighting models, including those of Kajiya and Kay (Kajiya and Kay, 1989), and Banks (Banks, 1994) with excess brightness

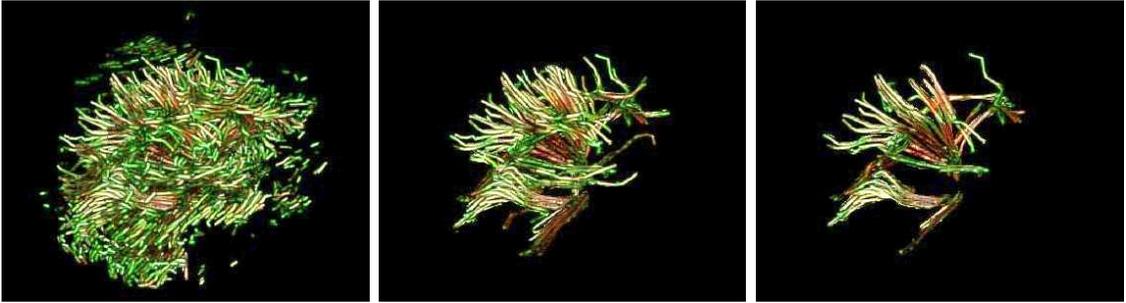


Figure 4.9: A sequence of renderings of a thread density volume with increasing length threshold from left to right. The rightmost image shows only long threads.

exponents of $p = 2$ and $p = 4$. Kajiya and Kay’s lighting model is similar to a Banks model with $p = 1$. Banks actually uses a value of around 4.8. With a small exponent, the threads become brighter and the lighting is less dramatic. We found $p = 2$ to be a good value for our applications and also a speedy one, since it does not require a square-root calculation.

4.2.4 Layering Volumes

Our volume-renderer implementation uses a single stack of view-aligned texture-mapped slices, rendered back to front and blended with weights $1 - \alpha$ and α . Each slice is rendered multiple times, once for each volume layer. Layers of each slice are also blended with the weights $1 - \alpha$ and α . We render our direct volume rendered layers first, the halos second, and the threads third. For our applications, this is equivalent to rendering the layers in the order of the scale of their largest structures. The direct-volume rendered layers tend to reveal large structures and surfaces that are easy to make out underneath the more finely detailed thread and halo layers. Our intuition is that in most cases this represents an acceptable heuristic to use when determining layer ordering.

For direct volume rendered layers, a Phong lighting model (Eq. (4.2)) is used. As in the thread lighting model described above, I_t is the combined output color from the transfer functions, N is the normal, L the light vector, H the halfway vector, k_a

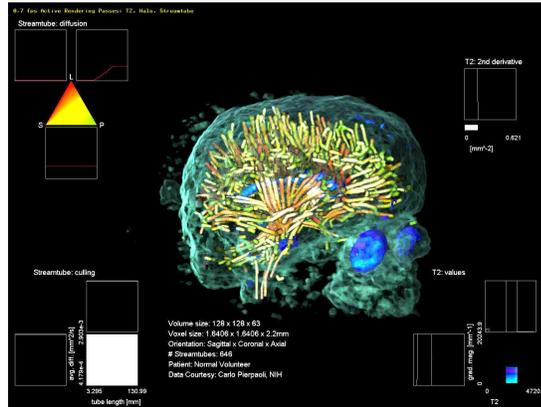


Figure 4.10: The interactive exploration tool. Clockwise from upper left are a 2D barycentric widget, a 1D widget, a 2D Cartesian widget, and a 2D Cartesian culling widget.

the ambient contribution, k_d the diffuse contribution, k_s the specular contribution, and n the specular exponent:

$$I = k_a I_t + k_d I_t (N \cdot L) + k_s (N \cdot H)^n \quad (4.2)$$

Visualizing multiple layers of volumetric data requires an extensive amount of texture memory. We utilize the OpenGL extension `ARB_texture_compression`, which provides a 4:1 compression ratio. Thus, a 256^3 eight bit per channel RGBA texture can be reduced from 64MB to 16MB. With this scheme, we can fit the multiple volume textures required in memory on commodity graphics cards.

4.2.5 Exploratory Culling and Transfer Function Manipulation

Interactive editing of transfer functions has become commonplace in volume rendering applications. We describe several application specific transfer function modes in the following section.

Exploratory culling applies a transfer function to an entire thread and halo based on attributes of the linear structure it represents. This is the approach that Wei *et al.*

take in (Wei et al., 2001) to cull discrete particle tracks based on the track’s energy. Doleisch *et al.* (Doleisch et al., 2003) introduce a more general framework in which degree-of-interest functions are defined to perform similar tasks. In our approach the entire thread and its halo are classified according to a metric (e.g. thread length) which is mapped to a $0 \dots 255$ domain. Every voxel in the data volume belongs to one class of threads. The cost of classifying the threads and halos for this example is an additional byte per voxel and a 1D transfer function that takes up 256 additional bytes of texture memory.

In our brain visualizations, both long and short threads are important and provide different types of insight into the data. We use this culling feature to interactively select a subset of threads to display based on their average diffusion rate or their length, as seen in Figure 4.9. This approach is a significant advance over the state of the art in this application area. Similar culling in Zhang *et al.*’s (Zhang et al., 2003) approach required an entire preprocessing step taking between several minutes to several hours.

We provide several on-screen widgets, shown in Figure 4.10, to control transfer functions in the form of 1D, 2D, and 2D barycentric controls. Colors are manipulated within the hue, saturation, value, and transparency ($HSV\alpha$) color space. For the 1D and 2D transfer function widgets, color and opacity can be controlled interactively along each axis. In the multidimensional cases, the colors of the axes are averaged, while the opacities are combined multiplicatively. In our informal trials, these combination methods seemed most intuitive. The 2D barycentric manipulation widget, shown in the top left of Figure 4.10, is ideal for the brain visualization application since the space maps naturally to the anisotropy metrics defined in (Westin et al., 1997b).

4.2.6 Results and Discussion

The data are acquired using magnetic resonance imaging (MRI) and are of two types: second-order tensor-valued water-diffusion-rate images and scalar-valued anatomical images. At each point in a volume, the tensor-valued data capture the rate at which water is diffusing through tissues. That rate is different in different areas – in regions of pure fluid, it is fast; in tissues like bone, it is slow. The rate of diffusion can also be directionally dependent, particularly in fibrous tissues like axon tracts and muscles, diffusing more quickly along the fibers than across them. The scalar-valued data are typical T2-weighted MR images.

Within the second-order tensor field measuring the water diffusion rate, each value D is a symmetric tensor with real, positive eigenvalues. From D we derive several other measures. First, three scalar anisotropy measures introduced by Westin (Westin et al., 1997b), c_l , c_p , and c_s , describe how close to a line, a plane, or a sphere the corresponding ellipsoid shape is for a given measurement. Second, the trace of D , $\text{Tr}(D)$, is equivalent to the sum of the eigenvalues of D and gives a scalar measure of the overall diffusion rate. Third, the gradient of the trace, $\nabla\text{Tr}(D)$ and its magnitude, $|\nabla\text{Tr}(D)|$, describe how the diffusion rate is changing and in what direction; we use these quantities in lighting calculations for the direct volume rendered layer i.e. in Figure 4.4.

The fourth category of derived data is a set of threads and halos through the tensor field that represent the directions of diffusion. These are calculated and distributed within the volume as described by Zhang *et al.* (Zhang et al., 2000)(Zhang et al., 2001a)(Zhang et al., 2003). They follow the direction of fastest diffusion in linear regions. In planar regions, they stay within the plane formed by the major and medium eigenvectors, following whichever is more consistent with the path to that point. They are not present in isotropic regions.

From the T2-weighted image scalar field we derive the gradient of the value and the gradient magnitude, which help define how fast the value is changing and in which

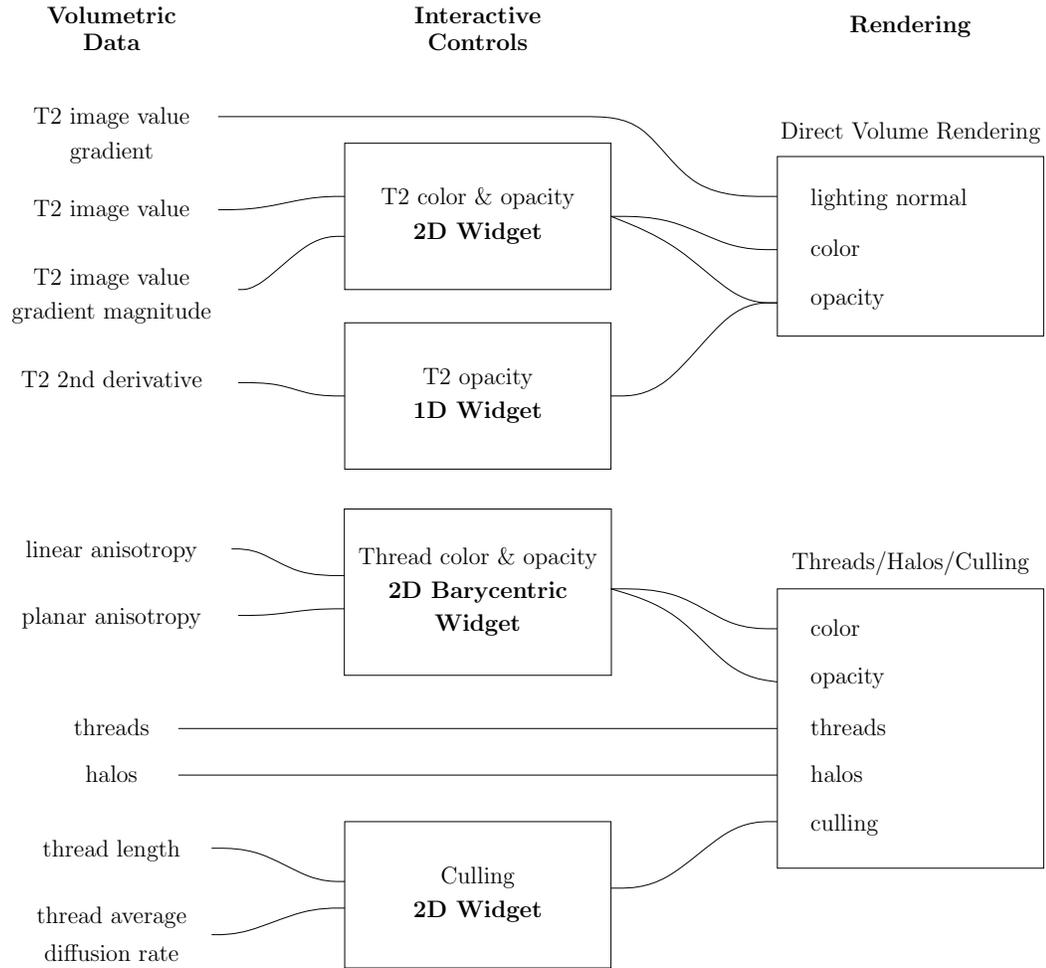


Figure 4.11: A data-flow diagram of the rendering setup for Figure 4.10.

directions. We use these quantities in lighting calculations. We also derive the second directional derivative to help define boundaries between homogeneous regions.

Figure 4.11 shows the mapping from the scalar- and tensor-valued volumes onto a direct volume rendered layer, a thread layer, and a halo layer as seen in Figure 4.10. The first layer directly renders the T2-weighted image. The hyper-intense ventricle regions were selected by interactively editing two transfer functions that are combined multiplicatively to yield α values for the layer. Color is specified only through the transfer function controlling the T2 image value and the gradient magnitude.

The second layer renders halos for the threads, and the third layer renders the

threads. The visible portions for these layers are interactively selected via three criteria. First, a transfer function maps the anisotropy metrics to an α value. For this rendering, areas of all types of anisotropy are shown. Second, each thread and halo can be selected via exploratory culling based on the thread’s length and on the average diffusion rate along it. In this rendering all threads are shown. Third, the thread density is provided directly by the precalculated thread volume. Likewise, halo density is provided directly by the halo volume. Each of the results for this dataset is rendered with 256^3 volume textures.

Our neuroscientist collaborators gained several insights into their data through these visualizations. Figure 4.3 shows detail of a diffusion dataset from a normal volunteer. A number of large white-matter structures are clearly visible, including the corpus callosum, internal capsule, corona radiata, and cingulum bundle.

Figure 4.4 shows a dataset from a patient with a brain tumor. Direct volume rendering captures the tumor as an opaque mass and threads show the diffusion direction and variation in anisotropy around it. Note the cradle of threads surrounding the tumor. Using this exploratory visualization has enabled our collaborators to discover a relationship between the different types of anisotropy around tumors. In particular, there is a notable increase in planar anisotropy (shown as green) in the area around the tumor (Zhang et al., 2004).

4.3 Rapid Prototyping

While surgeons and neurologists can use virtual reality applications to visualize different aspects of the brain data, the physical representations of those virtual models allows them to review the data with a very robust, natural, and fast haptic interface: their own hands. In this section we describe work toward creating color rapid prototyping (RP) plaster models as visualization tools for diffusion-tensor (DT) MRI analysis. Our initial results are encouraging, and end users are excited about the

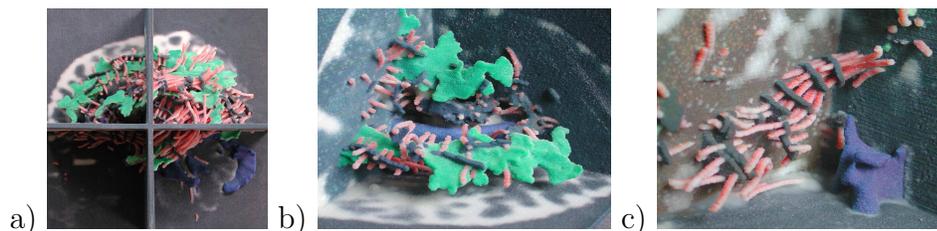


Figure 4.12: (a,b) A plaster model showing areas of linear and planar water self-diffusion obtained from DT-MR images. (c) Detail of support structures (dark gray surfaces around tubes) for the streamtubes; these surfaces are created using the second and third eigenvectors of the tensors that produce the tubes.

possibilities of this technique. For example, using these models in conjunction with digital models on the computer screen or VR environment provides a static frame of reference that helps keep users oriented during their analysis tasks.

RP has been used in visualization largely for building molecular models to test assembly possibilities (Bailey et al., 1998). Nadeau *et al.* (Nadeau and Bailey, 2000) created models of the human brain surface with the same RP techniques we use. Our approach, however, enables us to build inner brain structures.

We examine the streamtubes, streamsurfaces and ventricles models generated from a DTI dataset using tractography (Zhang et al., 2003). The images on the three orthogonal planes show slices of T2-weighted images collected with DTI.

To create our color models we use Z-Corp's Z406 printer. The digital model, in VRML format, is subdivided into horizontal layers by the printer software. These layers are then manufactured by putting down a thin layer of plaster powder and dropping colored binder at the boundaries of the model at that level. Once all the layers are built, the powder outside the boundaries of the model is vacuumed out and loose powder is removed using a fine blower. Finally, the piece is bathed in hot wax to strengthen it and enhance its colors.

The structures in the DTI models require very careful treatment. Because the long thin streamtubes often fail to support themselves during powder removal, we

inserted some supporting surfaces that interconnect neighboring streamtubes without occluding interesting features. These supports are created from the second and third eigenvectors of the diffusion tensor that creates the streamtubes, so they are perpendicular by definition (see Figure 4.12(c)). We arrived at this methodology after several tests, including building thicker tubes and increasing their number so they supported one another. Using information already present in the DTI data, we have been able to create models with better structural stability.

These early stages of development have highlighted some important issues. For example, our visualizations involve organic, free-form shapes, whereas current RP technology is designed for models with more regular shapes, such as mechanical parts and molecular models. Also, the printing and cleaning process can take as much as 12 hours for complicated brain models measuring up to $8'' \times 8'' \times 10''$. However, our initial experiments suggest that this technology has the great advantage of exploiting users' familiarity with physical models: they recognize the utility of holding them in their hands when studying them. Providing scientists with these models enhances the use and analysis of their digital counterparts. To quote one of the doctors who experimented with these models: *"These physical models complement displays in digital format by providing a hard-copy frame of reference that can be touched and manipulated with optimum hand-eye coordination and immediate results."* We believe this type of physical model can be very useful in preoperative planning when used as a quick reference for structure identification.

4.4 Conclusions

The streamtubes, streamsurfaces and the anatomical landmarks are complex 3D shapes that are difficult to grasp in a still picture. In this Chapter, I present three different display and interaction schemes for exploring these complicated brain models.

For surface models, virtual reality systems help conveying the 3D information. Fish tank VR system is considered by most users to be superior to the CAVE.

Volume rendering can support multiple components for visualization of multi-valued 3D scientific data. Combining the volume rendering of the scalar data sets with the thread and halo technique for interactive volume rendering of thin linear structures provides a comprehensive visualization of the models and the data.

The physical models complement displays in digital format by providing a hard-copy frame of reference that can be touched and manipulated with optimum hand-eye coordination and immediate results.

Chapter 5

Fiber Bundle Model

When using diffusion tensor models of the brain, it is often helpful to be able to group white matter into anatomically relevant fiber bundles to study individual brain structures. Traditionally, this task is accomplished by manually selecting DTI fiber tracts. Such a process requires a large time investment and a good deal of expert knowledge. This chapter presents an automatic proximity-based algorithm for sampling, culling, clustering diffusion tensor imaging integral curves into anatomically relevant bundles. I will demonstrate the accuracy of the method to form bundles of curves that are anatomically consistent when guided and evaluated by an expert rater. I will then discuss an automated labeling algorithm that models the expert rater's decisions on labeling anatomical structures and evaluate the results of this algorithm. The method potentially provides an efficient and reliable alternative to manual selection of DTI fiber tracts. The method may advance the ability of neuroscientists to test hypotheses about the functional correlates of specific white matter pathways.

5.1 Introduction

Diffusion tensor imaging (DTI) provides 3D information about cerebral white matter anatomy based on the Brownian motion of water molecules in tissue. The diffusion data for each image voxel is summarized as a tensor which is then decomposed to its three eigenvectors and eigenvalues which represent, respectively, the principle axes of diffusion and the rate of diffusion along them. The principle diffusion direction is known to correlate with neural fiber pathways in coherent white matter (Scollan et al., 1998). The eigenvectors and eigenvalues can then be used to produce *in vivo* 3D models that correspond closely to known cerebral white matter fiber structure – a technique called tractography (Basser et al., 2000) (or fiber tracking (Mori and van Zijl, 2002)). Tractography models can then be interactively segmented to isolate white matter tracts of interest (TOIs) on the basis of anatomical knowledge (Akers et al., 2004). Scalar metrics such as average anisotropy or length (Lee et al., 2006) can then be obtained that provide information about the structural integrity of the TOI. These quantitative tractography methods potentially allow neuroscientists to develop and test hypotheses about the functional correlates of specific white matter tracts in health and disease. However, the process of segmenting TOIs is not trivial and an optimal method for doing so has not been established. One approach is to place within the entire tractography model, 3D voxels-of-interest (VOIs) at various points along the presumed path of the desired TOI and then have the computer program display only the fibers that connect intersect the VOIs. This approach has certain advantages in the hands of an experienced and anatomically knowledgeable operator. Examples include the abilities to resolve quickly anatomical ambiguity as might occur in areas where two different white matter tracts are in close spatial alignment, and to make expert decisions about the veracity of anatomically implausible fibers that are sometimes generated in error by the tractography generation algorithm. The VOI approach also has a number of limitations. For example, it assumes considerable operator knowledge of white matter anatomy, it is prone to rater error from

misidentification of tracts or improper decisions about whether to include anatomically ambiguous fibers in a specific tract, and it is susceptible to experimenter bias.

Automated clustering may provide a means of overcoming some of these limitations.

Clustering methods are used extensively in pattern recognition, machine learning, computer vision, data mining, and in many other related fields. In a rather informal sense, the goal of clustering is to partition a given set of data into homogeneous groups. A number of approaches have been attempted to cluster DTI integral curves and the results differ across methods depending on the similarity measure (i.e., a measure of the variable used to partition the groups), the specifics of the clustering algorithm, and the parameters employed by the method. Moreover, the results may be affected by image artifacts that interact with the clustering method. I argue that the accuracy of results of a clustering algorithm will be improved if the choices made in generating the integral curves and selecting clustering algorithms are anatomically motivated. We assume that in cerebral white matter, there exists a natural, anatomically-determined proximity threshold \mathbb{T} , whereby any neural fiber A in a white matter fiber bundle has at least one neighboring neural fiber B in the same bundle and the proximity between A and B is less than \mathbb{T} , while the proximity between any two fibers in different fiber bundles is larger than \mathbb{T} . In reality, noise, motion, and other imaging artifacts can often perturb the proximity between these integral curves. Without an accurate model for these artifacts and with no “gold standard” for determining the accuracy of tractographic representations of white matter anatomy, it can be difficult to find an optimal value for \mathbb{T} . To address this problem, we designed a system for interactively selecting the proximity threshold for clustering.

In section 5.2, we review prior work in this area. In section 5.3, we describe the data acquisition and preparation methods. In section 5.4, we describe our sampling method, constraints and culling methods for generating a set of DTI integral curves. In section 5.5, we use a clustering algorithm and a proximity threshold to group these

integral curves into anatomically-related bundles. To validate our clustering method, in section 5.6 we interactively select proximity thresholds to achieve visually optimal clustering in models from 4 data sets, assigning confidence ratings for each of 12 target white matter fiber bundles of varying calibers and type (i.e., commissural, association, projection) in each data set. These interactively selected proximity thresholds and confidence ratings were used to guide the development of a fiber-bundle template for automated clustering and labeling a subset of the 12 selected fiber bundles on two additional subjects.

Preliminary results demonstrate that when our proximity-based clustering algorithm was guided (i.e., interactively thresholded to obtain optimal clustering) and evaluated by an expert rater, it produced fiber bundles that were qualitatively consistent with known white matter anatomy. However, bundles with smaller caliber or with low directional coherence are identified with lower confidence. We then used the results of the expert rater study to build a fiber bundle template and to develop an automated program for labeling the bundles that models the rater’s manual thresholding. When we applied the template and automatic labeling program to two new subjects (i.e., not used in developing the expert rater study), the algorithm automatically and rapidly produced labeled fiber bundles that are consistent with the target anatomy. These results demonstrate that it is possible to perform automated clustering and labeling of DTI integral curves into bundles that are qualitatively consistent with white matter anatomy and that successfully model expert rater’s decisions. We believe this work can be used to perform fast, unbiased, and accurate identification of neural fiber bundles and has a wide range of potential applications in white matter disease and abnormality studies.

5.2 Related Work

A number of methods have been developed for classifying clustering and labeling these DTI integral curves. Ding *et al.* defined a similarity measure between integral curves on their corresponding segments and grouped together similar integral curves with seeding points proximal to the original integral curve seeding point (Ding et al., 2003); Corouge *et al.* (Corouge et al., 2004) proposed an algorithm that propagates a cluster to neighboring fibers using three distance measures derived from paired points on the streamlines. The distance measures were the minimal distance between two paired points on the two curves, the mean proximity between all paired points, and the Hausdorff distance. Brun *et al.* used a spectral clustering method to cluster the integral curves (Brun et al., 2004). O’Donnell *et al.* extended this approach to a set of integral curves in data from a sample of brains (O’Donnell and Westin, 2005). Maddah *et al.* used hand-selected ROIs in white matter to construct a bundle template in which curves from a new subject can be registered (Maddah et al., 2005). Moberts *et al.* evaluated different fiber clustering methods by constructing a “gold standard” and used a Rand index and its derivatives for comparing the results of different similarity measures and clustering algorithms to this gold standard (Moberts et al., 2005).

The current study extends this previous work by implementing a method that generates DTI integral curves that incorporate anatomically motivated constraints (see section 5.4 below) and then choosing a clustering algorithm based on our anatomical assumptions about the proximity of neural fibers within and between bundles. We then developed a template-driven, automated labeling method that reliably models an expert rater’s decisions.

5.3 Data Acquisition and Preparation

The Siemens MDDW protocol was used to collect three co-registered sagittal double spin-echo, echo-planar diffusion-weighted volumes of the entire brain in 5 healthy adults (2 males, 3 females; mean (\pm standard deviation) age = 56.8 ± 10.23). All subjects provided written informed consent to participate in a DTI research project approved by the Institutional Review Board at Butler Hospital in Providence, RI.

The volumes were spatially offset in the slice direction by $0.0mm$, $1.7mm$ and $3.4mm$. Parameters for each acquisition were as follows: 5mm thick slices, 0.1mm inter-slice spacing, 30 slices per acquisition, $matrix = 128 \times 128$, $FOV = 21.7cm \times 21.7cm$, $TR = 7200$, $TE = 156$, no partial echoes, $NEX = 3$. Diffusion encoding gradients ($b = 0, 1000s/mm^2$) were applied in 12 non-collinear directions. Total time for the three acquisitions was slightly less than 15 minutes. A vacuum pillow was used to limit subject motion. The three acquisitions were interleaved to achieve true $1.7mm^3$ resolution images and then up-sampled (equivalent to zero-filling) to $0.85mm^3$ isotropic voxels for analysis.

Diffusion tensors were calculated with a non-linear sequential quadratic programming (SQP) method (Ahrens et al., 1998).

5.4 Generating DTI Integral Curves

As discussed in Chapter 3, the integral curves that comprise our tractography models were generated by solving the following equation:

$$p(t) = \int_0^t \vec{v}(p(s)) ds \quad (5.1)$$

where $p(t)$ is the generated streamline and \vec{v} corresponds to the vector field generated from the major eigenvector \vec{e}_1 of the diffusion tensor D . $p(0)$ is set to the initial point of the integral curve, often called the seed point.

We then apply three dimensional coloring and rendering of the resultant curves

and superimpose them on background anatomical images to provide context. The curves directly represent the diffusion information, not the underlying anatomy and therefore, the correspondence between the curves and the true underlying white matter is imperfect. Mismatches are most likely to occur in regions where the principal direction of diffusion is ambiguous. Examples include voxels with high planar anisotropy, (which is thought to be an indicator of reduced directional coherence in the true fiber structure), and voxels contaminated by imaging artifact such as the partial volume effect, which may occur at the boundaries between white matter and grey matter, cerebrospinal fluid, or lesions. Curves passing through such ambiguous regions may end up following an anatomically incorrect path. Such curves could mistakenly merge with other nearby fiber bundles, in which case they might be missed on visual inspection, or the error may be more dramatic producing an anatomically implausible curve. We call an integral curve that does not represent the shape and location of a neural fiber a “spurious curve”.

To limit the number of the spurious curves and maximize the anatomical correctness of our integral curves, we set thresholds on the minimum linear anisotropy value on each point of the curve, on the minimum average linear anisotropy of the entire curve, on the minimum value of signal-to-noise ratio (derived from the non-diffusion-weighted image, $b = 0$) and on the minimum length of the integral curve. We also set a constraint that an integral curve should project into the gray matter. We accomplish this by segmenting the brain into white matter, gray matter and CSF compartments using the FAST (Zhang et al., 2001b) segmentation tool. The segmentation is performed using the non-diffusion-weighted image as well as scalar maps of the trace of the diffusion tensor and fractional anisotropy. We then discard the integral curves whose ending points do not project into the gray matter.

The measure of the spatial proximity between these integral curves is of fundamental importance in both integral curve sampling and clustering and is described here.

We first distinguish between a proximity measure and a metric (also called a distance). A metric $d(x, y)$ must be nonnegative $d(x, y) \geq 0$, reflexive $d(x, x) = 0$, symmetric $d(x, y) = d(y, x)$ and it must obey the triangle inequality $d(x, y) + d(y, z) \geq d(x, z)$. A proximity measure is required to satisfy all of the above conditions except for the triangle inequality. It would be incorrect to call a proximity measure that does not satisfy the triangle inequality a distance. However, a proximity measure does represent the similarity between two data elements and thus can be used in clustering methods.

There are a number of proximity measures between two curves, which can be roughly categorized as the Euclidean distance between two selected points on each curve, such as the closest point measure and the Hausdorff distance (Corouge et al., 2004), the Fréchet distance (Alt and Godau, 1995); or the mean Euclidean distance along the run lengths of the curves. Examples of the latter could be the mean distance of closest distances defined by (Corouge et al., 2004), the mean thresholded closest distances defined by (Zhang et al., 2003); or the distance between two Euclidean space embeddings of the curves, such as the one used in (Brun et al., 2004). Of all these proximity measures, only Hausdorff distance, Fréchet distance and the distance between two Euclidean space curve embeddings are metrics.

The integral curves in the cerebral white matter can be quite long. An ideal proximity measure should consider the matching on the whole run lengths of the two curves. Measuring proximity by the distance between two selected points on two curves ignores all but these two points on those two curves (see Fig. 5.1(a)). Therefore it is not surprising that the mean of closest distances is found to give better clustering results than closest points and Hausdorff distance (Moberts et al., 2005). The mean of the closest distances measure is defined between curves A and B as:

$$d_{MC}(A, B) = \frac{d_m(A, B) + d_m(B, A)}{2},$$

$$\text{where } d_m(A, B) = \text{mean}_{a \in A} \min_{b \in B} \| a - b \|, \quad (5.2)$$

d_m itself is not a proximity measure since it is not symmetric. We can also define

the shorter mean distance of closest distances $d_{SC}(A, B) = \min(d_m(A, B), d_m(B, A))$ and longer mean distance of closest distances $d_{LC}(A, B) = \max(d_m(A, B), d_m(B, A))$ based on d_m .

When two curves are on the boundary of two branching fiber bundles, they might run very closely together for a long course and then diverge abruptly for a relatively short distance. Although two such fibers should be grouped into different clusters, the mean closest distances will be low due to their long overlapping course and they may end up being considered to be part of the same cluster. We address the problem by using a threshold for the mean distance of closest distances (see Fig. 5.1(b)).

$$d_{Mt}(A, B, t) = \frac{d_t(A, B, t) + d_t(B, A, t)}{2},$$

where $d_t(A, B, t) = \text{mean}_{a \in A, \min_{b \in B} \|a-b\| > t} \min_{b \in B} \|a - b\|,$ (5.3)

Similarly, we can define the shorter mean distance of thresholded closest distances as $d_{St}(A, B, t) = \min(d_t(A, B, t), d_t(B, A, t))$, and the longer mean distance of thresholded closest distances as $d_{Lt}(A, B, t) = \max(d_t(A, B, t), d_t(B, A, t))$.

We also use a proximity measure in our integral curve sampling scheme. One widely used sampling scheme for generating integral curves in the whole brain is to sample on a regular grid in the data volume. This sampling scheme biases the results such that long, thin bundles contain more curves in a unit cross-sectional area, while shorter fat bundles tend to fewer curves because they have fewer seed points. This is not desirable because there is no physical evidence to support this relationship between the fiber density and the shape of a bundle. To avoid this bias, we sample curves using a combination of jittered dense regular grid seeding and culling. We first generate seed points on a dense regular grid in the data volume and then jitter their locations. This produces a dense set of seed points that ensures that there is no place in the image volume that is under-sampled, thereby avoiding the bias mentioned above. During the integral curve generation, we cull the shorter member of any two pairs of curves that are too close to each other using d_{St} as the proximity measure. We choose d_{St} because we assume that a short curve staying close to a long

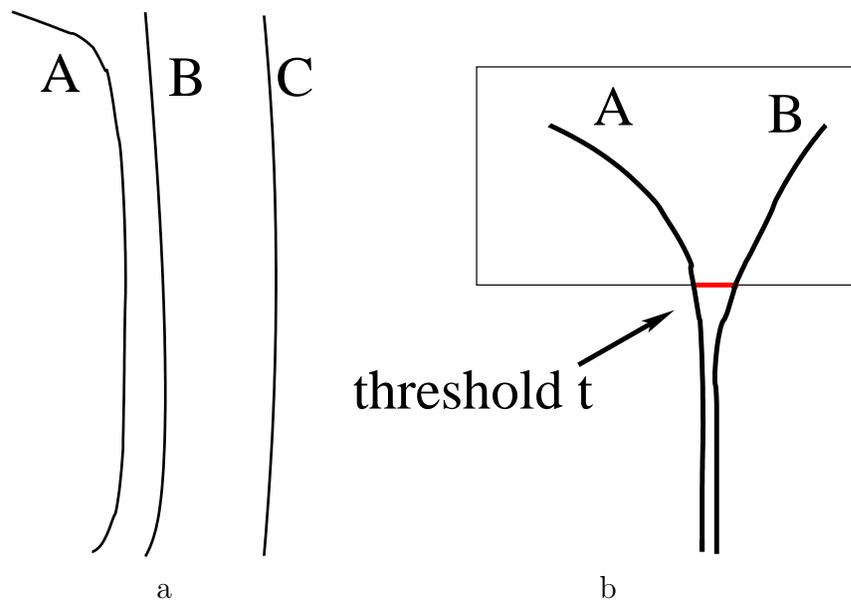


Figure 5.1: (a) shows the shortcoming of defining the proximity between two curves as the distance between two selected points. The Hausdorff and Fréchet distances of A and B are close to those of B and C, even if B is considered more similar to A than C in certain contexts. (b) shows the motivation behind the mean thresholded closest distances. A and B are considered different if they branch for a portion of their lengths. Without the threshold, the mean closest distances between A and B is low if they stay close for a large part of their lengths.

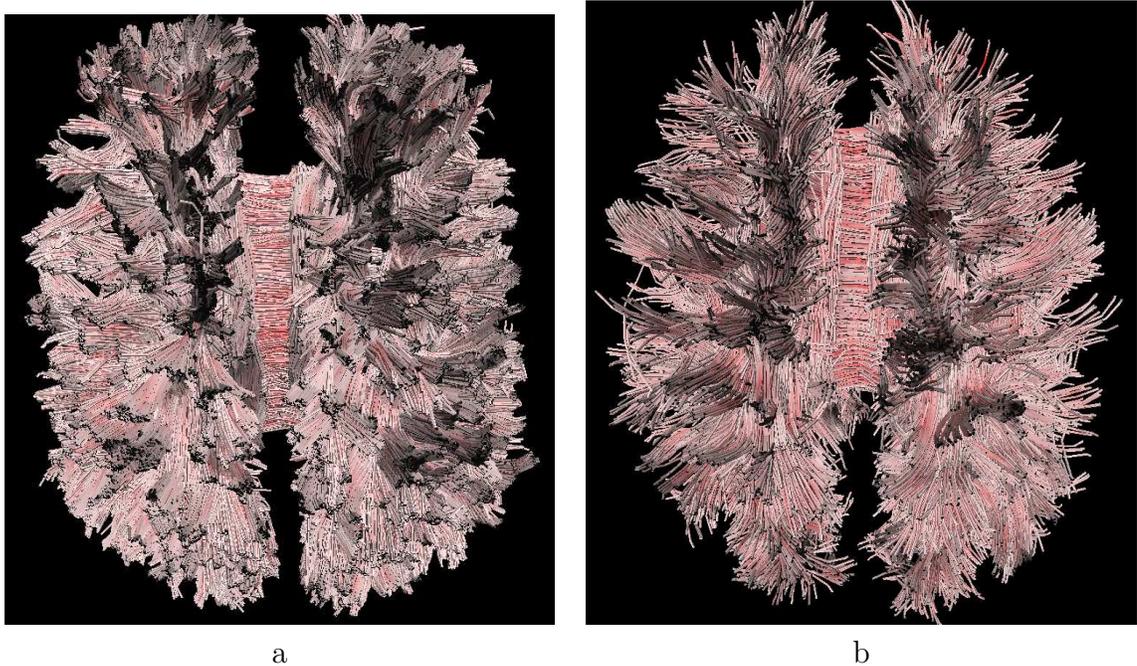


Figure 5.2: Integral curve models generated from a $128 \times 128 \times 90$ data set with a voxel size of $1.7\text{mm} \times 1.7\text{mm} \times 1.7\text{mm}$. (a) Shows 438,042 integral curves generated from a jittered regular grid of $256 \times 256 \times 180$, with minimum length set to 13mm ; minimum linear anisotropy set to 0.15 and minimum average linear anisotropy set to 0.1. (b) Shows the model generated from the same data set and the same jittered regular grid sampling, and with the added culling and gray matter projection constraints. The culling threshold on d_{st} is set to 1.0mm with $t = 0.5\text{mm}$. The model in (b) has 6,113 integral curves. Color is mapped to the linear anisotropy value. Redder means higher linear anisotropy.

curve in its entire length is likely the result of partial volume effect (e.g., encountered a tissue boundary) that prohibits the curve from running its full length, and thus should be considered “close” to its longer neighbor and culled from the model. In the resulting set of integral curves, no two curves are closer than the threshold on d_{st} . This minimum-distance spacing is a property of the Poisson-disk distribution.

Fig. 5.2 shows the integral curves generated on a $128 \times 128 \times 90$ data set with a voxel size of $1.7\text{mm} \times 1.7\text{mm} \times 1.7\text{mm}$. Fig. 5.2(a) shows 438,042 integral curves generated from a jittered regular grid of $256 \times 256 \times 180$, with minimum length set to 13mm ; minimum linear anisotropy set to 0.15 and minimum average linear anisotropy

set to 0.1. Fig. 5.2(b) shows the model generated from the same data set with culling and gray matter projection constraints. The culling threshold on d_{St} is set to $1.0mm$ with $t = 0.5mm$. This model has 6,113 integral curves.

5.5 Single Linkage Algorithm

For clustering the integral curves, we chose d_{Lt} as the proximity measure, used agglomerative hierarchical clustering methods (Duda et al., 2000) and defined the minimum proximity value (i.e., closest) between any two curves from two clusters as the distance between these clusters. This algorithm is also called the single-linkage algorithm (Duda et al., 2000) because any two points with a single path of links between them are clustered together.

The rationale for using the single linkage algorithm to achieve optimal clustering is rooted in our anatomical assumption that in cerebral white matter, there exist an implicit physical proximity measure between neural fibers within a cluster and between clusters. Successful modelling of these implicit physical proximity measures should provide a means for parsing individual curves into bundles that accurately reflect the true state of the underlying white matter. Moreover, the method should be as accurate or potentially superior to other optimization-based clustering methods such as K-means and spectral clustering in that the single linkage algorithm ensures that the physically-based proximity threshold is always satisfied. Optimization-based clustering often involves trade-offs among clusters so that although the goal function is optimized, the proximity between a specific pair of clusters may not satisfy a physically based threshold. Our method circumvents this trade-off. However, in reality, noise, motion and partial volume effects limit our ability to meet our anatomical assumptions. Therefore, our method incorporates an automatic adjustment of the proximity threshold to maximize the accuracy of the clusters in the face of these artifacts (see section 5.6.3).

We apply the single linkage algorithm to our curves using a proximity measure d_{Lt} . The algorithm follows the three general broad steps:

1. Given a set of n singleton clusters.
2. Merge the two nearest clusters
3. Repeat step 2 until the specified number of clusters are generated.

We choose d_{Lt} as the proximity measure between two integral curves. We use the longer one of the mean closest distances to make sure any difference between two curves are captured. Note that in section 5.4, we use d_{St} to get rid of the short curves that are likely to result from the partial volume effect. Here, we assume all those short, spurious curves have already been removed from the tractography in the model and therefore choose d_{Lt} for clustering.

If the maximum distance between any two curves from two clusters is used to measure the proximity between clusters, the algorithm is called the farthest-neighbor clustering algorithm, or “maximum algorithm”. The farthest-neighbor algorithm discourages the growth of elongated clusters. In white matter, some fiber bundles spread out over a thin wide strip, such as the corpus callosum or tapetum and the farthest algorithm might fail to cluster such bundles. The single-linkage algorithm facilitates clustering of such bundles.

The number of clusters rendered in the model is inversely related to the proximity threshold. That is, the greater the distance selected for identifying a single cluster, the fewer the number of clusters. When the nearest-neighbor threshold is set to a high distance (i.e., low proximity) there will be fewer but larger clusters; when the threshold is set to a low distance (i.e., high proximity) the number of clusters increases.

Fig. 5.3 shows a clustering result of the model in Fig. 5.2(b) using single linkage algorithm and d_{Lt} proximity measure with proximity threshold set to $3.5mm$.

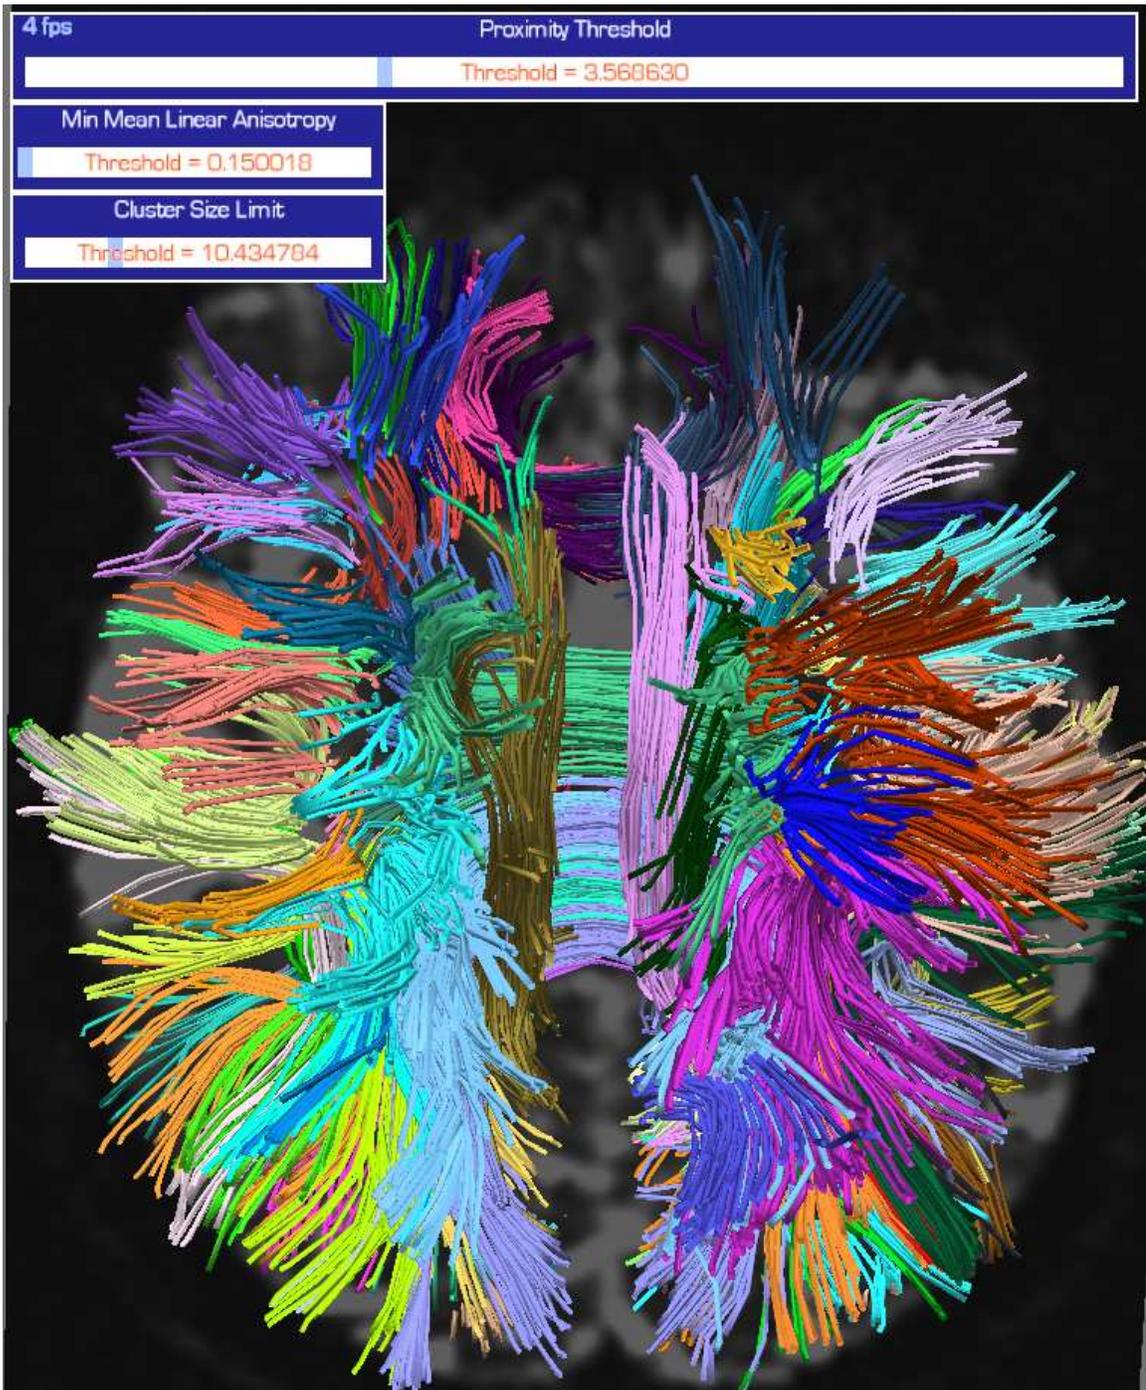


Figure 5.3: The clustering result of the model in 5.2(b) using single linkage algorithm and d_{Lt} as proximity measure with a threshold of $3.5mm$. The top left corner of the interactive interface contains three sliders for adjusting the proximity threshold, minimum mean fractional anisotropy values and minimum number of curves in a cluster.

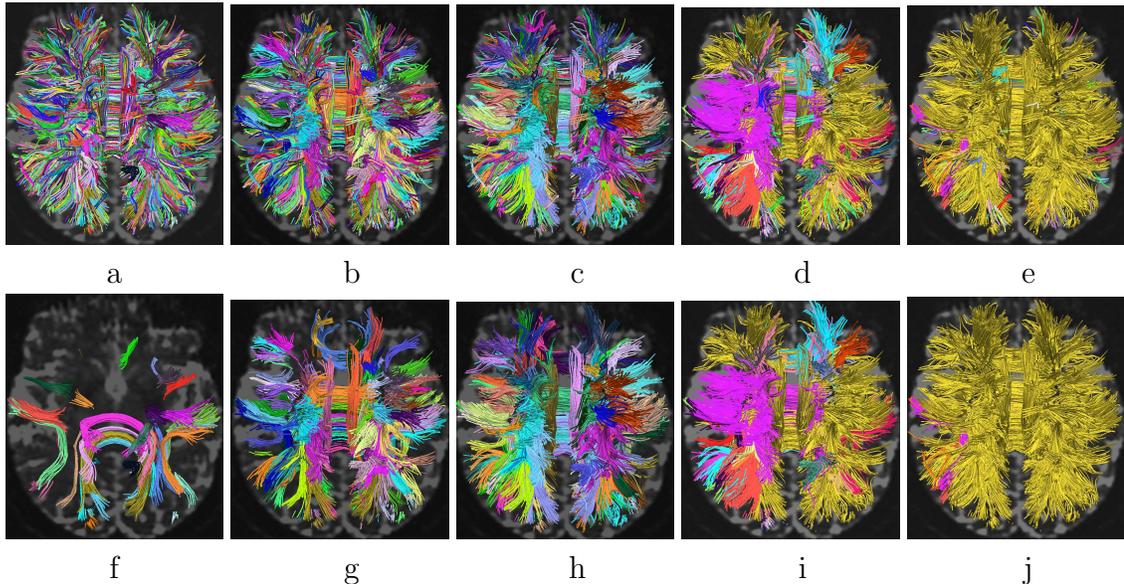


Figure 5.4: Snapshots of exploring the clustering models shown in Fig. 5.3. a,b,c,d and e show the clustering results with proximity threshold set to $1.5mm$, $2.5mm$, $3.5mm$, $4.5mm$ and $5.5mm$ respectively. f,g,h,i and j show the same model/threshold with minimum cluster size threshold set to 10.

5.6 Exploration, Evaluation and Automatic Labeling

5.6.1 Interactive Exploration of Clustering Models

We built an interactive interface for exploring the clustering results and searching for an optimal representation of white matter fiber bundles.

The interface is shown in Fig. 5.3. The model in the picture is generated with methods described in section 5.4. Single linkage clustering is precomputed on the model with the proximity threshold set from $0mm$ to $10mm$ in $0.1mm$ intervals. In the top left corner, three sliders are used to change the values of the proximity threshold, minimum mean linear anisotropy along an integral curve and the cluster size threshold. The interactive system runs at about 10 frames-per-second with a $30MB$ VRML model on a desktop PC with a $2GHz$ AMD CPU and $1GB$ memory.

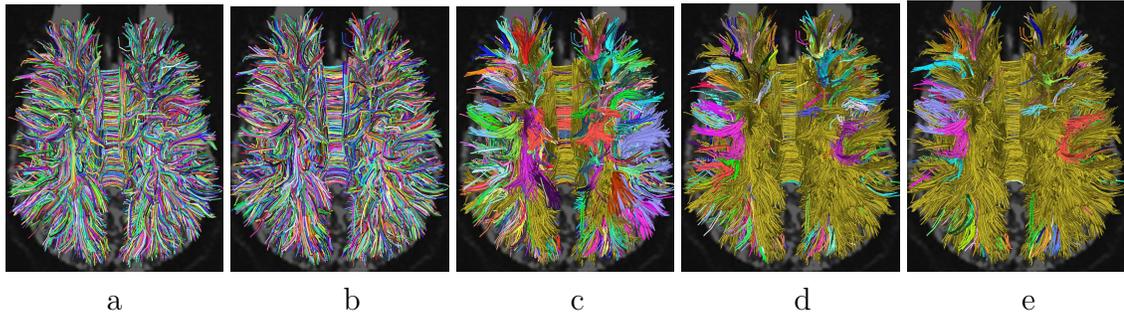


Figure 5.5: The interactive interface can be used to quickly evaluate a new clustering algorithm or proximity measure. a,b,c,d show the same model used in Fig. 5.4 with the same single linkage algorithm, but different proximity measure (d_{St} instead of d_{Lt}). a, b, c, d and e show the clustering results with the proximity threshold set to $0.5mm$, $1.4mm$, $1.5mm$, $2.1mm$ and $2.5mm$ respectively. The abrupt change of cluster structures around proximity threshold $1.5mm$ suggests that the d_{St} is more prone to this kind of jump than d_{Lt} .

Fig. 5.4 shows snapshots of clustering results of the model shown in Fig. 5.3 with different proximity thresholds and cluster size thresholds.

This system is convenient for qualitatively evaluating new clustering algorithms or proximity measures. This capability is shown again in Fig. 5.5, where d_{St} instead of d_{Lt} is chosen as the proximity measure. The four pictures show the resulting clusters with proximity threshold set to $0.5mm$, $1.4mm$, $1.5mm$, $1.9mm$. It is clear that using d_{ST} as the similarity measure results in an abrupt change in cluster structures around the proximity threshold of 1.5. This is because d_{ST} deems short and long curves similar provided that the short curve is close to a portion of the long curve, usually resulting in big clusters with curves of different lengths.

5.6.2 Evaluation

To evaluate our clustering method, we interactively selected a proximity threshold to achieve visually optimal clustering in models from 3 subjects S1, S2, S3 and S4. Subject S2 and S4 are the same individual scanned on different occasions. An expert rater determined the threshold that produced a set of color-labeled bundles that maximally

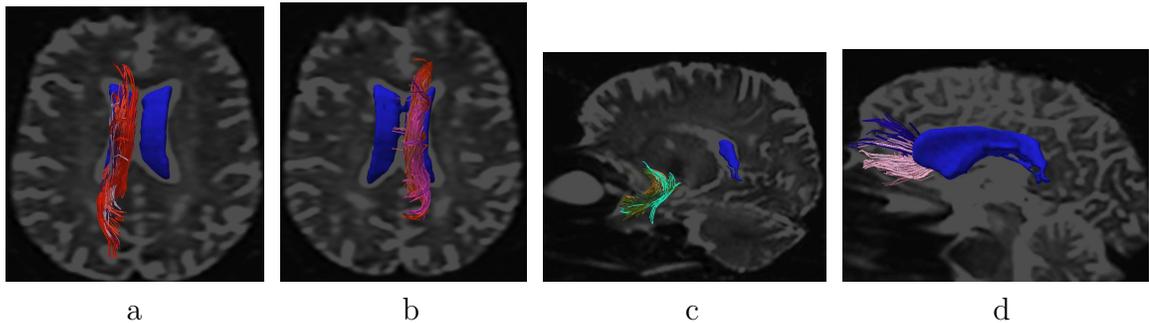


Figure 5.6: Four of the fiber bundles in the template. (a) left cingulum bundle, (b) right cingulum bundle, (c) left uncinate, (d) forceps minor. Blue surface represents the ventricles and a slice of T_2 image is shown as context. Colors are randomly generated for each bundle of each subject.

facilitated identification of 12 pre-specified target white-matter tracts of varying calibers and type (i.e., commissural, association, projection) in both hemispheres (where applicable) in each subject. After finding an optimal threshold, the rater used a 4-point ordinal Likert-type scale to assign confidence ratings for identifying each of the 12 target white matter tracts in each hemisphere in the thresholded bundled models. Points on the scale were u=unlikely located or not located, p=possibly located, l=likely located, v=very likely located.

The results are shown in Table 5.1. Each row represents one of the target white-matter tracts. Each rating pair within a row shows the ratings for that feature in both hemispheres of the brain. A comparison across the subjects shows that long, thick, coherent white matter tracts like the uncinate fasciculus, cingulum bundles, and superior and middle cerebellar peduncles are identified with high confidence in the bundled model. In contrast, shorter, thinner white matter tracts like the anterior commissure are almost completely missing from these optimally thresholded models. The difficulty in identifying these smaller structures in the clustered models may be due to the combined effects of their small caliber and the limited resolution of the images. Other tracts were identified with intermediate confidence. For example, some tracts, like the frontal forceps, are rated with reasonably good confidence in

most subjects but are missing or rated somewhat lower in one subject, suggesting possibly an error in the model due to image artifacts. Some tracts, such as the superior longitudinal fasciculus, inferior longitudinal fasciculus, arcuate fasciculus and inferior cerebellar peduncle, received inconsistent ratings. It is important to note that some clusters observed by the rater but not included among our 12 target ones, were thought to be possibly spurious or at least anatomically questionable.

WMF	S1	S2	S3	S4
fminor	v	v	l	v
fmajor	v	v	v	v
slf	l,v	l,v	l,v	p,p
ilf	l,p	l,v	p,l	l,l
af	l,v	u,v	p,v	p,p
uf	v,v	v,v	v,v	v,v
ac	p	u	l	u
cb	v,v	v,v	v,v	v,v
scp	u,u	p,p	l,l	p,p
mcp	l,l	v,v	v,v	v,v
icp	p,p	u,u	u,p	p,p
cst	v,v	v,v	v,v	v,v

Table 5.1: Confidence ratings for the 12 white matter features in the clustering models of four data sets, S1-S4 . WMF – white matter fiber tract; ff – frontal forceps; slf – superior longitudinal fasciculus; ilf – inferior longitudinal fasciculus; af – arcuate fasciculus; uf – uncinate fasciculus; ac – anterior commissure; cb – cingulum bundle; scp – superior cerebellar peduncle; mcp – middle cerebellar peduncle; icp – inferior cerebellar peduncle; cst – cortico–spinal tract; v indicating very likely located, l indicating likely located, p indicating possibly located, and u indicating unlocated. A rating pair shows the ratings for the same tract in each cerebral hemisphere.

These results demonstrate that the proximity-based clustering method is capable of bundling and labeling fibers that in the view of an expert rater, have high correspondence with known white matter anatomy, at least for large distinct tracts. The method produces less consistent results for some smaller tracts that may be either inherently more difficult to cluster accurately and consistently and, because of their size or location, may be more vulnerable to image artifacts. Despite these limitations,

the results show that our method has promise as a means for interactive clustering of DTI integral curves into anatomically-relevant clusters. Improvement in image quality, motion reduction, and in the algorithm itself may provide better results in the future.

5.6.3 Automatic Fiber Bundle Labeling

One potentially very useful goal is to implement a method that performs automated clustering and labeling. We explored this potential by constructing a bundle template based on the interactive results presented in the previous section, and then attempting to match integral curve bundles from a new subject to this template.

We built a fiber bundle template by matching the clustering results on 4 data sets (S1, S2, S3, and S4) using the expert-defined optimal proximity threshold. First, we registered the non-diffusion-weighted image from every data set to one particular data set chosen at random using FLIRT (Jenkinson and Smith, 2001). The registration was constrained to translation, rotation and scaling operations only. We then used the transformation matrices to register all the cluster models. For each integral curve cluster, the centroids of the starting points C_s , middle points C_m , and end points C_e were calculated. Integral curve clusters from the two subjects were then aligned and grouped according to the sum of the distances between these centroids

$$d_M(A, B) = \| C_s(A) - C_s(B) \| + \| C_m(A) - C_m(B) \| + \| C_e(A) - C_e(B) \| \quad (5.4)$$

To be matched together, clusters are required to be mutually closest in the feature space as well as nearer than $40mm$.

These matched clusters are then used to construct a color-labeled fiber bundle template by picking six fiber bundle structures that are matched well across all four data sets.

A cluster A from a new subject can be matched to one of the fiber bundles B in

the template by a matching score of

$$\mathbb{M}(A, B) = \min_{i \in \mathbb{S}} \frac{d_M(A, B_i)}{|A|}, \quad (5.5)$$

where $|A|$ is the number of curves in A . \mathbb{S} is the set of subjects we used for building the template. B_i is the cluster in subject i that is labeled as B . \mathbb{M} is designed to favor big clusters in A that are close to B . \mathbb{M} is set to a large number if there is a non-match.

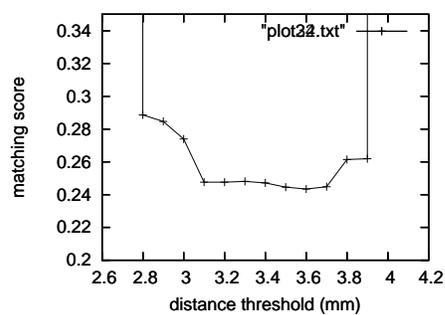
The algorithm then searches for a proximity threshold on the new subject that minimizes the sum of the matching scores for all the clusters from the new subject that are matched to one of the fiber bundles in the template. The matching scores for two new subjects S5 and S6 are shown in Fig. 5.7. The search is easy since the range of clustering proximity thresholds for a good match is rather narrow.

We then match the automatically-thresholded clustering result to the fiber bundle template. Fig. 5.8 shows the matched bundles for S5 and S6. We rate these bundles as very likely or likely to be labeled correctly.

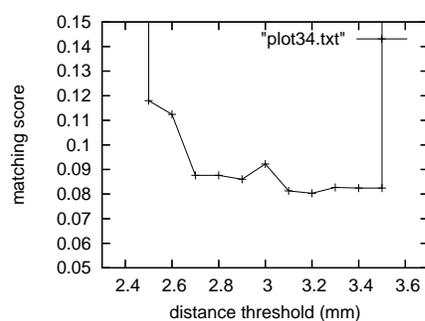
5.7 Discussion and Conclusion

White matter abnormalities caused by acquired or inherited disease or from congenital conditions can impact specific white matter pathways differentially. Automatic classification and labeling of white matter fiber bundles in DTI tractography models potentially provides an efficient method for rapidly and reliably identifying specific bundles across individual subjects. Once clustered, white matter bundles can be measured using any number of scalar metrics that lend themselves to statistical analysis. Automated clustering therefore provides neuroscientists with an important tool for generating and testing hypotheses about the functional correlates of specific white matter tracts in both normal and abnormal brains.

The results of this proof-of-concept study demonstrate that anatomically-motivated methods for integral curve generation and clustering provide can reliably identify large

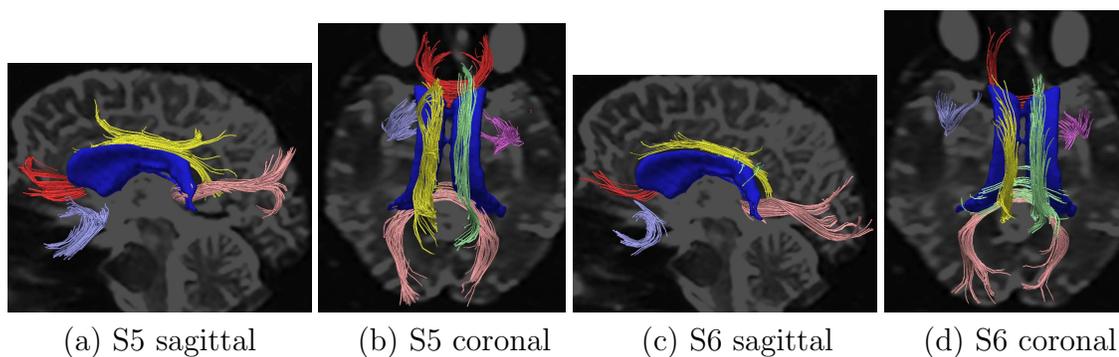


(a) S5



(b) S6

Figure 5.7: The sum of matching scores for all the fiber bundles in the template. (a) shows the result of S4. Proximity threshold is 3.6mm for the optimal match, while the expert-selected threshold is 3.9mm . (b) shows the result of S5. Proximity threshold is 3.2mm for the optimal match, while the expert-selected threshold is 3.5mm .



(a) S5 sagittal

(b) S5 coronal

(c) S6 sagittal

(d) S6 coronal

Figure 5.8: The matched bundles for subject 4 and 5. Color is fixed on each label. Blue surface represents the ventricles. Note the correspondence between the same bundles across the two subjects.

fibers bundles that bear close qualitative correspondence with known white matter anatomy across subjects. The results for these large tracts are good both when the algorithm is guided with interactive input from an expert rater and when it is implemented automatically.

At this point, however, our method works best for large, distinct white matter tracts such as the cingulum bundles and uncinate fasciculus; and works less well, and in some cases quite poorly, for smaller or less anatomically distinct tracts such as the inferior longitudinal fasciculus or the anterior commissure. We believe that the method will likely produce more consistent results for a wide set of target tracts with further refinements and with improvements in image acquisition, post-processing (e.g., motion correction), and tractography algorithms. Improvement can be made in building a high resolution, verified fiber-bundle template, reducing noise, motion and other artifacts in image collection, and incorporating more anatomical knowledge into integral-curve generation, clustering and labeling.

Chapter 6

Quantitative Study of the Streamtubes Model

In terms of revealing information from high-dimensional data sets, qualitative and quantitative methods are often complementary and both important. So far we have mostly discussed qualitative methods. For example, Chapter 4 presents several visualization schemes for qualitative data exploration. Chapter 5 expedites this exploration by automatically identifying major fiber bundle structures. On the other hand, quantitative methods are needed for fast and unbiased assessment of the models and the data for both developing reliable biomarkers and improving the accuracy of the models.

In this chapter, we first design a number of metrics on streamtubes for assessing white matter integrity. Then we build a forward model from the streamtubes model to diffusion MRI signals and use this model to optimize the streamtubes model.

6.1 Quantitative Tract-of-Interest Metrics for White Matter Integrity

In (Lee et al., 2005), we proposed several quantitative tract-of-interest metrics for assessing white matter integrity. That work is excerpted here.

The majority of studies using DTI to assess white matter microstructure in clinical samples have been based on two-dimensional grayscale maps of scalar values such as mean diffusivity (MD), a measure of the magnitude of diffusion in each image voxel, and fractional anisotropy (FA), a measure of the extent to which that diffusion is directionally restricted. Generally, these basic scalar measures are derived from the eigenvalues of the multi-valued tensor data and do not incorporate eigenvector information. The scalar values in each image voxel are then mapped to two-dimensional grayscale images. An exception is the use of eigenvector information to produce two-dimensional FA maps in which fiber orientation is mapped to color (e.g., (Pajevic and Pierpaoli, 1999); (Wakana et al., 2004)). However, our work is based on linear anisotropy (c_l), a scalar parameter related to FA and proposed by (Westin et al., 1997b).

Tractography methods complement scalar methods by providing detailed information about the orientation and curvature of white matter pathways as they course through the brain. These methods utilize both the tensor eigenvalues and the eigenvectors to calculate trajectories in the direction of fastest diffusion. The trajectories are then portrayed graphically using curved lines (Xue et al., 1999) or glyphs such as hyperstreamlines, which were initially proposed by (Delmarcelle and Hesselink, 1993) as a means of visualizing other types of second-order tensor fields and then subsequently applied to DT-MRI data by (Zhang et al., 2003). To date, tractography has gained widest acceptance in neuroscience studies exploring white matter connectivity, effects of pathologies on connectivity, and improvement of data acquisition and visualization methods (e.g., (Huang et al., 2005; Wakana et al., 2004)).

Only a few preliminary methodological studies have explored the utility of combining tractography with quantitative scalar measures (i.e., “quantitative tractography”) for clinical research where group comparisons are important. For example, Ciccarelli et al. (Ciccarelli et al., 2003a), studied the reproducibility of tract-“normalized” volume (NV) and FA in three white matter pathways traced by the fast marching tractography (FMT) algorithm (Parker et al., 2002a; Parker et al., 2002b). The results (Ciccarelli et al., 2003a) showed variability in measures of tract volume and fractional anisotropy across different fiber bundles suggesting that fiber organization has an impact on the reproducibility of tractography algorithms. (Ciccarelli et al., 2003b) also examined the extent of intersubject variability in the anterior corpus callosum, optic radiations, and pyramidal tracts. They found that the tractography maps corresponded well to known anatomy and that there was greater intersubject variability at the terminal ends of tracts adjacent to cerebral cortex, but lower variability in the core of tracts and no right-left differences in variability. (Ding et al., 2003) also demonstrated good reproducibility of tractography-based metrics such as curvature, torsion, parallel diffusivity, and perpendicular diffusivity along bundle length. (Huang et al., 2005) have also used quantitative methods for parcellating projections from the corpus callosum to cortical regions. Each of these studies demonstrated the utility of using quantitative tractography.

The present study proposes several new quantitative tractography metrics for quantifying cerebral white matter integrity in whole brain white matter and in specific white matter tracts of interest (TOIs). Some preliminary results are published in (Lee et al., 2006). In this section, we define the new metrics and demonstrate their stability across multiple streamtube models derived from a single dataset for one healthy control. Then, as an initial and preliminary validity test, we compare these metrics between a small cohort of patients with known vascular white matter injury (i.e., vascular cognitive impairment, VCI) and a demographically matched cohort of healthy controls. VCI provides a good model for assessing the clinical and research

utility of quantitative DTI since ischemic white matter injury is associated with a characteristic pattern of increased diffusivity and decreased anisotropy (Jones et al., 1999). These diffusion changes are thought to reflect axonal loss (Beaulieu, 2002) with possible contributions from demyelination (Bihan et al., 2001), gliosis (Larsson et al., 2004), or other pathological processes.

6.1.1 Tract of Interest (TOI) Selection

We implemented a method for interactive TOI selection similar to the volume of interest (VOI) approach of (Akers et al., 2004), which allows users to define box-shaped VOIs and Boolean expressions to select streamtubes of interest. Our approach is similar to a more recent method published by Mori and colleagues (Jiang et al., 2006). For each subject, we quantified the streamtube data for the whole brain and three TOIs: interhemispheric streamtubes passing through the corpus callosum, and the left and right cingulum bundles. TOI selection was performed by an experienced rater (SYL) trained and supervised by two faculty with good knowledge of white matter anatomy (SC and DT). Interhemispheric fibers were defined as all streamtubes passing through the corpus callosum. Fibers that extended inferiorly through a plane defined by the inferior boundary of the splenium and rostrum in mid-sagittal view were removed from this TOI. This conservative approach may have removed anatomically correct streamtubes from the TOI (e.g., the tapetum) but was designed to ensure that anatomically implausible streamtubes were not included in the analysis. The left and right cingulum bundles were identified visually in the models and were segmented by the intersection of voxels placed in the posterior third and anterior third of each bundle using a pre-specified algorithm for voxel placement. Given the potential errors that can occur in DTI fiber tracking (Jiang et al., 2006; Mori and van Zijl, 2002), we adopted a conservative approach to identifying and culling anatomically questionable fibers (e.g., those that seemed possibly spurious based on trajectory). The final TOI results were inspected by two faculty for appropriateness.

6.1.2 Proposed Metrics

We developed the following several new quantitative tractography metrics: 1) total length (TL), 2) total weighted length (TWL), 3) normalized total length (NTL), 4) normalized total weighted length (NTWL), 5) number of streamtubes (NS), and 6) normalized number of streamtubes (NNS). We also developed a seventh metric by dividing TL by NS to yield average TL (ATL). These metrics should provide indirect markers of the structural integrity of white matter TOIs. The metrics are defined conceptually and mathematically as follows:

Total length of a TOI was defined as:

$$TL = \sum_{\forall s \in S} L_s \quad (6.1)$$

where L_s is the length in millimeters of a single streamtube s contained within the set of streamtubes S in the given TOI. This metric should provide a coarse marker of the overall microstructural integrity of a TOI. Total weighted length (TWL), weights TL by the average linear anisotropy of the streamtubes comprising the TOI. Weighting by the average linear anisotropy provides different information about the integrity of the tract. For example, TWL might be sensitive to differences in tract integrity across groups (i.e., differences in the average linear anisotropy of the tract) that are insufficient to impact TL, but nonetheless reflect a disease state. TWL is defined as follows:

$$TWL = \sum_{\forall s \in S} C_s \times L_s \quad (6.2)$$

where C_s is the average linear anisotropy within streamtube s . We would expect TWL to be inversely correlated with TOI injury.

NS quantifies the number of streamtubes in a TOI:

$$NS = \|S\| \quad (6.3)$$

NS may provide useful information about white matter health not captured by the length metrics and may be sensitive to white matter changes, particularly in shorter, denser tracts. NS might provide a compliment to the length metrics. For example, lower TL could occur because of broken fibers due to local minima in linear anisotropy or to loss of fibers. In the former situation, lower TL might be accompanied by higher NS, whereas in the latter situation, both TL and NS might be decreased.

The value of TL, TWL, and NS, however, are likely influenced by brain size, and thus may require further correction. Therefore, we normalized these metrics by approximate intracranial volume. We chose to normalize by intracranial volume instead of brain volume because it likely provides a better index of brain size prior to the impact of age and pathology.

We approximated total intracranial volume by first defining a box enclosing the brain. The box boundaries were determined in sagittal and coronal views of the T2 (i.e., $b=0$) volume and were placed at the most superior, inferior, anterior, posterior, lateral points of the cranial vault. The inferior boundary was the foramen magnum. We then measured the volume of a mathematically generated three-dimensional ellipsoid circumscribed within the box. To account for extra-cranial tissue captured by the ellipsoid, we multiplied the ellipsoid volume for each individual by a correction factor based on normative intracranial volume estimates reported in four previous studies . The correction factor was the ratio of a mean intracranial volume estimate calculated across these four studies to the mean brain box volume in our healthy controls and VCI subjects combined. Given the small sample size, we did not derive gender-specific intracranial volume correction factors. Application of this correction factor brought our intracranial volume estimates in line with previously reported intracranial volume estimates derived with more exacting methods.

We normalized TL, TWL, and NS by dividing our length metrics by the ratio of each participant's approximate intracranial volume to the average intracranial volume of all participants. That is, we computed normalized total length (NTL) and

normalized total weighted length (NTWL) as follows:

$$NTL = \frac{TL}{V/\bar{V}} \quad (6.4)$$

$$NTWL = \frac{TWL}{V/\bar{V}} \quad (6.5)$$

$$NNS = \frac{NS}{V/\bar{V}} \quad (6.6)$$

where V is approximate intracranial volume, and \bar{V} is the mean approximate intracranial volume for the healthy controls and VCI subjects combined.

Note that with our seeding and culling scheme, the streamtubes fill the white matter volume and keep certain distance from each other. Thus the total length of the streamtubes and its weighted derivatives can be regarded as being roughly proportional to the volume of the white matter, prompting us to use V/\bar{V} for normalization.

Lastly, our seventh metric, ATL is a non-normalized measure of the average length of streamtubes.

$$ATL = \frac{TL}{NS} \quad (6.7)$$

ATL may provide slightly different information from TL. For example, ATL could help discriminate between individuals with similar TL values but different white matter integrity. For example, an individual with good white matter integrity reflected in long streamtubes vs. one with poor white matter integrity reflected by many shorter streamtubes. We additionally calculated average values for our other length metrics by dividing each by NS, however they were omitted from analyses partly to reduce the number of variables and because they were highly correlated ($r > .72$) with the metric in the numerator.

6.1.3 Reproducibility Study

Two important factors potentially affect the reliability of our metrics: 1) the degree to which our streamtube generation algorithm yields consistent results; and 2) the degree to which our DTI acquisition parameters yield consistent results across repeated scans where the head position is not identical.

Data Acquisition

We examined the consistency of our metrics across data collected from a single healthy 45-year-old, right-handed control subject scanned on three different occasions — the second and third data sets were collected on the same day with removal from the scanner between acquisitions; the first data set was collected approximately 4 weeks earlier. Data from this control subject was not included in other data analysis in this report.

We used percent differences to evaluate the consistency of our streamtube model in datasets of a single healthy control subject collected at three separate time points. Between the three pairs of datasets (i.e., Time 1 vs. Time 2, and so on) the percent difference between our metrics was $\leq 0.51\%$ for NS , $\leq 1.59\%$ for TL , and $\leq 2.6\%$ for TWL .

Streamtube Generation

Streamtubes are generated using algorithms that place randomly jittered seed points at regular intervals within the imaging grid. This jittering method reduces the likelihood of the output model being overly influenced by the grid. The cost for this small random jitter is that each run of the algorithm on a single dataset produces a slightly different streamtube model, even if the input data and parameters are held constant. These subtle inconsistencies should roughly mirror differences in streamtube models across multiple image acquisitions of the same subject using strict seed point placement. However, these subtle inconsistencies could also translate into additional error

Table 6.1: Results of the reproducibility study (mean \pm SD) and coefficient of variance (COV; means and SD are rounded)

Seed	NS	TL(mm)	TWL(mm)
$1.7mm^3$	7895 ± 11 COV = .1%	108978 ± 844 COV = .8%	26874 ± 256 COV = 1.0%
$0.85mm^3$	11856 ± 28 COV = .2%	138752 ± 823 COV = .8%	33639 ± 245 COV = .9%
$0.64mm^3$	13406 ± 40 COV = .3%	151180 ± 703 COV = .8%	36381 ± 157 COV = .4%

variance in our metrics which, if too great, could mask true group differences (i.e., increased risk of Type 2 error).

To address this concern, we assessed the consistency of our metrics in multiple streamtube models derived from one dataset of a single healthy control subject. Seven streamtube models were produced for each of three different seeding parameters that varied by coarseness (i.e., $1.7mm^3$, $0.85mm^3$, and $0.64mm^3$) for a total of 21 models. Other streamtube parameters were held constant as described previously. We then assessed the percent difference in our metrics across the eight whole brain models produced at each seeding parameter. This method allowed us to determine the impact of both random jitter and different seeding parameters on the consistency of our metrics.

In the whole brain models, coefficients of variance did not exceed 1.0% for either NS, TL, or TWL across seven models of a single brain produced at the three seeding densities of $1.7mm^3$, $0.85mm^3$, and $0.64mm^3$ seeding (see Table 6.1).

6.1.4 Whole Brain Metrics

We used MANCOVA (controlled for age) to test for significant differences in our metrics between the VCI and healthy control groups in whole brain models. We expected that patients with VCI would have lower values on all our metrics than

Table 6.2: Group comparison, whole brain corrected for age

Metric	Healthy Controls	VCI	F(1,15)	p
Whole brain (n=18 for the HC group; n=15 for the VCI group)				
<i>NS</i>	12647 ± 2076	11385 ± 2571	5.42	.027
<i>NNS</i>	12720 ± 2286	11425 ± 2611	5.38	.027
<i>TL(mm)</i>	396453 ± 84619	327379 ± 97501	7.70	.009
<i>ATL(mm)</i>	31 ± 3	28 ± 3	7.70	.007
<i>TWL(mm)</i>	103883 ± 24499	80186 ± 26256	8.04	.008
<i>NTL(mm)</i>	399021 ± 91869	328055 ± 94979	10.20	.003
<i>NTWL(mm)</i>	104475 ± 25946	80482 ± 26256	10.43	.003
Interhemispheric fibers (n=17 for the HC group; n=14 for the VCI group)				
<i>NS</i>	804 ± 151 650 ± 210	9.23	.005	
<i>NNS</i>	801 ± 137	650 ± 214	10.21	.003
<i>TL(mm)</i>	44900 ± 12150	32485 ± 15533	12.50	.001
<i>ATL(mm)</i>	55 ± 6	48 ± 9	14.76	.001
<i>TWL(mm)</i>	15284 ± 4204	10353 ± 5502	14.75	.001
<i>NTL(mm)</i>	44835 ± 11922	32546 ± 15612	13.27	.001
<i>NTWL(mm)</i>	15256 ± 4099	10406 ± 5567	15.36	.001

healthy controls. Our decision to use MANCOVA for this analysis reflected our desire to control for age. Given the results of our Shapiro-Wilks test, we repeated the analyses using non-parametric Mann-Whitney U tests to evaluate the extent to which violations of normality may have influenced our results.

Figure 6.1 shows the 3-dimensional views of whole brain streamtube models for one healthy control (a) and one VCI patient (b).

In the healthy control group, correlations between age and each of the metrics were negative and significant with all $r \geq -0.57$ and all $p \leq 0.1$, except between age and ATL ($r = -0.17$; $p = .51$). In the VCI group, age was not significantly correlated with any of the metrics. Nonetheless, we controlled for age in our main whole brain analysis as discussed above.

A MANCOVA was conducted to evaluate group differences for the seven metrics with age as a covariate. The overall model showed a non-significant trend for age (F

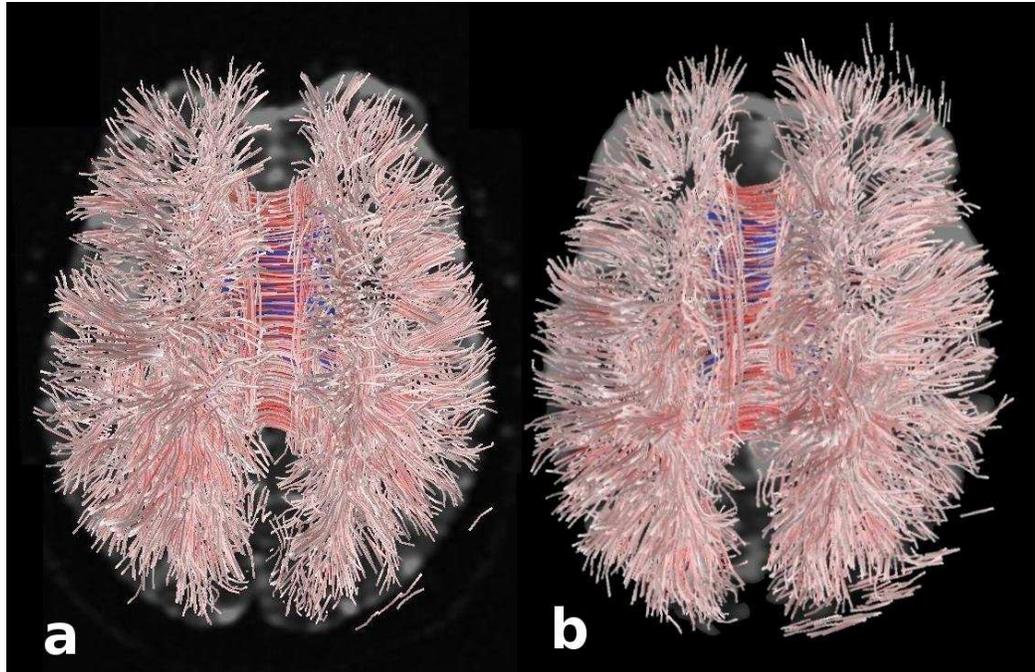


Figure 6.1: The 3D views of whole brain streamtube models for one healthy control (a) and one VCI patient (b).

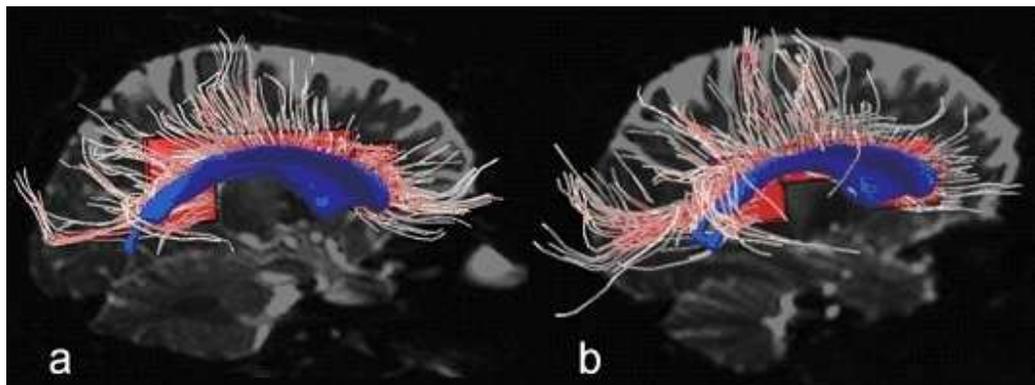


Figure 6.2: The IHF selection in one healthy control (a) and one VCI patient (b).

= 2.24; $p = .07$) and a significant effect of group ($F = 2.93$; $p = .02$). After controlling for age, the VCI group had significantly lower values on all seven quantitative tractography metrics compared to the healthy controls (Table 6.2). Effect sizes across the seven metrics were small ($\eta_p^2 = .15$ to $.23$); observed power ranged from $.62$ to $.89$.

A Pearson bivariate correlation matrix revealed strong and statistically significant intercorrelations among all seven metrics with r values ranging $.523$ for *ATL* vs. *NNS* to $.99$ for *TWL* vs. *NTWL* (all $p < .001$).

6.1.5 Interhemispheric Fibers

Figure 6.2 shows the IHF selection in one healthy control (a) and one VCI patient (b).

The overall MANCOVA model revealed a non-significant trend for age ($F = 2.00$; $p = .10$); and a significant effect of group ($F = 2.92$; $p = .03$). After controlling for age, the VCI group had significantly lower values on all seven quantitative tractography metrics compared to the healthy controls with the exception (Table 6.2). Effect sizes across the seven metrics were small to moderate ($\eta_p^2 = .25$ to $.35$); observed power ranged from $.83$ to $.97$. The group differences for six of the seven metrics were also significant when tested using the non-parametric Mann-Whitney U test (all $p < .05$). The sole exception was for *NNS* for which there was a non-significant trend ($p = .07$). Overall the results of the non-parametric analysis suggest that the violations of normality for these variables were not sufficient to invalidate the parametric MANCOVA procedure.

6.1.6 Conclusion

This proof-of-concept study demonstrates that metrics can be derived from DTI tractography data that provide clinically-relevant information about the microstructural integrity of cerebral white matter. The validity of our seven metrics is evidenced

by their significantly different values in patients with known vascular white matter injury vs. controls and their correlations with age, and cognitive test performance. Our metrics can be obtained for specific TOIs in individual datasets for group or individual comparisons. Currently, TOIs must be interactively selected which could lead to operator error and selection of incorrect fibers. However, a potential advantage of our approach is that since the models are defined in a 3D environment, TOI selection can occur without sophisticated image registration. Our quantitative tractography metrics are potentially a powerful alternative to the non-tractography scalar-based approaches to DTI analysis commonly used in current clinical studies. Our methods provide researchers with new tools for assessing the health of specific white matter pathways and its relation to cognition, behavior, and motor function.

6.2 Inverse Modeling

The scientific study of any physical system can be roughly divided into three steps:

1. Model design: define a model that describe the system and represent it by a list of parameters.
2. Forward modeling: discover the physical laws which given the model parameters will predict the values of the measurements on some observable parameters.
3. Inverse modeling: infer the model parameters with the measurements on some observable parameters.

These three steps are closely related. Progress in one of them often leads to advances in the other two.

In diffusion MRI, a great number of models have been proposed and studied. We can roughly separate them into local models and global models. Local models independently explain the measurements in a single location, namely a voxel in diffusion

MRI. Diffusion tensor, multi-tensor (Tuch, 2002), and PASMRI (Jansons and Alexander, 2003) are examples of local models. They can all be regarded as approximators of the diffusion propagator at a point location. Global models represent the diffusion MRI on the whole brain and incorporate the correlations between the measurements in different locations. The most salient feature of this correlation is the connectivity by the neural fibers across different regions of the gray matter. Streamline-based models such as the streamtubes model introduced in Chapter 3 are integrated in the first eigenvector field of the diffusion tensor field. Other methods used for generating global models include the level set method (Parker, 2000) and the importance sampling method (Brun et al., 2002).

Currently, the emphasis on the global models has been the inference of the connectivity, and rightly so, since it is probably the single most important feature of the white matter structure. But by focusing only on the fiber orientations, these global models tend to simplify or ignore other physical laws governing the imaging process such as the partial volume effect. As a result, inconsistency and spurious structures can appear. In this chapter, we place emphasis on both the forward modeling and inverse modeling steps to ensure the consistency between the model and the data.

We can formulate the inverse problem of global DWI modeling with Bayes' Theorem:

$$P(\omega|D, \mathcal{H}) = \frac{P(D|\omega, \mathcal{H})P(\omega|\mathcal{H})}{P(D|\mathcal{H})} \quad (6.8)$$

where $D = \{I^q\}$ is the collection of volumetric DWI images, I^q is the DWI image encoded by b-vector q , \mathcal{H} is the model assumption, and ω is the parameters of the model.

In section 6.2.1, we list our models and assumptions in detail. In section 6.2.2, we present a forward model that calculates $P(D|\omega, \mathcal{H})$. In section 6.2.3, we implement two optimization algorithms to improve the accuracy of the inverse modeling.

6.2.1 Model

As discussed in Chapter 2, in the human brain white matter, gray matter and cerebral-spinal fluid are isotropic materials that pose different constraints on the diffusion of water and thus produce different diffusion MR signals. Gray matter and CSF are isotropic materials. White matter, on the other hand, is anisotropic and produces different diffusion MR signals with different orientations and magnetic gradient directions. Therefore the diffusion MR signal in a voxel depends on both the fraction of the volume for each material and the geometric construction of the white matter in the voxel. If the white matter structure is coherent within the voxel, then we only need one compartment to model the white matter, otherwise we divided the white matter into individually coherent compartments. Our model of the diffusion imaging data thus has two components: the geometry of the white matter and the partial volume model.

The geometry of the white matter includes a collection of fibers

$$G = \{F_0, \dots, F_K\}$$

which are 3D curves.

Note that we do not model the apparent junction structures because to our knowledge there is no such structure in the anatomy. We intend to simulate merging, kissing, and crossing structures with the fiber model.

The partial volume model for each pixel denotes the percentage of the volume in the voxel occupied by each compartment:

$$V_p(i) = \{v_{GM}(i), v_{CSF}(i), v_{F_1}(i), \dots, v_{F_m}(i)\}$$

m is the number of fibers intersecting the voxel. i is the voxel index. The partial volume of the white matter v_{WM} is the sum of all v_{F_j} :

$$v_{WM}(i) = \sum_{j=1}^m v_{F_j}(i)$$

All the partial volumes in the voxel add up to 1:

$$v_{GM}(i) + v_{CSF}(i) + \sum_{j=1}^m v_{F_j}(i) = 1$$

Together, the model is represented by

$$W = (K, F_1, \dots, F_K, \{v_{GM}(i), v_{CSF}(i), v_{F_1}(i), \dots, v_{F_K}(i)\}; i = 1, \dots, M)$$

where K is the number of fibers, M is the total number of voxels.

Our model assumptions (\mathcal{H}) in equation 6.8 include:

1. Brain consists of gray matter, white matter and cerebral-spinal fluid (CSF).
2. Gray matter and CSF result in constant DWI signals along each gradient direction.
3. White matter consists of a group of cylindrical neural fibers that run from gray matter to gray matter without breaking or merging into other neural fibers along their paths.
4. Within each voxel, gray matter, CSF and each neural fiber are well separated into independent compartments. There is no interaction among different compartments. And within each compartment, there is free Gaussian diffusion.

Assumptions 1 and 3 are based on our anatomy knowledge. Assumptions 2 says that CSF and gray matter are uniform throughout the whole brain which is a simplification in the case of the gray matter. Assumption 4 is a rather arbitrary assumption that simplifies our calculation of the forward modeling given our current ignorance about the effect of mixing materials on the MR signal.

6.2.2 Forward Modeling

The forward modeling problem in diffusion imaging modeling can be stated as: given the model V_p, G , how do we get the diffusion MR signal for each voxel in the data volume? We begin by the equation that relates DWI signal to the tensor model (Basser

et al., 1994):

$$I = I_0 e^{-b:D}$$

I_0 is the intensity with no encoding (for a given material type typically a constant within a dataset), b is the diffusion encoding tensor, D is the tensor model for the diffusion.

Integrating it over the whole voxel, we get:

$$I_v = \int_{x \in v} I_x dx = \int_{x \in v} I_{0x} e^{-b:D_x} dx \quad (6.9)$$

With assumptions listed in the last section, we can rewrite equation 6.9 as:

$$I_v = I_{0csf} e^{-b:D_{csf}} v_{csf} + I_{0gm} e^{-b:D_{gm}} v_{gm} + \sum_{i=1}^K I_{0wm} e^{-b:D_i} v_{F_i} \quad (6.10)$$

where i indexes over the fibers in the voxel. D_i is calculated from F_i by rotating a constant linear tensor sampled from coherent white matter to the direction of F_i .

We assume the data measurements are normally distributed around the true value with the same variance. And we assume that the data measurements are independent in different locations. In order to get the maximum likelihood of the data, we need to minimize the quantity:

$$\chi^2 = \sum_{q=1}^m \sum_{v=1}^n \left(\frac{I_{v,q} - \hat{I}_{v,q}}{\sigma} \right)^2 \quad (6.11)$$

called χ^2 (Press et al., 1992), where q denotes the set of magnetic gradient directions, I and \hat{I} are the scanned and simulated signals, and σ is the variance of the distribution of the error between the two signals. We use the noise value in the scanned image for σ .

We demonstrate our forward modeling on the central part of a mid-coronal section through a normal human brain. The section includes the corpus callosum (red in Fig. 6.3(a)), cingulum bundles (green), and ventricle (white). Simulated DWI data

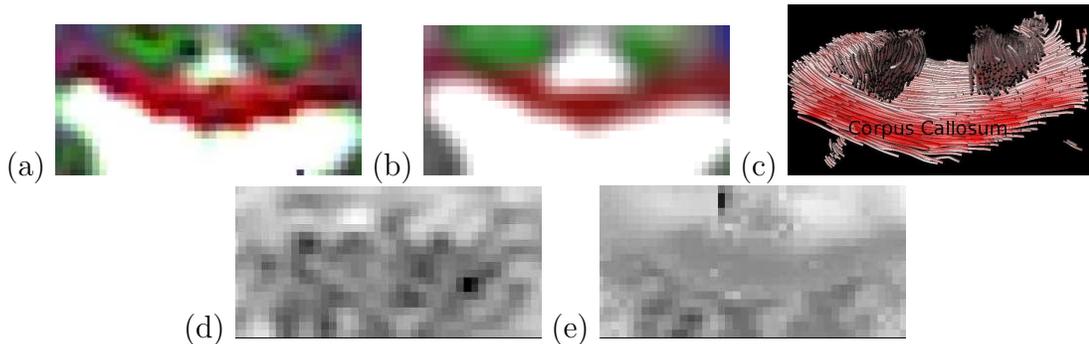


Figure 6.3: This figure shows the results of the forward imaging model on the central portion of a mid-coronal section through a human brain. The section includes parts of the cingulum bundles (green in (a)) and the corpus callosum (red in (a)). (a) and (b) show DTIs generated from the original DWIs and simulated DWIs, with the diagonal tensor elements mapped to red, green, and blue; (d) and (e) show a section in one of the original DWIs and the simulated DWIs; (c) shows the fiber model used in this forward modeling.

were created using a combination of the fiber model shown in Fig. 6.3(c) and partial volume classification calculated with FAST (Zhang et al., 2001b). Because of the relatively high noise in the original DWIs, it is difficult to visually compare the original DWI (Fig. 6.3(d)) and the simulated DWI (Fig. 6.3(e)). However, χ^2 between the real and simulated voxels in one DWI averages 1.6 over this section; this is close to the optimal value of 1.0 and supports that the model is tenable. A tensor-valued image was fit to the simulated DWIs. The diagonal elements of the original tensor image and the simulated tensor image are displayed on red, green, and blue in Figs. 6.3(a) and (b), showing good agreement and further supporting the forward imaging model.

6.2.3 Optimization

To solve the problem with a fully Bayesian approach, we need to set the prior information on $P(\omega|\mathcal{H})$ based on our knowledge in anatomy, such as that neural fibers project into the gray matter and have smooth curvatures, then search the parameter space for the optimum posterior probability $P(\omega|D, \mathcal{H})$. The forward modeling step applies in calculating the likelihood $P(D|\omega, \mathcal{H})$. Currently the forward modeling step takes about 15 minutes to finish for a $128 \times 128 \times 90$ diffusion MRI data set with 12

magnetic gradient directions. As a first experiment with optimization, we designed two optimization schemes. The results in this section are published in (Zhang and Laidlaw, 2006).

Two Optimization Schemes

For a fixed number of fibers N , we optimize the configuration of the fibers by minimizing the goal function χ^2 .

We define three kinds of rearrangements:

- Addition: add a fiber with a randomly selected seed point P .
- Removal: remove an randomly selected fiber C from the current configuration.
- Addition and removal: a combination of one addition and one removal.

The χ^2 difference $\Delta\chi^2 = \chi_2^2 - \chi_1^2$ is calculated after each step.

The greedy algorithm works as follows:

1. If the number of fibers is less than N , try addition. If $\Delta\chi^2 < 0$, accept the addition.
2. If the number of fibers is equal to N , try addition and removal. If $\Delta\chi^2 < 0$, accept the addition and removal.
3. repeat 2 until the decrease in χ^2 becomes discouraging.

To avoid letting the greedy algorithm get trapped in a local minimum, we also implemented a simulated annealing algorithm as follows:

1. If the number of fibers is less than N , try addition. If $\Delta\chi^2 < 0$, accept the addition. If $\Delta\chi^2 \geq 0$, accept with probability $\exp(\frac{-\Delta\chi^2}{kT})$.

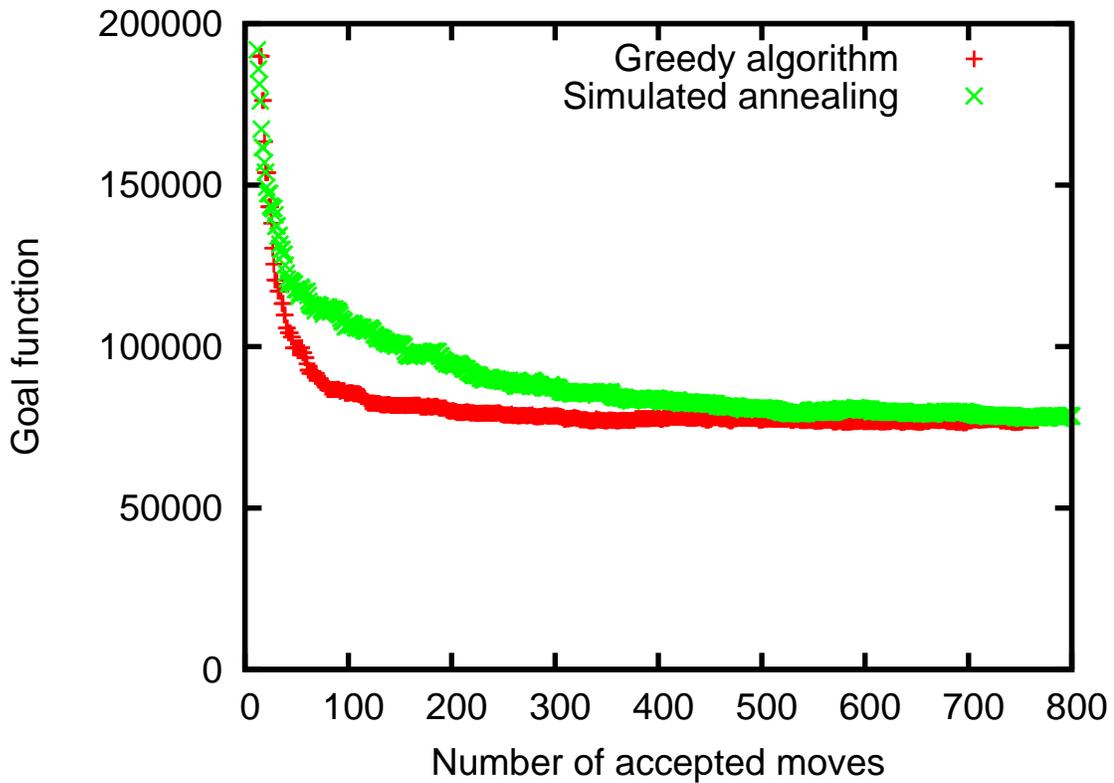


Figure 6.4: This plot shows the fibers optimization processes for the greedy algorithm and the simulated annealing. For both algorithms the χ^2 value converges to around 77,600. The greedy algorithm converges faster, and it does not appear to get trapped in a local minimum.

2. If the number of fibers is equal to N , try addition and removal. If $\Delta\chi^2 < 0$, accept the addition and removal. If $\Delta\chi^2 \geq 0$, accept with probability $\exp(\frac{-\Delta\chi^2}{kT})$.
3. repeat 2 until the number of reconfigurations reaches N_R or the number of steps reaches N_S . Reduce T by 10% if any of the two conditions are met.
4. repeat until T becomes close to 0.

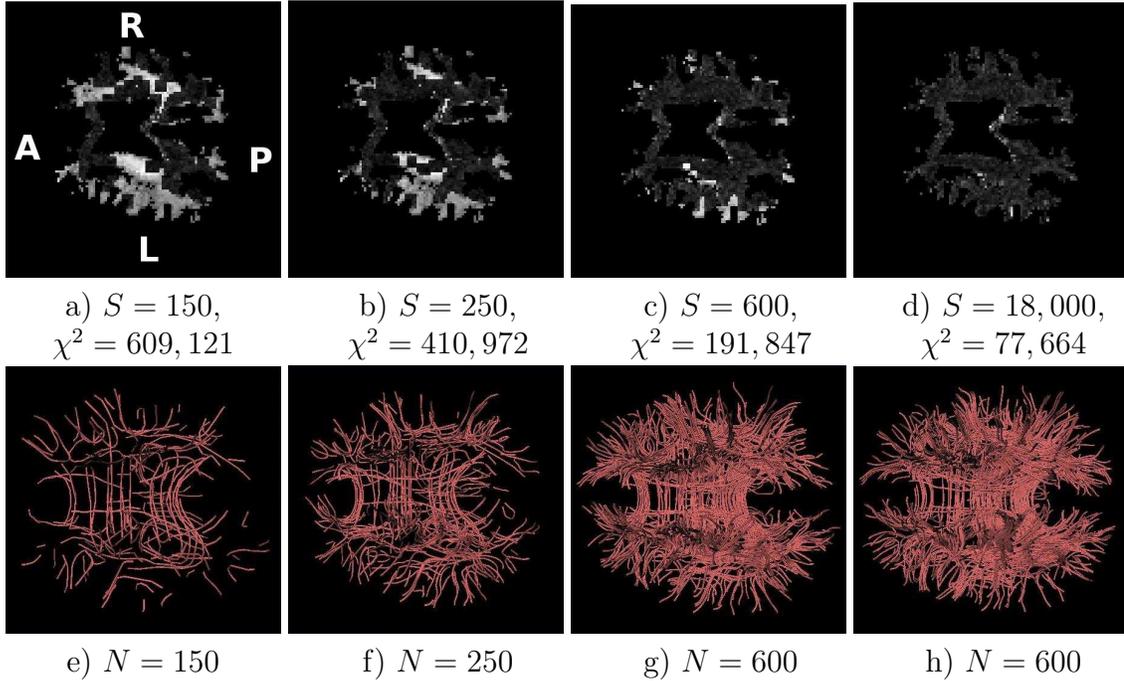


Figure 6.5: a), b), c) and d) show the snapshots of the value $\chi_v^2 = \sum_{q=1}^m \left(\frac{I_{v,q} - \hat{I}_{v,q}}{\sigma} \right)^2$ for each voxel v in the data volume during different stages of the simulated annealing. S is the number of accepted moves. e), f), g) and h) show the fiber model generated at these stages. N is the number of fibers. Note that g) and h) both have 600 fibers, but the optimization of the configuration lowered the χ^2 value from g) to h) significantly.

Results and Discussion

We ran the two optimization algorithms on the brain data set. Since we are interested in the white matter structures, we limit our χ^2 calculation to the voxels with more than 50% white matter. There are 84,025 such voxels. The seed point for a new fiber was selected within these voxels. We set the target number of the fibers to 600. The k in simulated annealing was set to 1. The temperature T was originally set to 1000. N_R was set to 300 and N_S was set to 600.

Fig. 6.4 shows the optimization processes for the greedy algorithm and the simulated annealing. For both algorithms the χ^2 value converges to around 77,600. The greedy algorithm converges faster, and it does not appear to get trapped in a local minimum.

Fig. 6.5 shows the snapshots of the value $\chi_v^2 = \sum_{q=1}^m \left(\frac{I_{v,q} - \hat{I}_{v,q}}{\sigma} \right)^2$ for each voxel v in the data volume during different stages of the simulated annealing. While the addition of the fibers quickly fills in large high χ_v^2 area, the repeated addition and removal steps gradually decrease the overall χ^2 value by reconfiguration.

For comparison, we also performed random seed point sampling for generating 600 fibers and calculated χ^2 values for these models. The experiment was repeated five times and $\chi^2 = 197,840 \pm 9,166$.

Conclusions

We presented a forward-modeling-based sampling of diffusion-tensor imaging (DTI) fibers. We built a forward model that generates DWIs from the DTI fibers based on multi-tensor modeling. We employ the sum of the difference between the simulated DWIs and the acquired DWIs as the goal function and optimize the placement of the DTI fibers with a greedy algorithm and a simulated annealing approach. The results show that with the same number of fibers, the optimized set of DTI fibers fit better to the data than randomly seeded DTI fibers. With the proposed moves of random addition, random removal and their combination, the greedy algorithm converges faster and does not seem to get trapped in local minimums.

This work has the potential to improve the accuracy of the white matter fiber models with an economic number of fibers.

Note that we made several simplifications in our forward model such as the white matter partial volumes calculation and the constant radius of the fiber bundles. Future improvement on the precision of the forward model may help increase the accuracy of the optimized model.

Chapter 7

Applications

Diffusion imaging is especially useful in studying fibrous structures such as white matter and muscle: the anisotropy information it provides reveals the fiber orientation in the tissue and can be used to map the white-matter anatomy and muscle structure in vivo (Basser et al., 2000). The diffusion coefficient measures a physical property of the tissue and the measurements can be compared across different times, locations, and subjects. Therefore, DTI has frequently been used to identify differences in white matter due to a variety of conditions. Normal conditions such as age and gender have been reported to affect anisotropy and diffusivity (Naganawa et al., 2003), and neural developments such as myelination (Klingberg et al., 1999), physical trauma such as brain injury (Nakayama et al., 2006), and neurodegenerative diseases such as multiple sclerosis (Simon et al., 2005) and HIV (Paul et al., 2006) have all been indicated by DTI studies to affect white-matter composition, location, or integrity.

The variety of DTI applications provides a valuable testbed for visualization methods. Indeed, without applications to guide the development of computational and visualization tools, these tools are far less likely to be useful. I introduce some of the applications of diffusion imaging in this section.

In section 7.1, we present a method to detect at-risk neuronal fibers in the corpus callosum related to distant focal demyelinating lesions in multiple sclerosis patients

by generating our streamtube models only in the neighborhood of these lesions. In section 7.2, we reveal the disparity in anisotropy values between the tumor and the non-tumor sides of the brain by using a quantification of white matter structural information contained within DTI data along with the qualitative visualization of the streamtubes and streamsurfaces models. In section 7.3, we evaluate the use of 3D DTI to follow white matter trajectories in intact fixed brains of several rare primate species.

7.1 Multiple Sclerosis

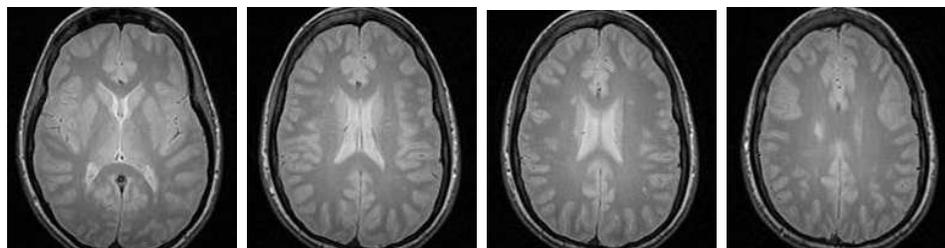


Figure 7.1: Fast-spin-echo images through the region of corpus callosum.

The work in this section is published in (Simon et al., 2005).

Multiple sclerosis (MS) is a chronic disease of the central nervous system that predominantly affects young adults during their most productive years. Pathologically, MS is characterized by the presence of areas of demyelination and T-cell predominant perivascular inflammation in the brain white matter. While demyelination is central to the onset of MS, focal inflammatory MS lesions are also characterized by axonal injury, including transections (Trapp et al., 1998). As a result, the study of axonal injury has become an important new area of MS research, and axonal injury is now considered an important contributor to irreversible injury, disability and possibly conversion to progressive stages of disease. In advanced stages of MS, there is

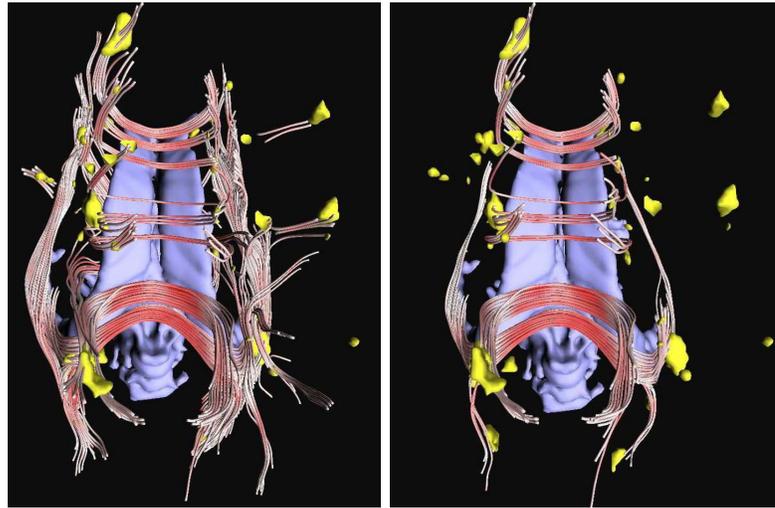


Figure 7.2: Left: Streamtube map culled to fibers intersecting lesions. Right: Streamtube map with additional culling to fibers passing through corpus callosum and intersecting lesion.

direct evidence for significant neuronal loss, especially in the corpus callosum (Evangelou et al., 2000). In early MS, there is circumstantial evidence for neuronal tract injury due to focal, inflammatory demyelinating lesions associated with a clinically isolated syndrome (CIS) based on neuronal tract (Wallerian) degeneration patterns in corticospinal tract (Simon et al., 2000) and across the corpus callosum (transcallosal bands) (Simon et al., 2001). The presence of such axonal injury is supported by an informative case study where confocal microscopy revealed empty myelin cylinders



Figure 7.3: Left: T2W image showing few focal T2-hyperintensities. Right: Sagittal fiber map showing in red the location of fibers in the midline corpus callosum that connect to the focal MS lesions.

in spinal cord of an MS patient with a distant subacute brainstem lesion (Bjartmar et al., 2001).

We report an MRI strategy to detect at-risk neuronal fibers in the corpus callosum related to distant focal demyelinating lesions. This strategy is necessary since a priori assumptions about the anatomical location of lesions relative to fiber tracts is often misleading.

We conducted a prospective longitudinal study including 18 CIS patients with an MRI. These patients have at least 2 characteristic T2-lesions, placing them at high risk for the development of MS (Group, 2002). MRI acquisition was at 3T and included 3mm thick non-gapped proton density/T2 series and sagittal fast spin echo (FSE) T2-weighted imaging with 3mm non-gapped slices (see Figure 7.1). A diffusion tensor imaging sequence is run in the axial plane with sets of diffusion tensor images acquired by echo-planar technique using a slice thickness of 5.1 mm and each series shifted from the prior series by 1.7 mm inferiorly, in 25 gradient directions with maximal b-value of 1000. Lesions are segmented by in-house semi-automated segmentation routines based on the FSE image data. A mutual information algorithm was used to determine the optimal 12 DOF affine transformation registering T2 axial with B0 diffusion data. This transformation was then applied to the segmented T2 data. A set of streamtubes (Zhang et al., 2003) is generated by seeding from within the lesions. Each streamtube follows the fastest direction of diffusion until it transcends the dataset boundary, hits a region of low linear anisotropy, or curves excessively. The redness on the streamtubes represents linear anisotropy. Cerebrospinal fluid in the ventricles is represented by a blue surface, and lesions are shown as yellow surfaces.

Case CIS 2 is shown as a representative example. This individual presented unilateral optic neuritis, and had a positive MRI. Structural images suggest lesions potentially related to the corpus callosum. The streamtube tractography suggests lesions intersecting numerous fiber pathways, and provides (best viewed in 3D-not shown) the explicit pathways that are potentially affected by the lesion. Further culling to

fibers that intersect the corpus callosum shows a large potential volume of involvement (see Figure 7.2). Unedited initial runs contain assignment errors, for example some fibers running anterior-posterior. These can be removed through further culling using objective template driven criteria or based on expert review of the images in 3D. The sagittal fiber-at-risk map shows that the distribution of the at-risk fibers does not correspond to the limited areas of abnormal appearing white matter conventionally determined by T2-hyperintensity (see figure 7.3).

Our streamtube strategy was devised to allow us to interrogate voxels at-risk by quantitative MRI techniques. Using our methodology risk is assessed based on connection via fiber tracts to typical focal MS lesions. Our ultimate goal is to test the hypothesis that focal inflammation results in Wallerian degeneration that can be explicitly assayed over time in vivo. Initial attempts to relate focal MS lesions to corresponding regions in callosum have been unsatisfactory to date, as fiber tracts from corpus callosum curve in 3D space, and lesions such as those located lateral to the mid-sagittal corpus callosum previously assumed to lie in transcallosal tracts often do not intersect the fibers of interest. Using the tractography strategy outlined here, we can explicitly determine the locations of voxels in corpus callosum that are most likely to be injured by secondary degeneration, and we can follow these at-risk locations prospectively in longitudinal MRI studies. In addition, the tractography approach provides a unique perspective into early MS pathology allowing us to visualize the full extent of potentially injured fibers from multiple lesions. This strategy is applicable to any neuronal tract that can be imaged with high quality diffusion tensor and structural MRI.

7.2 Tumor Analysis

The work in this section is published in (Zhang et al., 2004).

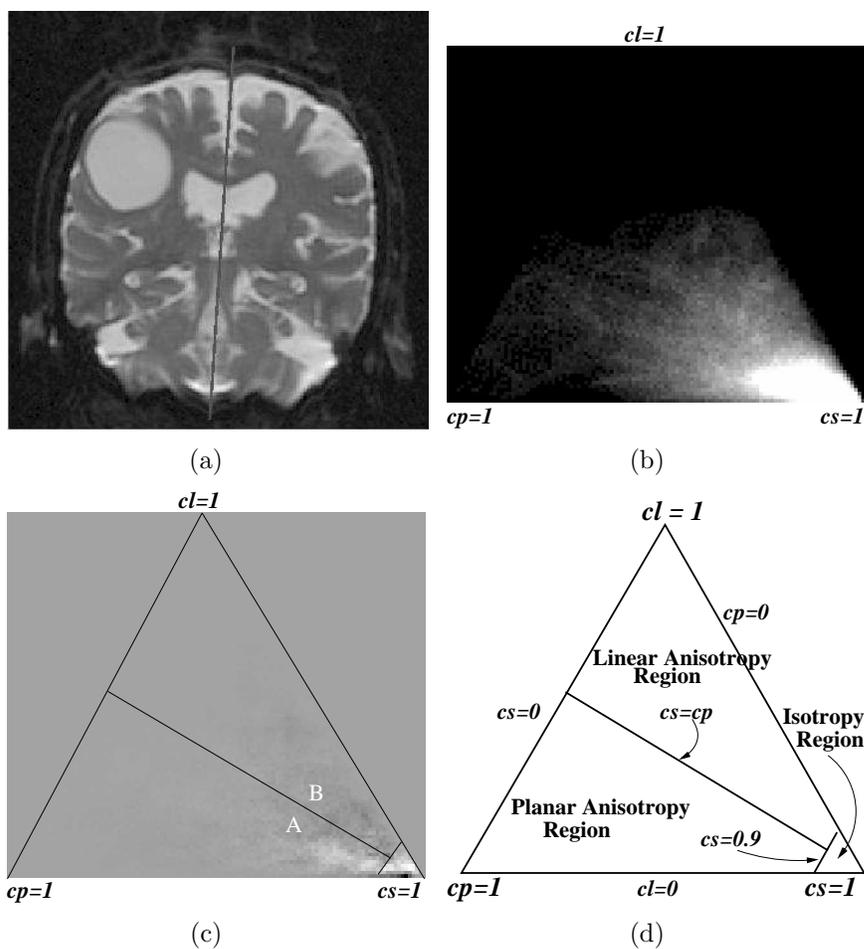


Figure 7.4: The process of creating 2D barycentric histograms from DTI data. a: For a given slice, the brain is divided in half with a straight line following the falx cerebri on the T_2 -weighted EP image. b: A barycentric histogram is generated for each side of the brain. In this example, the histogram comes from the tumor-bearing section shown in a. Note the large number of voxels near the $c_s = 1$ vertex, reflecting the cystic component of the tumor. c: The difference histogram obtained by subtracting normalized barycentric histograms calculated from tumor-bearing and contralateral sections. In this difference histogram, zero maps to medium gray because the difference is signed. Note that the most striking difference occurs near the $c_s = 1$ vertex. d: The definition of the diffusion anisotropy regions in the barycentric histogram. These regions are chosen to capture the differences between regions A and B in part c.

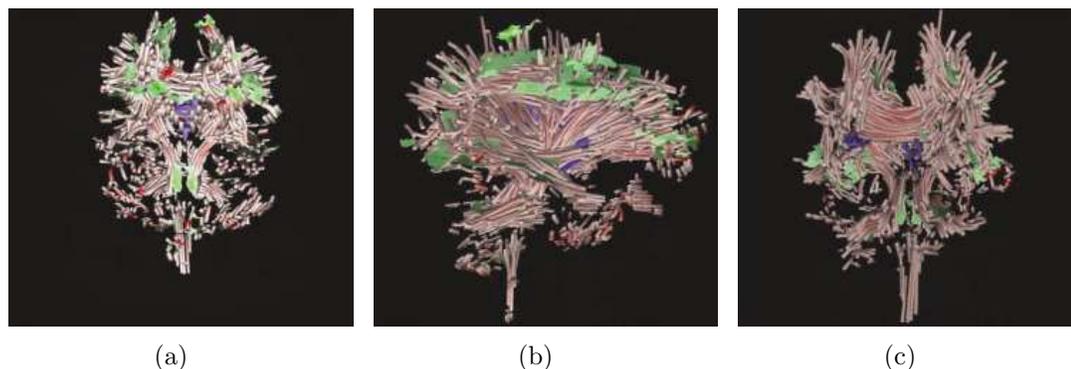


Figure 7.5: Geometric models representing the diffusion metric data obtained from a 33-year-old female volunteer viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. Linear diffusion is represented by red streamtubes, planar diffusion by green streamsurface fragments, and CSF by the blue surface. Note the symmetry between two sides of the brain in this normal subject.

One particularly interesting clinical application of the DTI visualization techniques is in the study of intracranial tumors and how they affect surrounding brain structure. Several groups have already studied patterns of white matter tract disruption adjacent to brain tumors using DTI data. Wieshmann *et al.* (Wieshmann et al., 2000) found evidence of displacement of white matter fibers of the corona radiata in a patient with low-grade glioma when compared with spatially normalized data collected from 20 healthy volunteers. The data obtained from this particular patient were consistent with displacement rather than destruction of fibers due to the mass effect of the tumor. Mori *et al.* (Mori et al., 2002) found evidence of displacement and destruction of the superior longitudinal fasciculus and corona radiata in two patients with anaplastic astrocytoma. Gossel *et al.* (Gossel et al., 2002) observed distortion of the pyramidal tract in a patient with a high-grade glioma. Finally, Witwer *et al.* (Witwer et al., 2002) found evidence of white matter tract edema, infiltration, displacement, and disruption in 10 tumor patients with a range of brain malignancies. These studies show that while mapping white matter fiber tract disruption patterns

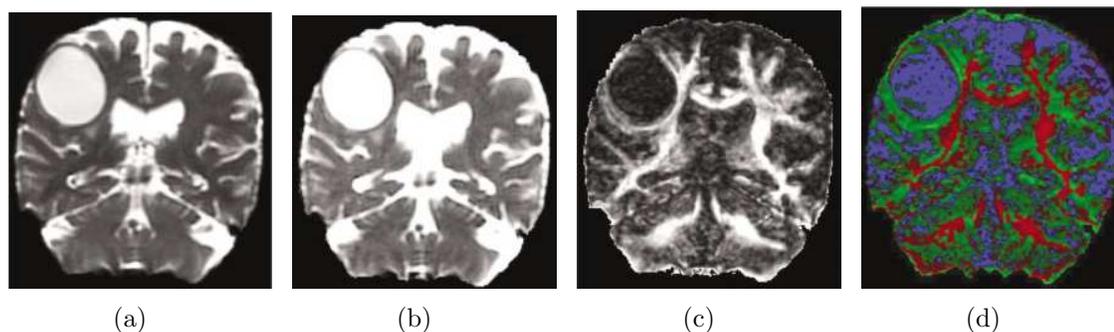


Figure 7.6: Maps of (a) T_2 -weighted signal intensity, (b) $\langle D \rangle$, (c) FA, and (d) color composite diffusion metric data obtained from a 72-year-old male patient with a right parietal metastasis of a lung adenocarcinoma. In d, maps of linear (red channel), planar (green), and spherical (blue) diffusion are overlaid so that the color in each voxel represents the relative magnitude of the three diffusion metrics. Note the high $\langle D \rangle$, low FA, and spherical (isotropic) diffusion of the tumor/edema volume. Also note that the tumor/edema volume is surrounded by a large region of planar diffusion.

in neoplasia is technically possible, more work is still required to validate DTI tractography methods if they are to become a reliable tool for investigating brain structure and the pathophysiology of disease. Achieving this goal will require the development of robust protocols for measuring white matter tract topology, quantifying the natural biological variation in tract orientation in the normal brain, and applying these analyses to carefully characterized patient groups.

As a first step toward developing such robust protocols, this work describes a method that permits the 3D visualization and quantification of the diffusion anisotropy information contained within DTI data. In this method, regions within the brain are defined as possessing linear, planar, or spherical (isotropic) diffusion, based on the geometric properties of D (Westin et al., 1997a). Visualization of these diffusion metrics is realized by generating streamtube and streamsurface models to represent regions of linear and planar diffusion (Westin et al., 2002). Quantification of differences in diffusion anisotropy between regions-of-interest (ROIs) is then achieved by analyzing 2D barycentric histograms created from the complete distribution of diffusion metric values measured in each region (Kindlmann and Weinstein, 1999; Alexander et al.,

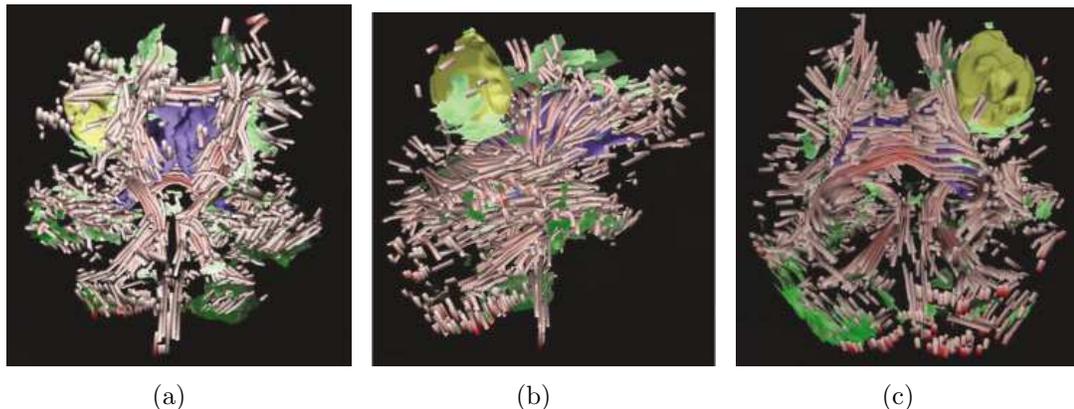


Figure 7.7: Geometric models representing the diffusion metric data obtained from the 72-year-old male patient viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. Linear diffusion is represented by red streamtubes, planar diffusion by green streamsurface fragments, and CSF by the blue surface. The yellow surface, which is rendered from the signal abnormality on the T_2 -weighted EP images, depicts the shape of the tumor/edema volume. Note that the tumor/edema volume is partly surrounded by regions of planar diffusion.

2000). Here the feasibility of using this approach in mapping normal brain structure and characterizing how neoplasia alters this structure is investigated.

7.2.1 Method

All MRI data were obtained using a GE Signa LX 1.5 T (General Electric, Milwaukee, WI) clinical scanner, equipped with a self-shielding gradient set ($23mTm^{-1}$ maximum gradient strength, $120Tm^{-1}s^{-1}$ slew rate, and 60 cm inner diameter horizontal bore) and a manufacturer-supplied birdcage quadrature head coil. The component diffusion-weighted (DW) images of the DTI data set were acquired using a single-shot, spin-echo, echo-planar imaging (EPI) sequence in which two symmetric trapezoidal gradient pulses were inserted around the 180° refocusing pulse in the required gradient channel. Sets of coronal DW-EP images ($b = 0$ and $1000smm^{-2}$) were collected with diffusion gradients applied sequentially along six non-collinear directions (13). Six acquisitions consisting of a baseline T_2 -weighted EP image (G^0) and six DW-EP

images (G^1 - G^6), for a total of 42 images, were collected per slice position. Parameters for the DW-EPI sequence included a repetition time (TR) of 15 s, an echo time (TE) of 102.5 ms, a field of view (FOV) of 200 200 mm, an acquisition matrix of 128×128 (zero-filled to 256×256), and 36 contiguous coronal slices of 5-mm thickness. Therefore, each voxel had dimensions of $0.78 \times 0.78 \times 5.0\text{mm}$, with an inherent voxel resolution of $1.56 \times 1.56 \times 5.0\text{mm}$.

Bulk patient motion and eddy current-induced artifacts were then removed from the component EP images using a 3D computational image alignment technique (Jenkinson and Smith, 2001). First, EP images with the same b-matrix were rigidly aligned to remove patient motion. Second, affine transformations were used to align the DW-EP images acquired with diffusion gradient directions $G^2 - G^6$ to the first gradient direction G^1 . Third, all DW-EP images were aligned using affine transformations to the T_2 -weighted (G^0) acquisition. Then the set of six component EP images for each gradient direction was averaged to give seven high signal-to-noise ratio (SNR) images for each slice. Finally, from this MRI data D was calculated in each voxel from the signal intensities in the component EP images (Basser et al., 1994).

The three diffusion metrics can also be employed to quantify diffusion anisotropy in different brain regions using a barycentric space histogram approach (11,12). In this method a 2D triangular barycentric space containing all possible diffusion anisotropy values is defined, with c_l , c_p , and c_s being the barycentric coordinates. Asymmetries in the distribution of diffusion anisotropy values between the right and left hemispheres can then be characterized by analyzing the barycentric histograms obtained from each side of the brain. Using data from the tumor patient, Fig. 7.4 illustrates how this process was implemented. First, for a given slice the brain was divided in half with a straight line following the first cerebri on the T_2 -weighted EP image coregistered with the DTI data (Fig. 7.4(a)). A barycentric histogram was then generated for each side of the brain using the T_2 -weighted EP image to identify voxels that were located within the cerebrum (Fig. 7.4(b)). To eliminate the effect of size differences

between the two ROIs, the histograms were normalized so that the sum of all values within them was unity. The histograms were then subtracted to show how the two sides of the brain differ in diffusion anisotropy (Fig. 7.4(c)). Regions of linear, planar, and spherical diffusion within the 2D barycentric space were then defined. In order to capture the differences between regions A and B in the difference histogram close to the $c_s = 1$ apex (Fig. 7.4(c)), the isotropic region was defined as $c_s > 0.77$. The remaining space was split along the $c_l = c_p$ line. Thus, linear diffusion was defined as $c_s < 0.77$ and $c_l > c_p$, while planar diffusion was defined as $c_s < 0.77$ and $c_l < c_p$ (Fig. 7.4(d)). The number of voxels that fell within each of these regions was determined for a given number of slices for both sides of the brain to provide a 3D analysis. In the case of the four healthy volunteers, the slices chosen covered the entire cerebrum. However, for the tumor patient the analysis was performed not only for all slices covering the cerebrum, but also on just those slices in which signal hyperintensity was visible on the coregistered T2-weighted EP images (i.e., the tumor/edema volume). To quantify the degree of asymmetry in the number of $c_l(nc_l)$, $c_p(nc_p)$, and $c_s(nc_s)$ voxels in each ROI between the right and left sides of the brain, an asymmetry index ε of the form

$$\varepsilon(nc_l, nc_p, nc_s) = 100 \frac{\text{right}(nc_l, nc_p, nc_s) - \text{left}(nc_l, nc_p, nc_s)}{\text{right}(nc_l, nc_p, nc_s) + \text{left}(nc_l, nc_p, nc_s)} \quad (7.1)$$

was evaluated (17). Finally, the percentage of voxels for each ROI characterized as possessing linear, planar, or spherical diffusion was determined.

7.2.2 Results

Figure 7.5 displays a set of geometric models representing the diffusion metric data obtained from volunteer 4, a 33-year-old female, viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. In this figures the red streamtubes represent structures characterized by linear diffusion, the green streamsurface fragments represent structures characterized by planar diffusion, and the blue surface represents

Table 7.1: Comparison of Diffusion Metric Data Obtained From Right and Left Hemispheres for Four Normal Healthy Volunteers

Volunteer 1 (male/43)			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	459.84	459.45	0.04
Planar	956.96	956.41	0.03
Spherical	3,352.18	3,308.82	0.65
Percentage	Right	Left	
Linear	9.6	9.7	
Planar	20.1	20.2	
Spherical	70.3	70.0	
Volunteer 2 (male/43)			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	623.60	597.47	2.14
Planar	1,221.74	1,262.98	-1.66
Spherical	3,507.74	3,355.15	2.22
Percentage	Right	Left	
Linear	11.7	11.5	
Planar	22.8	24.2	
Spherical	65.5	64.3	
Volunteer 3 (female/26)			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	407.85	390.49	2.17
Planar	952.51	997.77	-2.32
Spherical	3,148.14	3,056.85	1.47
Percentage	Right	Left	
Linear	9.0	8.8	
Planar	21.1	22.4	
Spherical	69.8	68.8	
Volunteer 4 (female/33)			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	526.12	551.75	-2.38
Planar	1,003.63	1,045.22	-2.03
Spherical	3,371.16	3,378.21	-0.10
Percentage	Right	Left	
Linear	10.7	11.1	
Planar	20.5	21.0	
Spherical	68.8	67.9	
Mean \pm SD			
	$\varepsilon(\%)$		
Linear	0.49 \pm 2.16		
Planar	-1.50 \pm 1.05		
Spherical	1.06 \pm 1.01		
Percentage	Right	Left	
Linear	10.3 \pm 1.2	10.3 \pm 1.3	
Planar	21.1 \pm 1.2	22.0 \pm 1.8	
Spherical	68.6 \pm 2.2	67.8 \pm 2.5	

Table 7.2: Comparison of Diffusion Metric Data Obtained From Right and Left Hemispheres in a 33-Year-Old Female Volunteer Scanned on Three Separate Occasions

First exam			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	526.12	551.75	-2.38
Planar	1,003.63	1,045.22	-2.03
Spherical	3,371.16	3,378.21	-0.10
Percentage	Right	Left	
Linear	10.7	11.1	
Planar	20.5	21.0	
Spherical	68.8	67.9	
Second exam			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	513.92	521.88	-0.77
Planar	970.85	1,007.61	-1.86
Spherical	3,106.81	3,187.83	-1.29
Percentage	Right	Left	
Linear	11.2	11.1	
Planar	21.1	21.4	
Spherical	67.7	67.6	
Third exam			
Number($\times 1000$)	Right	Left	$\varepsilon(\%)$
Linear	534.78	555.55	-1.91
Planar	1,033.60	1,082.56	-2.31
Spherical	3,388.47	3,266.30	1.84
Percentage	Right	Left	
Linear	10.8	11.3	
Planar	20.9	22.1	
Spherical	68.4	66.6	
Mean \pm SD			
	$\varepsilon(\%)$		
Linear	-1.68 ± 0.83		
Planar	-2.07 ± 0.23		
Spherical	0.15 ± 1.58		
Percentage	Right	Left	
Linear	10.9 ± 0.3	11.2 ± 0.1	
Planar	20.8 ± 0.3	21.5 ± 0.5	
Spherical	68.3 ± 0.6	67.4 ± 0.7	

Table 7.3: Comparison of Diffusion Metric Data Obtained From Different Sized Regions-of-Interest Covering Tumor/Edema (Ipsilateral) and Contralateral Volumes in a 72-Year-Old Male Patient With a Right Parietal Metastasis of a Lung Adenocarcinoma

Hemisphere			
Number($\times 1000$)	Ipsilateral	Contralateral	$\epsilon(\%)$
Linear	392.04	463.48	-8.35
Planar	1,164.38	1,084.59	3.55
Spherical	3,129.50	3,112.84	0.27
Percentage			
Linear	8.4	9.9	
Planar	24.8	23.3	
Spherical	66.8	66.8	
Slices covering tumor/edema volume			
Number($\times 1000$)	Ipsilateral	Contralateral	$\epsilon(\%)$
Linear	255.85	332.64	-13.05
Planar	722.63	633.67	6.56
Spherical	1,513.74	1,429.83	2.85
Percentage			
Linear	10.3	13.9	
Planar	29.0	26.4	
Spherical	60.7	59.7	

structures that are hyperintense on the coregistered T2-weighted EP images, namely cerebrospinal fluid (CSF) in the ventricles. (Note that the colors in these figures should not be confused with the standard color scheme used to represent the direction of the principal eigenvector of D in two dimensions.) The red streamtubes follow a number of major white matter fiber tracts (e.g., the corpus callosum, internal capsule, corona radiata, and optic radiations), while there are a number of small green streamsurfaces interspersed among them.

The results from the analysis of the 2D barycentric histograms obtained from the right and left hemispheres in the four normal volunteers are presented in Table 7.1. For all subjects the spherical diffusion class, which predominantly reflects gray matter and CSF, contains the greatest number of voxels (approximately 70% of all the voxels in the brain), while the linear diffusion class contains the fewest voxels (approximately 10% of all voxels). The planar diffusion class contains approximately 20% of all voxels. In these four volunteers there is very little asymmetry between the number of

linear, planar, and spherical diffusion voxels found in each hemisphere, with ε varying between 2.38% and 2.22%. Furthermore, the average value of ε for the three diffusion classes is in the range of 1.5%. Table 7.2 displays the results of the analysis of the diffusion metric data obtained from volunteer 4 on three separate occasions. The values of ε for linear and especially planar diffusion voxels are reasonably reproducible, with ε being 2.38%, 0.77%, and 1.91% for linear diffusion, and 2.03%, 1.86%, 2.31% for planar diffusion. Furthermore, the values for the percentage of voxels assigned to each diffusion class are very similar in the three examinations. Figure 7.6 shows maps of (a) T2-weighted signal intensity, (b) $\langle D \rangle$, (c) FA, and (d) color composite diffusion metric data for a coronal section through the tumor/edema volume in the 72-year-old male patient. Figure 7.7 presents a set of geometric models representing the diffusion metric data obtained from this patient viewed from the (a) anterior, (b) right lateral, and (c) posterior directions. In both these figures the growing metastasis, which probably originated at the gray/white matter junction of the brain, is clearly seen to affect the surrounding white matter fibers, thus altering the measured diffusion anisotropy. Figure 7.6(c) shows the total loss of white matter fiber structure within the tumor/edema volume, while Figure 7.6(d) shows that the tumor/edema volume is surrounded by an extensive region of planar diffusion. This shell-like pattern in planar diffusion (green surface) surrounding the tumor/edema volume (yellow surface) is also evident in Figure 7.7. The results obtained from the analysis of the 2D barycentric histograms calculated for different-sized ROIs covering the ipsilateral and contralateral regions in this patient are shown in Table 7.3. The first set of three columns shows data obtained from an ROI covering both hemispheres, which is the ROI used in the volunteer studies above. The second set of three columns shows data obtained from an ROI formed from just those slices that cover the tumor/edema volume on the coregistered T2-weighted EP images. For both ROIs there is a significant asymmetry between the number of linear and planar diffusion voxels identified in the ipsilateral and contralateral regions, and this asymmetry becomes more pronounced

as the ROI becomes smaller. As suggested by Fig. 7.6(d) and 7.7, this asymmetry arises due to a reduction in the number of linear diffusion voxels and an increase in the number of planar diffusion voxels in the ipsilateral compared with the contralateral side of the brain. Thus, while the largest values of ε seen in the volunteer data sets is 2.38%, the effect of the tumor is to increase this right/left asymmetry to 8.35% for linear and 3.55% for planar diffusion. This asymmetry is even greater for the ROI formed from those slices that cover the tumor/edema volume, with ε being 13.05% for linear diffusion and 6.56% for planar diffusion.

7.3 The Aye-aye Brain

The work in this section is published in (Kaufman et al., 2005).

We have evaluated the use of 3D DTI to follow white matter trajectories in intact fixed brains of several rare primate species. This work is a close collaboration among researchers at the Pittsburgh NMR Center, Brown University and Caltech. These data show that 3D DTI at microscopic resolution in fixed samples is highly feasible, but computationally challenging due to the massive datasets involved. After the DTI datasets have been acquired in the intact specimen, it is processed for histology and the sections are stained for myelin. The histology slices are then aligned with the corresponding DTI 'slices' in order to provide feedback about the fidelity of DTI-based models of fiber tracts.

Figure 7.8 shows three orthogonal slices taken from a 3D DTI dataset in the intact fixed brain of a mouse lemur (*Microcebus murinus*) which is one of the world's smallest primates (< 90 grams total body weight). The total brain size is only slightly bigger (~25%) than a mouse brain. Data was acquired at 70 μm isotropic resolution and the 3D DTI was computed from a total of 28 diffusion-weighted volumes. The data was rendered in such a way that the diagonal elements of the effective diffusion tensor are displayed as a single composite image, where each color channel (red, green, blue) is

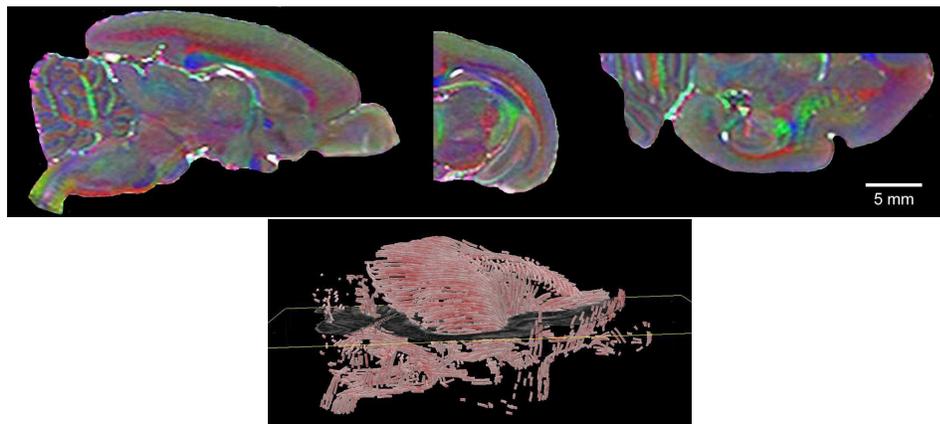


Figure 7.8: 3D DTI in the intact fixed brain of the mouse lemur. Shown on the left are hemi-segments of para-sagittal, coronal, and horizontal slices taken from the 3D volume. A composite RGB image was formed from the diagonal tensor elements, where red=rostral-caudal, green=dorsal-ventral, and blue=lateral. A dominant color in a region indicates the directionality of a track. Diffusion weighted volumes (28) were acquired using a 3D PGSE sequence. The diffusion gradients were applied along 7 different directions, and four values of b-matrix were used along each direction (maximum $b \sim 2000 \text{ mm}^2/\text{s}$). The data were acquired with $512 \times 256 \times 256$ points and a $\sim 70 \text{ um}$ isotropic resolution. On the right a streamtube representation of the dataset. For anatomical context, a T2-weighted grayscale para-sagittal slice from the 3D volume is shown in the yellow plane.

assigned to a diagonal tensor element. Regions showing predominantly a single color represent fiber tracts with a high degree of coherence. Voxels containing splayed or crossed tracts cannot be resolved using DTI and is one of the limitations of the technique; in this case voxels may appear gray or faintly-colored. Gray matter also appears gray due to the diffusion isotropy in this tissue. Figure 7.8(right) shows a streamtubes visualization of the major fiber tracts.

Figure 7.9 shows a direct comparison among anatomical MRI, DTI, and histology in the same image plane. The data in Figure 7.9 was acquired in the brain of a larger ($\sim 1 \text{ kg}$) prosimian, the Aye-Aye lemur.

In the proposed animal work we focused on large coherent tracks that are readily identifiable in DTI data sets; the same tracks will be identified in “best-fit” histology slices. The DTI-apparent tracts can then be viewed histologically with high magnification microscopy to describe the fiber trajectory. A DTI slice from the Aye-Aye



Figure 7.9: Comparison of different imaging modalities in the brain of the Aye-Aye lemur. 3D DTI data were acquired in a fixed, intact, brain at ~ 100 μm resolution. Shown is the caudal portion of a single para-sagittal slice from the 3D dataset. The panels from left-to-right represent the T2-weighted anatomical image, an RGB composite image of the diagonal elements of the diffusion tensor (see above), and the corresponding fiber-stained histology section. Note that the anatomical image and histology essentially provide scalar information about the location of the fiber tracts, whereas the DTI provides fiber location as well as directionality of the fibers. Scale bar=1cm.

is shown in Figure 7.10(left), where two major fiber pathways, the cingulum bundle and the splenium of the corpus callosum, are indicated. The histology section that best corresponds to this DTI slice is identified (Figure 7.10(center)), and the regions of interest are imaged at different magnifications using optical microscopy. From the optical micrographs, the tracts of interest can be qualitatively described (i.e. cingulum bundle fibers along image plane, and fibers in the splenium of the corpus callosum perpendicular to the image plane). With this qualitative description at hand, the representation accuracy of DTI-based fiber analyses can readily be evaluated. Figure 7.10(right) shows one such DTI rendering of the cingulum bundle and the splenium of the corpus callosum using a normalized diffusion ellipsoid representation (Ahrens et al., 1998; Laidlaw et al., 1998). Note that the prolate ellipsoid orientation closely matches the qualitative description of the fiber tracts obtained from the optical micrographs of the cingulum bundle and the splenium of the corpus callosum. Thus, DTI in animal models, in conjunction with fiber-stained histology in the same brains, helps to validate fiber bundle models. Direct correlation between fibers tracts observed in histology and tracts generated by DTI models will be an

integral part of our initial validation studies. We have recently devised a powerful analysis approach for making quantitative comparisons between histology and DTI. As a pilot study, we have prepared a complete histological series from the same hemisphere of the Aye-Aye brain that was used for DTI (See Figs. 7.9 and 7.10). Using the Gallyas silver stain for myelinated fibers (Gallyas, 1979), it is possible to trace individual axons at high magnification and fiber bundles at low magnification. These tracings can then be analyzed using directional statistics to quantify fiber orientation and coherence. We applied a polar histogram or "rose diagram" analysis, previously described by Wu et al. (Wu et al., 2004), to measure the mean fiber orientation and spread over numerous microscopic regions of interest. The polar histograms can then be compared to DTI. Additional details concerning these methods are provided in Section D.

Fig. 7.11 show results from the posterior portion of the cingulum bundle. This is the same region that is rendered using diffusion ellipsoids in Fig. 7.10. The mean fiber orientation progresses from a rostral-caudal orientation at the top of this region, to a superior-inferior arrangement below (Fig. 7.11A). These results are in good agreement with fiber orientations computed from DTI, where Fig. 7.11B shows a gradual progression from red voxels to green voxels in an RGB-pseudocolor image.

7.4 Conclusions

In this chapter, I presented three applications of DTI visualization and analysis techniques.

In section 7.1, we utilize streamtube diffusion tractography to identify neuronal fibers that intersect focal lesion and pass through a region of interest, in this case corpus callosum, where distal (to focal lesion) interrogation can be accomplished. A new class of tissue is defined called fibers-at-risk (FAR) through connectivity, which is distinct from the abnormal appearing white matter and comprises an impressive

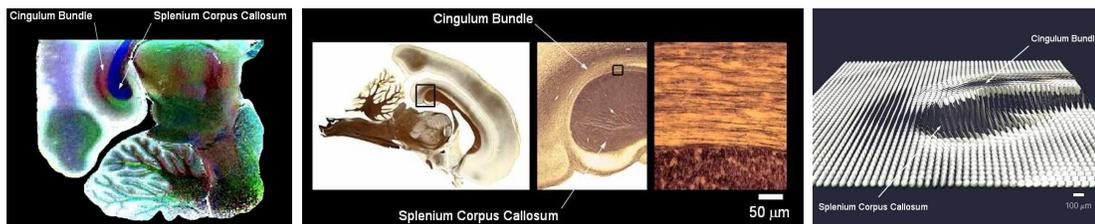


Figure 7.10: Left: DTI slice from the Aye-Aye data set showing targeted fiber tracts for histological comparisons. DTI slice is a RGB composite image of the diagonal elements of the diffusion tensor (see above). Center: ‘Best-match’ histology slice to the DTI image shown on the left. Micrographs at several magnifications are shown; the highest magnification (right panel) shows the boundary between the cingulum bundle and the splenium of the corpus callosum. Histology was performed on the same brain as the MRI by embedding in celloidin, serially sectioning at 40 μm using a sliding microtome, and staining with the Gallyas technique. Right: Diffusion ellipsoid rendering showing the relative diffusion anisotropy on a voxel-by-voxel basis. The observed orientation of the diffusion ellipsoids is agrees qualitative with the histology.

fraction of NAWM in the early stages of disease.

In section 7.2, we combine the qualitative visualization of the brain with quantification of differences in diffusion anisotropy between different regions of interest. we found that in a patient with a metastatic brain lesion there was marked asymmetry in both linear and planar diffusion between comparable ipsilateral and contralateral regions, with a significant reduction in the number of linear diffusion voxels and an increase in the number of planar diffusion voxels in the tumor-bearing hemisphere. These results demonstrate the potential of our approach to characterize brain structure in both healthy and diseased subjects.

In section 7.3, analysis of DTI data in the anterior cingulum bundle of an aye-aye brain demonstrates a strong correlation between fiber spread as measured from histological sections and fiber spread as measured from DTI.

These applications provide guidance, feedback and validation for the development of computational and visualization tools.

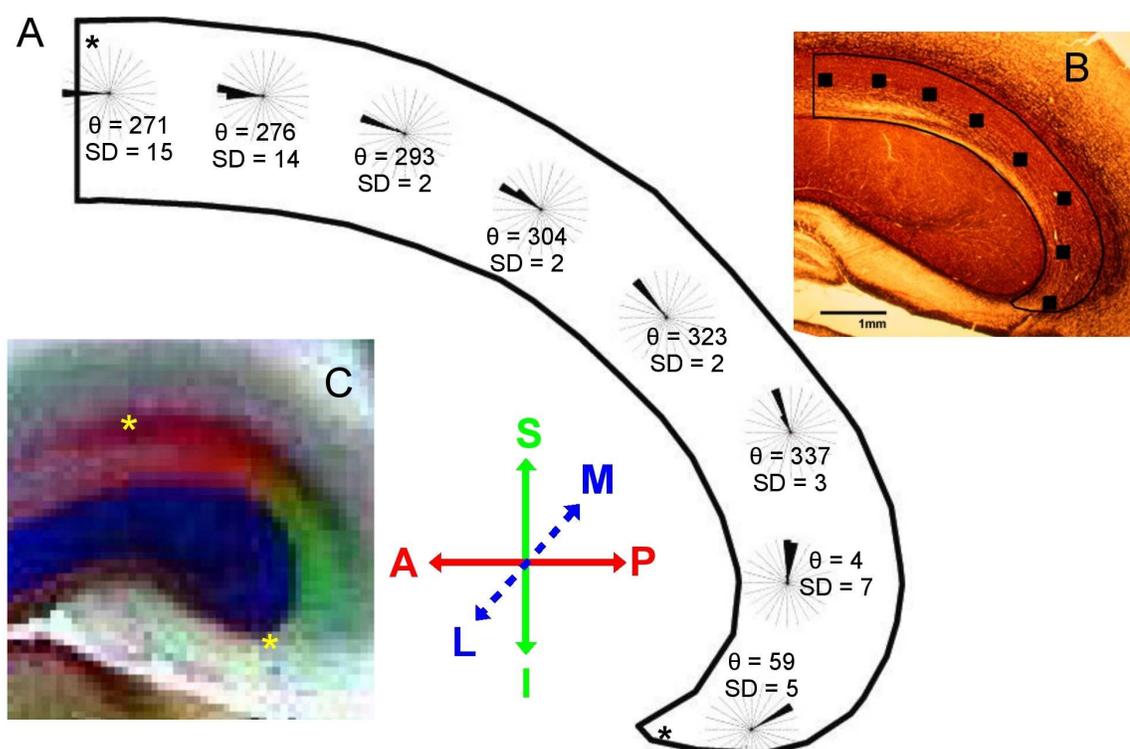


Figure 7.11: Fiber tracings from a sequence of sites in the cingulum bundle in the Aye-Aye brain. (A) depicts polar histograms centered on each sample site depicting fiber orientation and spread; the histograms encode the mean polar angle (θ) of the fiber orientation and the standard deviation (SD). Data from (A) were extracted directly from microscopic images of the histology slice shown in (B). (C) shows the corresponding DTI slice of the cingulum bundle rendered in an RGB format (red=anterior-posterior, green=superior-inferior, blue=medial-lateral).

Chapter 8

Summary and Conclusion

This dissertation presents several diffusion MRI visualization and modeling methods designed to reveal the white matter structures in the brain. The thesis of this dissertation is that the diffusion imaging data can reveal white matter structures with visualization and data analysis techniques.

We used a model composed of streamtubes and streamsurfaces to visualize linear and planar structures respectively. The shape and color of the streamtubes and streamsurfaces represent implied fiber connectivity and diffusion anisotropy. To maximize the information we can show while avoiding visual clutter, we define a proximity function between the models; we also design seeding and culling strategies that keep a minimum distance between any two models.

To effectively explore the complicated models and data, we develop three different display and interaction schemes.

We place the surface models for the streamtubes, streamsurfaces and the ventricles into the CAVE and fishtank VR system and design interaction schemes respectively. A qualitative study shows that stereo helps understanding, and the users preferred the small but crisp display from the fishtank VR.

We also develop a thread-and-halo technique suitable for interactive volume rendering of thin linear structures together with a number of components for exploring

the brain models and data. Volume rendering has the defining property of mapping from the tensor field attributes to a rendered image, without introducing geometry. However, combining the volume scene with fiber tracts together creates a more informative image.

We present work toward creating color rapid-prototyping plaster models as visualization tools for the brain models. The physical representations of these virtual models allow the users to review the data with a very robust, natural and fast haptic interface: their hands.

Along with clear visual representation, we extract structures from our model to facilitate anatomical studies. We present a proximity-based algorithm sampling, culling, clustering diffusion tensor imaging (DTI) integral curves into anatomically relevant bundles. We then demonstrate the accuracy of the method to form bundles of curves that are anatomically consistent when guided and evaluated by an expert rater. We then developed an automated labeling algorithm that models the expert rater's decisions and evaluated the results.

The results of this proof-of-concept study demonstrate that anatomically-motivated methods for integral curve generation and clustering provide can reliably identify large fibers bundles that bear close qualitative correspondence with known white matter anatomy across subjects. The results for these large tracts are good both when the algorithm is guided with interactive input from an expert rater and when it is implemented automatically.

Apart from qualitative data exploration, quantitative methods are also important in revealing information from high-dimensional data sets.

We present new quantitative tractography-based metrics for assessing group differences in cerebral white matter health. We also report on initial validation of the new metrics in elderly healthy controls and in non-demented patients with vascular cognitive impairment. The study demonstrates that metrics can be derived from DTI

tractography data that provide clinically-relevant information about the microstructural integrity of cerebral white matter. The validity of our seven metrics is evidenced by their significantly different values in patients with known vascular white matter injury vs. controls and their correlations with age, and cognitive test performance. Our metrics can be obtained for specific TOIs in individual datasets for group-wise or individual comparisons.

The lack of quantitative evaluation on existing global tractography models has limited the medical and clinical applications of the diffusion MRI. To build a quantitative model that is consistent with the data, we pose an inverse problem for modeling diffusion MRI on the whole brain. This includes strategies for defining a parametric model of the brain, calculating synthetic images from that model and optimizing the parameters. We build a forward model that generates DWIs from the DTI integral curves based on multi-tensor modeling. We employ the sum of the difference between the simulated DWIs and the acquired DWIs as the goal function and optimize the placement of the DTI integral curves with a greedy algorithm and a simulated annealing approach. The results show that with the same number of curves, the optimized set of DTI integral curves fit better to the data than randomly seeded integral curves. With the proposed moves of random addition, random removal and their combination, the greedy algorithm converges faster and does not seem to get trapped in local minimums.

We work on applications that utilize both qualitative exploration and quantitative evaluation of the diffusion MRI data. With imaging on multiple sclerosis patients, we were able to identify fiber at risk by seeding the streamtubes in the lesion area. Tracing these streamtubes to midbrain reveals the fibers in corpus callosum affected by the MS lesion. For tumor patients, we observe the disparity in tumor versus non-tumor side of the brain with the increased presence of streamsurfaces around the tumor edema. Then we quantify this difference with a barycentric map for the diffusion anisotropy.

Diffusion MRI acquisition, visualization, analysis, and application have spurred numerous multidisciplinary efforts. The problem of modeling and visualization is especially intriguing because diffusion MRI data sets are large, three-dimensional and multivariate. The methods prototyped in this dissertation will facilitate more applications of diffusion MRI and help people realize its great potential for revealing the brain white matter structures.

Bibliography

- Ahrens, E. T., J.M.Allman, E.Bush, Laidlaw, D. H., and Zhang, S. (2004). Comparative 3d anatomy of the prosimian brain:DTI and histological studies. In *Proceedings of the Scientific Meeting and Exhibition of the International Society for Magnetic Resonance in Medicine*.
- Ahrens, E. T., Laidlaw, D. H., Readhead, C., Brosnan, C. F., and Fraser, S. E. (1998). MR microscopy of transgenic mice that spontaneously acquire experimental allergic encephalomyelitis. *Magnetic Resonance in Medicine*, 40(1):119–132.
- Akers, D., Sherbondy, A., Mackenzie, R., and Robert (2004). Exploration of the brain’s white matter pathways with dynamic queries. In *IEEE Visualization*, pages 377–384.
- Alexander, A. L., Hasan, K., Kindlmann, G., Parker, D. L., and Tsuruda, J. S. (2000). A geometric analysis of diffusion tensor measurements of the human brain. In *ISMRM*, pages 86–86.
- Alt, H. and Godau, M. (1995). Computing the frechet distance between two polygonal curves. *International Journal of Computational Geometry and Applications*, 5:75–91.
- Assaf, Y. and Cohen, Y. (2000). Assignment of water slow diffusing component in CNS using q-space diffusion MRS: Implications for fiber tract imaging. *Magnetic Resonance in Medicine*, 43:191–199.

- Bailey, M. J., Schulten, K., and Johnson, J. E. (1998). The use of solid physical models for the study of macromolecular assembly. *Current Opinion in Structural Biology*, 8(2):202–208.
- Banks, D. C. (1994). Illumination in diverse codimensions. In Glassner, A., editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 327–334. ACM SIGGRAPH, ACM Press. ISBN 0-89791-667-0.
- Basser, P. J., Mattiello, J., and Lebihan, D. (1994). Estimation of the effective self-diffusion tensor from the NMR spin-echo. *J. Magnetic Resonance B*, 103(3):247–254.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., and Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632.
- Beaulieu, C. (2002). The basis of anisotropic water diffusion in the nervous system - a technical review. *NMR in Biomedicine*, 15:435–455.
- Bihan, D. L. (1991). Molecular diffusion nuclear magnetic resonance imaging. *Magnetic Resonance Q.*, 7:1–.
- Bihan, D. L., Mangin, J.-F., Poupon, C., Clark, C. A., Pappata, S., Molko, N., and Chabriat, H. (2001). Diffusion tensor imaging: Concepts and applications. *JOURNAL OF MAGNETIC RESONANCE IMAGING*, 13:534–546.
- Bishop, R. and Crittenden, R. (1964). *Geometry of manifolds*. Academic Press, New York.
- Bjartmar, C., Kinkel, R., Kidd, G., Rudick, R., and Trapp, B. (2001). Axonal loss in normal-appearing white matter in a patient with acute MS. *Neurology*, 57(7):1248–52.

- Bloc, F. (1946). Nuclear induction. *Physical Review*, 70:460–474.
- Brooks, F. P. (1996). The computer scientist as a toolsmith II. *Communications of the ACM*, 39(3):61–68.
- Brun, A., Bjornemo, M., Kikinis, R., and Westin, C.-F. (2002). White matter tractography using sequential importance sampling. In *ISMRM 2002*, page 1131.
- Brun, A., Knutsson, H., Park, H.-J., Shenton, M., and Westin, C.-F. (2004). Clustering fiber traces using normalized cuts. In *MICCAI*, pages 368–375.
- Buxton, R. B. (2002). *Introduction to Functional Magnetic Resonance Imaging: Principles and Techniques*. Cambridge University Press, Cambridge, UK.
- Cabral, B., Cam, N., and Foran, J. (1994). Accelerated volume rendering and tomographic reconstruction using texture mapping hardware. In *ACM Symposium On Volume Visualization*, pages 91–98.
- Callaghan, P. T. (1993). *Principles of nuclear magnetic resonance microscopy*. Oxford, Great Britain.
- Ciccarelli, O., Parker, G. J. M., Toosy, A. T., Wheeler-Kingshott, C. A. M., Barker, G. J., Boulby, P. A., Miller, D. H., and Thompson, A. J. (2003a). From diffusion tractography to quantitative white matter tract measures: a reproducibility study. *NeuroImage*, 18:348–359.
- Ciccarelli, O., Toosy, A., Parker, G., Wheeler-Kingshott, C., Barker, G., Miller, D., and Thompson, A. (2003b). Diffusion tractography based group mapping of major white-matter pathways in the human brain. *NeuroImage*, 19:1545–1555.
- Cook, R. L. (1986). Stochastic sampling in computer graphics. *ACM Transactions on Graphics*, 5(1):51–72.

- Cormans, J., Luybaert, R., Verhelle, F., Stadnik, T., and Osteaux, M. (1994). A method for myelin fiber orientation mapping using diffusion-weighted MR images. *Magnetic Resonance Imaging*, 44:443–454.
- Corouge, I., Gouttard, S., and Gerig, G. (2004). Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *International Symposium on Biomedical Imaging*, pages 344–347.
- Cruz-Neira, C., Sandin, D. J., and DeFanti, T. A. (1993). Surround-screen projection-based virtual reality: The design and implementation of the cave. In *Proceedings of ACM SIGGRAPH*, volume 27, pages 135–142. ACM.
- Delmarcelle, T. and Hesselink, L. (1993). Visualizing second-order tensor fields with hyperstream lines. *IEEE Computer Graphics and Applications*, 13(4):25–33.
- Demiralp, C., Jackson, C., Karelitz, D., Zhang, S., and Laidlaw, D. H. (2006). Cave and fishtank virtual-reality displays: A qualitative and quantitative comparison. *IEEE Transaction on Visualization and Computer Graphics*, 12(3):323–330.
- Demiralp, C., Laidlaw, D. H., Jackson, C., Keefe, D., and Zhang, S. (2003). Subjective usefulness of CAVE and fish tank VR display systems for a scientific visualization application. In *IEEE Visualization Poster Compendium*, Seattle, WA.
- Dickinson, R. (1989). A unified approach to the design of visualization software for the analysis of field problems. In *Three-dimensional Visualization and Display Technologies (Proceedings of SPIE)*, pages 173–180.
- Ding, Z., Gore, J. C., and Anderson, A. W. (2003). Classification and quantification of neuronal fiber pathways using diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49:716–721.
- Dippé, M. A. Z. and Wold, E. H. (1985). Antialiasing through stochastic sampling.

- In Barsky, B. A., editor, *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 69–78.
- Doleisch, H., Gasser, M., and Hauser, H. (2003). Interactive feature specification for focus+context visualization of complex simulation data. In *Proceedings of the Symposium on Data Visualisation*, pages 239–248.
- Duda, R. O., Hart, P. E., and Stork, D. G. (2000). *Pattern Classification*. Wiley-Interscience Publication.
- Einstein, A. (1905). Über die von der molekularkinetischen Theorie der Wärme geforderte Bewegung von in ruhenden Flüssigkeiten suspendierten Teilchen. *Annalen der Physik*, 17:549–560.
- Enchanted Learning (2006). <http://www.zoomschool.com/subjects/anatomy/brain/neuron.shtml>.
- Evangelou, N., Konz, D., Esiri, M., Smith, S., Palace, J., and Matthews, P. (2000). Regional axonal loss in the corpus callosum correlates with cerebral white matter lesion volume and distribution in multiple sclerosis. *Brain*, 123(9):1845–9.
- Frank, L. R. (2002). Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 47:1083–1099.
- Gallyas, F. (1979). Silver staining of myelin by means of physical development. *Neurology Research*, 1(2):203–9.
- Gossl, C., Fahrmeir, L., Putz, B., Auer, L., and Auer, D. (2002). Fiber tracking from DTI using linear state space models: detectability of the pyramidal tract. *Neuroimage*, 16(2):378–88.
- Group, C. S. (2002). Mri predictors of early conversion to clinically definite MS. *Neurology*, 59(7):998–1005.

- Hahn, E. (1950). Spin echoes. *Physical Review*, 80:580–594.
- Hauser, H., Mroz, L., Bisch, G. I., and Groeller, M. E. (2001). Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 7(3):242–252.
- Hsu, E., Muzikant, A., Matulevicius, S., Penland, R., and Henriquez, C. (1998). Magnetic resonance myocardial fiber-orientation mapping with direct histological correlation. *American Journal of Physiology*, 274:1627–1634.
- Huang, H., Zhang, J., Jiang, H., Wakana, S., Poetscher, L., and Miller, M. I. (2005). DTI tractography based parcellation of white matter: Application to the mid-sagittal morphology of corpus callosum. *NeuroImage*, 26(1):195–205.
- Interrante, V. and Grosch, C. (1997). Strategies for effectively visualizing 3d flow with volume lic. In *IEEE Visualization '97*, pages 421–424.
- Jansons, K. M. and Alexander, D. (2003). Persistent angular structure: new insights from diffusion magnetic resonance imaging data. *Inverse Problems*, 19:1031–1046.
- Jenkinson, M. and Smith, S. (2001). A global optimisation method for robust affine registration of brain images. *Medical Image Analysis*, 5:143–156.
- Jiang, H., Zijl, P. V., Kim, J., Pearlson, G., and Mori, S. (2006). DTIstudio: Resource program for diffusion tensor computation and fiber bundle tracking. *Comput. Methods Programs Biomed.*, 81(2):106–116.
- Jones, D., Lythgoe, D., Horsfield, M., Simmons, A., Williams, S., and Markus, H. (1999). Characterization of white matter damage in ischemic leukoaraiosis with diffusion tensor MRI. *Stroke*, 30(2):393–397.
- Jones, J. K. (2003). Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49:7–12.

- Kajiya, J. T. and Kay, T. L. (1989). Rendering fur with three dimensional textures. In Lane, J., editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 271–280.
- Kaufman, J. A., Ahrens, E. T., Laidlaw, D. H., Zhang, S., and Allman, J. M. (2005). Anatomical analysis of an aye-aye brain (*daubentonia madagascariensis*, primates: Prosimii) combining histology, structural magnetic resonance imaging, and diffusion-tensor imaging. *Anatomical Record*, 287A(1):1026–1037.
- Kindlmann, G., Weinstein, D., and Hart, D. (1999). Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. *VIS '99: Proceedings of the conference on Visualization '99*, pages 183–189.
- Kindlmann, G. L. and Weinstein, D. M. (1999). Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization '99*, pages 183–190.
- Klingberg, T., Vaidya, C. J., Gabrieli, J. D. E., Moseley, M. E., and Hedehus, M. (1999). Myelination and organization of the frontal white matter in children: a diffusion tensor MRI study. *NeuroReport*, 10:2812–2821.
- Kniss, J., Kindlmann, G., and Hansen, C. (2002a). Multidimensional transfer functions for interactive volume rendering. *Transactions on Visualization and Computer Graphics*, 8(3):270–285.
- Kniss, J. M., Hansen, C., Grenier, M., and Robinson, T. (2002b). Volume rendering multivariate data to visualize meteorological simulations: A case study. In *Proceedings of Eurographics - IEEE TCVG Symposium on Visualization*, pages 189–194.
- Laidlaw, D. H., Ahrens, E. T., davidkremers, Avalos, M. J., Readhead, C., and Jacobs, R. E. (1998). Visualizing diffusion tensor images of the mouse spinal

- cord. In *Proceedings of IEEE Visualization 1998*, pages 127–134. IEEE Computer Society Press.
- Larsson, E. M., Englund, E., Sjobeck, M., Latt, J., and Brockstedt, S. (2004). MRI with diffusion tensor imaging post-mortem at 3.0 t in a patient with frontotemporal dementia. *Dement. Geriatr. Cogn. Disord.*, 17(4):316–319.
- Lauterbur, P. (1973). Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature*, 242:190–191.
- Lee, S., Correia, S., Tate, D., Paul, R., Zhang, S., Salloway, S., Malloy, P., and Laidlaw, D. H. (2005). Quantitative tractography metrics for white matter integrity based on diffusion tensor MRI data. *Neuroimage*. In Review.
- Lee, S., Correia, S., Tate, D., Paul, R., Zhang, S., Salloway, S., Malloy, P., and Laidlaw, D. H. (2006). Quantitative tract-of-interest metrics for white matter integrity based on diffusion tensor mri data. In *ISMRM*.
- Lengyel, J. E. (2000). Real-time fur. In *Eurographics Rendering Workshop*, pages 243–256.
- Li, G., Bordoloi, U., and Shen, H. (2003). Chameleon: An interactive texture-based rendering framework for visualizing three-dimensional vector fields. In *Proceedings IEEE Visualization*, pages 241–248.
- Lum, E. B. and Ma, K.-L. (2002). Hardware-accelerated parallel non-photorealistic volume rendering. In *Proceedings of the Non-Photorealistic Animation and Rendering Conference*, pages 67–74.
- Maddah, M., Mewes, A. U. J., Haker, S., Grimson, W. E. L., and Warfield, S. K. (2005). Automated atlas-based clustering of white matter fiber tracts from DTMRI. In *MICCAI*.

- Moberts, B., Vilanova, A., and van Wijk, J. J. (2005). Evaluation of fiber clustering methods for diffusion tensor imaging. In *IEEE Visualization*, pages 65–72.
- Mori, S., Crain, B., Chacko, V., and van Zijl, P. (1999). Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45(2):265–9.
- Mori, S., Frederiksen, K., Zijl, P. v., Stieltjes, B., Kraut, M., Solaiyappa, M., and Pomper, M. (2002). Brain white matter anatomy of tumor patients evaluated with diffusion tensor imaging. *Ann. Neurol.*, 51(3):337–8.
- Mori, S. and van Zijl, P. (2002). Fiber tracking: principles and strategies a technical review. *NMR in Biomedicine*, 15(7-8):468–480.
- Nadeau, D. and Bailey, M. (2000). Visualizing volume data using physical models. In Ertl, T., Hamann, B., and Varshney, A., editors, *Proceedings IEEE Visualization 2000*, pages 497–500. IEEE Computer Society Press. Case Studies - Physics, Information and Fields.
- Naganawa, S., Sato, K., Katagiri, T., Mimura, T., and Ishigaki, T. (2003). Regional ADC values of the normal brain: differences due to age, gender, and laterality. *European Radiology*, 13(1):6–11.
- Nakayama, N., Okumura, A., Shinoda, J., Yasokawa, Y.-T., Miwa, K., Yoshimura, S.-I., and Iwama, T. (2006). Evidence for white matter disruption in traumatic brain injury without macroscopic lesions. *J. Neurol. Neurosurg. Psychiatry*, 77:850–855.
- ODonnell, L. and Westin, C.-F. (2005). White matter tract clustering and correspondence in populations. In *MICCAI*.
- Pajevic, S. and Pierpaoli, C. (1999). Color schemes to represent the orientation of

anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42:526–540.

- Parker, G. (2000). Tracing fibre tracts using fast marching. In *ISMRM*, page 85.
- Parker, G. J. M., Stephan, K. E., Barker, G. J., Rowe, J. B., MacManus, D. G., Wheeler-Kingshott, C. A. M., Ciccarelli, O., Passingham, R. E., Spinks, R. L., Lemon, R. N., and Turner, R. (2002a). Initial demonstration of in vivo tracing of axonal projections in the macaque brain and comparison with the human brain using diffusion tensor imaging and fast marching tractography. *NeuroImage*, 15(797-809).
- Parker, G. J. M., Wheeler-Kingshott, C. A. M., and Barker, G. J. (2002b). Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE TRANSACTIONS ON MEDICAL IMAGING*, 21(5):505–512.
- Parker, S., Parker, M., and and, Y. L. (1999). Interactive ray tracing for volume visualization. *IEEE Transactions on Visualization and Computer Graphics*, 5(3):238–250.
- Paul, R. H., Laidlaw, D. H., Tate, D. F., Lee, S., Hoth, K., Gunstad, J., Zhang, S., Cohen, R. A., and Flanigan, T. (2006). Neuropsychological and neuroimaging outcome of HIV-associated progressive multifocal leukoencephalopathy in the era of ART: a case report. *HIV Medicine*. In Press.
- Pierpaoli, C. and Basser, P. J. (1996). Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance in Medicine*, 36:893–906.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, New York, second edition.

- Purcell, E., Torrey, H., and Pound, R. (1946). Resonance absorption by nuclear magnetic moments in a solid. *Physical Review*, 69:37–43.
- Scollan, D., Holmes, A., Winslow, R., and Forder, J. (1998). Histological validation of myocardial microstructure obtained from diffusion tensor magnetic resonance imaging. *American Journal of Physiology*, 275:2308–2318.
- Simon, J. H., Jacobs, L., and Kinkel, R. (2001). Transcallosal bands: a sign of neuronal tract degeneration in early MS? *Neurology*, 57(10):1888–90.
- Simon, J. H., Kinkel, R., Jacobs, L., Bub, L., and Simonian, N. (2000). A wallerian degeneration pattern in patients at risk for MS. *Neurology*, 54(5):1155–60.
- Simon, J. H., Zhang, S., Laidlaw, D. H., Miller, D. E., Brown, M., Corboy, J., and Bennett, J. (2005). Identification of fibers at risk for degeneration by diffusion tractography in patients at high risk for MS after a clinically isolated syndrome. *Journal of Magnetic Resonance Imaging*. In Press.
- Stanisz, G. J. and Henkelman, R. M. (2002). Tissue compartments, exchange and diffusion. In *Workshop on Diffusion MRI: Biophysical Issues*, pages 34–37, Saint-Malo, France.
- Stompel, A., Lum, E. B., and Ma, K.-L. (2002). Visualization of multidimensional, multivariate volume data using hardware-accelerated non-photorealistic rendering techniques. In *10 th Pacific Conference on Computer Graphics and Applications (PG'02)*, page 394.
- Trapp, B., Peterson, J., Ransohoff, R., Rudick, R., Mork, S., and Bo, L. (1998). Axonal transection in the lesions of Multiple Sclerosis. *New England Journal of Medicine*, 338:278–285.
- Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., and Wedeen, V. (2002).

- High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48(4):577–582.
- Tuch, D. S. (2002). *Diffusion MRI of complex tissue structure*. PhD thesis, Harvard.
- Upson, C., Jr., T. F., Kamins, D., Laidlaw, D. H., Schlegel, D., Vroom, J., Gurtwitz, R., and van Dam, A. (1989). The application visualization system: A computational environment for scientific visualization. *Computer Graphics and Applications*, 9(4):30–42.
- Wakana, S., Jiang, H., Nagae-Poetscher, L. M., van Zijl, P. C. M., and Mori, S. (2004). Fiber tractbased atlas of human white matter anatomy. *Radiology*, 230:77–87.
- Wei, X., Kaufman, A., and Hallman, T. (2001). Case study: Visualization of particle track data. In Ertl, T., Joy, K., and Varshney, A., editors, *Proceedings IEEE Visualization'01*, pages 466–468. IEEE Computer Society Press. Case Studies - Physical Science and Engineering.
- Weinstein, D. M., Kindlmann, G. L., and Lundberg, E. C. (1999). Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *IEEE Visualization '99*, pages 249–254.
- Wenger, A., Keefe, D., Zhang, S., and Laidlaw, D. H. (2004). Interactive volume rendering of thin thread structures within multivalued scientific datasets. *IEEE Transactions on Visualization and Computer Graphics*, 10(6):664–672.
- Westin, C., Peled, S., Gubjartsson, H., Kikinis, R., and Jolesz, F. (1997a). Geometrical diffusion measures for MRI from tensor basis analysis. In *Proceedings of ISMRM*, page 1742, Vancouver, Canada.
- Westin, C.-F., Maier, S., Mamata, H., Nabavi, A., Jolesz, F., and Kikinis, R. (2002). Processing and visualization for diffusion tensor mri. *Medical Image Analysis*, 6(93-108).

- Westin, C.-F., Peled, S., Gudbjartsson, H., and R. (1997b). Geometrical diffusion measures for MRI from tensor basis. In *ISMRM '97*, page 1742, Vancouver Canada.
- Wieshmann, U., Symms, M., Parker, G., Clark, C., Lemieux, L., Barker, G., and Shorvon, S. (2000). Diffusion tensor imaging demonstrates deviation of fibres in normal appearing white matter adjacent to a brain tumour. *J Neurol Neurosurg Psychiatry*, 68(4):501–3.
- Witwer, B., Moftakhar, R., Hasan, K., Deshmukh, P., Haughton, V., Field, A., Arfanakis, K., Noyes, J., Moritz, C., Meyerand, M., Rowley, H., Alexander, A., and Badie, B. (2002). Diffusion-tensor imaging of white matter tracts in patients with cerebral neoplasm. *J Neurosurg*, 97(3):568–75.
- Wu, Y., Bhatia, V., Lauer, H., and Seiler, L. (2003). Shear-image order ray casting volume rendering. In *Proceedings of the 2003 symposium on Interactive 3D graphics*, pages 152–162.
- Wu, Y.-C., Field, A. S., Chung, M. K., Badie, B., and Alexander, A. L. (2004). Quantitative analysis of diffusion tensor orientation: Theoretical framework. *Magnetic Resonance in Medicine*, 52:1146–1155.
- Xue, R., van Zijl, P. C., Crain, B. J., Solaiyappan, M., and Mori, S. (1999). In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magnetic Resonance in Medicine*, 42:1123–1127.
- Zhang, S., Bastin, M., Laidlaw, D. H., Sinha, S., Armitage, P. A., and Deisboeck, T. S. (2004). Visualization and analysis of white matter structural asymmetry in diffusion tensor MR imaging data. *Magnetic Resonance in Medicine*, 51(1):140–147.
- Zhang, S., Curry, C. T., Morris, D. S., and Laidlaw, D. H. (2000). Streamtubes and

- streamsurfaces for visualizing diffusion tensor MRI volume images. In *Visualization '00 Work in Progress*.
- Zhang, S., Demiralp, C., Keefe, D., DaSilva, M., Greenberg, B. D., Basser, P. J., Pierpaoli, C., Chiocca, E. A., Deisboeck, T. S., and Laidlaw, D. H. (2001a). An immersive virtual environment for DT-MRI volume visualization applications: a case study. In *Proceedings of IEEE Visualization 2001*, pages 437–440.
- Zhang, S., Demiralp, C., and Laidlaw, D. H. (2003). Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE Transactions on Visualization and Computer Graphics*, 9(4):454–462.
- Zhang, S. and Laidlaw, D. H. (2001). Elucidating neural structure in diffusion tensor MRI volumes using streamtubes and streamsurfaces. In *Proc. 9th International Society of MR in Medicine*.
- Zhang, S. and Laidlaw, D. H. (2004). DTI fiber clustering in the whole brain. IEEE Visualization 2004 Poster Compendium.
- Zhang, S. and Laidlaw, D. H. (2005). DTI fiber clustering and cross-subject cluster analysis. In *Proceedings of ISMRM*, Miami, FL.
- Zhang, S. and Laidlaw, D. H. (2006). Sampling DTI integral curves based on DWI forward modeling. In *IEEE EMBC*, New York, NY. In Press.
- Zhang, Y., Brady, M., and Smith, S. (2001b). Segmentation of brain mr images through a hidden markov random field model and the expectation maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57.
- Zöckler, M., Stalling, D., and Hege, H. (1996). Interactive visualization of 3D-Vector fields using illuminated streamlines. In *Proceedings IEEE Visualization*, pages 107–113.