

Exploration of the Brain's White Matter Pathways with Dynamic Queries

David Akers*

Anthony Sherbondy

Rachel Mackenzie

Robert Dougherty

Brian Wandell

Stanford University




Figure 1: **The corona radiata.** Our system uses dynamic queries to find structure in neural pathways suggested by MR tractography.

ABSTRACT

Diffusion Tensor Imaging (DTI) is a magnetic resonance imaging method that can be used to measure local information about the structure of white matter within the human brain. Combining DTI data with the computational methods of MR tractography, neuroscientists can estimate the locations and sizes of nerve bundles (white matter pathways) that course through the human brain. Neuroscientists have used visualization techniques to better understand tractography data, but they often struggle with the abundance and complexity of the pathways. In this paper, we describe a novel set of interaction techniques that make it easier to explore and interpret such pathways. Specifically, our application allows neuroscientists to place and interactively manipulate box-shaped regions (or volumes of interest) to selectively display pathways that pass through specific anatomical areas. A simple and flexible query language allows for arbitrary combinations of these queries using Boolean logic operators. Queries can be further restricted by numerical path properties such as length, mean fractional anisotropy, and mean curvature. By precomputing the pathways and their statistical properties, we obtain the speed necessary for interactive question-and-answer sessions with brain researchers. We survey some questions that researchers have been asking about tractography data and show how our system can be used to answer these questions efficiently.

*dakers|sherbond|rachelm|bobd|wandell@stanford.edu

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction Techniques; I.3.8 [Computer Graphics]: Applications

Keywords: Visualization, DTI, MR Tractography

1 INTRODUCTION

The brain is a massively interconnected organ. Individual neurons in the cortex typically connect to between 1,000 and 10,000 nearby neurons within the gray matter. The entire central core of the brain, known as the white matter, comprises relatively large fiber tracts that mediate communication between neurons at widely separated locations. Until recently, scientists have had limited ability to measure these white matter connections in human brains.

Knowledge about these white matter connections should enhance our understanding of normal brain function. Such knowledge should also help diagnose certain pathological disorders in patients. For example, recent research has found white matter pathway syndromes related to language deficits [13, 15, 7]. Furthermore, an understanding of white matter structure could help surgeons to avoid damaging important pathways.

Motivated by such concerns, a new technology called Diffusion Tensor Imaging (DTI) has emerged, providing a non-invasive way to measure properties of white matter pathways. Based on magnetic resonance imaging, DTI estimates the random diffusion of water molecules within biological tissue. It is widely believed that water diffuses fastest along the length of axons (rather than across their boundaries), which suggests that the principle direction of diffusion can be used to approximate the local orientation of nerve fiber bundles.

The inherent complexity of the diffusion data has motivated a variety of visualization algorithms designed to assist the researcher in analysis. One class of techniques known as MR tractography seeks to trace the principal direction of diffusion through the tensor field, connecting points together into pathways (also referred to in other literature as “fiber tracts”). As a visual representation, MR tractography is well-suited to the problem of determining white matter structure, since it implies possible anatomical connections between the endpoints of the pathways.

The pathways produced by tractography do not represent individual nerve fibers, nor do they represent bundles of these fibers. Rather, these pathways are abstract representations of possible routes through the white matter of the brain. While tractography algorithms typically produce tens of thousands of pathways, neuroscientists now believe that there are tens of millions of white matter nerve fibers grouped into hundreds of major fiber tracts. Nevertheless, the tractography estimates do have the potential to suggest real neural connections, especially when there are additional data to corroborate these estimates (e.g. from post-mortem dissections, animal studies, functional magnetic resonance imaging, etc.)

Our key contribution is a new interaction technique to assist in the exploration and identification of the pathways suggested by MR tractography. We precompute the pathways and their statistical properties and query the resulting database on-the-fly, allowing for easy exploration of tractography results using a direct-manipulation interface. We enable the specification and interactive manipulation of box-shaped regions of interest within the brain, making it possible to selectively display pathways that pass through specific anatomical regions. Querying by other path properties such as length and average curvature allows the researcher to further limit the data displayed simultaneously, making results more comprehensible. This dynamic query approach enables researchers to answer specific questions about brain connectivity with far less time or effort than is required by existing approaches.

2 RELATED WORK

A variety of techniques have emerged for the visualization of diffusion tensor data. Methods based on visual abstractions of the tensors have been used effectively to convey information about tensors at local scales within the volume (see Westin et al. [26] for a summary). Direct volume rendering techniques [11, 12] provide views of the larger trends in the data. These methods are not designed to extract or visualize estimated white matter pathways.

More relevant to our purpose of estimating white matter connectivity are MR tractography techniques that attempt to trace white matter pathways from DTI data. Streamline tracking (STT) techniques trace pathways by following the principle direction of diffusion [16, 5, 3]. Mori et al. [16] developed the FACT algorithm, a variable-step STT method that can change directions at the boundary of each voxel. Conturo et al. [5] used a constant step size, while Bassler et al. [3] suggested dynamically adjusting the step size to account for pathway curvature. Lazar et al. [14] described the tensor-deflection algorithm (TEND) based on previous work by Weinstein [25]. TEND may provide more accuracy in reconstructing certain anatomical features. Poupon et al. [19] developed a regularization technique for improved tracking, and suggested ways to model branching of nerve fiber bundles. Zhukov and Barr [28] have used regularization based on assumptions of anatomical smoothness to extract pathways in the presence of noisy data.

Many of these techniques have been criticized for their inability to handle branching or represent uncertainty [2], but they have been shown to be capable of recovering basic anatomical structures [28]. Zhang et al. [27] render the resulting pathways as streamtubes, where the cross-section of the streamtube is determined by the two smaller eigenvectors of the diffusion tensor. da Silva et al. [6] use

streamtubes to visualize differences between diffusion tensor data sets, comparing both tractography algorithms and data sets from multiple subjects.

Several groups have pointed out the potential value of filtering MR tractography data, both for rendering efficiency and simplicity of display. Zhang et al. [27] pre-filter streamtubes based on length, average linear anisotropy, and distance separating neighboring streamtubes. Conturo et al. [5] use volumetric regions of interest to select pathways that connect anatomically or functionally defined regions. Wakana et al. [24] have combined region-of-interest filters with AND, OR, and NOT operations to isolate particular neurological pathways. All three groups filter streamtubes as a pre-processing step; unlike the present application, those applications do not describe an interactive filtering technique.

While there has been significant progress on DTI visualization algorithms, surprisingly little has been written about interaction techniques. Zhang et al. [27] have been displaying streamtubes in a CAVE environment to explore the possibility that virtual reality will help doctors to make diagnoses.

Our interactive software is based on the principles of direct manipulation [22, 10] and dynamic queries [1]. An important motivation for our technique has been the development of recent methods for visual query and analysis. Hochheiser and Shneiderman [9] showed the power and simplicity of a visual query approach for answering specific questions about time-series data. By defining and manipulating rectangular regions of interest within the data set, a researcher could select quantities (e.g. stock prices) that followed certain patterns of behavior over time. Our software can be seen as an application of their 2-D spatio-temporal method to the 3-D spatial domain of the brain.

3 PREPROCESSING

3.1 Acquiring and Processing the Data

All DTI data were acquired from a neurologically normal male human subject aged 35. The DTI protocol involved eight three-minute whole-brain scans, averaged to improve signal quality. We acquired 38 axial slices for two b-values, $b = 0$ and $b = 800 \text{ s/mm}^2$ along 12 different diffusion directions. We used a 1.5T GE Signa LX and a diffusion-weighted single-shot spin echo, echo planar imaging sequence with $2 \times 2 \times 3 \text{ mm}$ voxel size (TE = 63 ms; TR = 12 s; NEX = 1; flip angle = 90° ; FOV = 260 mm; matrix size = 128×128 , bandwidth = $\pm 110 \text{ kHz}$, partial k-space acquisition).

We also collected high-resolution T1-weighted anatomical images using a sagittal 3D-SPGR sequence ($1 \times 1 \times 1 \text{ mm}$ voxel size). The T1 images were used to confirm the locations of the DTI measurements with respect to the anterior commissure (AC), posterior commissure (PC), and the mid-sagittal plane. With these landmarks, we computed a rigid-body transform from the native image space to the conventional AC-PC aligned space. The DTI data were then resampled to 2 mm isotropic voxels using a spline-based tensor interpolation algorithm [17], taking care to rotate the tensors to preserve their orientation with respect to the anatomy. We confirmed that this co-registration technique aligns the DTI and T1 images to within a few millimeters (except in regions prone to susceptibility artifacts, such as orbito-frontal and anterior temporal regions). The registration process took about twenty minutes.

We precomputed fractional anisotropy values for each diffusion tensor. Fractional anisotropy (FA) is derived from the normalized variance of the eigenvalues of each diffusion tensor [18]. FA is a scalar value that summarizes the anisotropy of the ellipsoid representation for diffusion. An FA of zero indicates spherical diffusion, as is found in the gray matter. As FA increases, the diffusion becomes more anisotropic. FA values near 0.5 indicate either linear (cigar-shaped) or planar (pancake-shaped) ellipsoids, as are

typically found in the white matter. As FA approaches 1, the diffusion becomes increasingly linear, indicated by long and thin ellipsoids. We use the precomputed FA values to establish termination criteria for path tracing algorithms (Section 3.2), to calculate the average FA along pathways for query purposes (Section 3.3), and in our interactive application to aid in navigation (Section 4.1). Our decision to use FA was motivated by its widespread adoption in the literature; there is reason to consider alternatives if the goal is to develop new tractography algorithms. See Westin et al. [26] for a good discussion of anisotropy measures and their uses in DTI.

3.2 Precomputing Pathways

Most existing tractography software traces pathways during interaction: the user selects a region of interest and the software traces pathways from seed points within this region. This approach has the disadvantage that path tracing can be time consuming, leading to frustrating delays during interaction. Instead, our approach is to precompute pathways that cover the entire white matter region of the brain, then use our software interface to efficiently “prune” these pathways to answer specific questions. Accordingly, we initialized seed points for path tracing at every other voxel in each dimension, evenly sampling the volume with seed points. (A similar seeding approach was described by Conturo et al. [5].) This sampling strategy ensured that each white matter region would have at least some pathways passing through it. However, because the pathway shapes cannot be predicted at seeding time, some regions contain more pathways than others. In the future we may explore other seeding methods that discard pathways that are too closely packed, as suggested by Zhang et al. [27] and Vilanova et al. [23].

We generated our pathways using two standard tractography methods. We chose these two algorithms because they are simple and have already been compared in the literature [14]:

- **STT:** This method follows the principal diffusion direction throughout the volume. We used a constant step size of 2 mm, an FA termination threshold of 0.15, and an angular threshold of 90°. The paths generated by STT often take sharp turns because they always follow the largest magnitude eigenvector, even in regions where the two or three largest eigenvalues are nearly identical.
- **TEND:** This method uses the tensor at each point to multiply the incoming path vector, resulting in a new vector that is deflected toward the principal direction of diffusion [14]. As with STT, we used a constant step size of 2 mm, an FA termination threshold of 0.15, and an angular threshold of 90°. The paths generated by TEND are relatively straight, since TEND avoids sharp turns when it encounters regions of low anisotropy.

To interpolate between tensors during tracing, we used a simple linear interpolation approach [28]. After thresholding by FA, our precomputation process produced about 26,000 pathways, including about 13,000 from each algorithm. All 26,000 pathways were computed in about five minutes on an Intel Pentium 1.6 GHz PC.

3.3 Precomputing Statistical Properties of Pathways

Besides precomputing the pathways, the system also precomputes statistics and other aggregate path information that can be used to specify queries. The statistical criteria we have chosen are meant as examples, and by no means represent an exhaustive set. Currently, we calculate and store the following properties for each pathway:

- **Length:** Longer paths are less likely to represent real anatomical connections, since error accumulates during path tracing. Additionally, very short paths are often distracting.

- **Average Fractional Anisotropy:** Pathways that pass through areas of low FA may be less likely to represent physical connections. (In these nearly isotropic regions, tractography algorithms differ greatly in how to proceed with path tracing.)
- **Average Curvature:** Pathways that make sharp turns are often suspect and may represent incorrect tracings. Neuroscientists often have prior knowledge about the shapes of pathways, and can use this property to remove pathways that do not follow expected shapes. Curvature is computed for each set of three consecutive points along the path, by using Heron’s formula to find the osculating circle, then computing the reciprocal of its radius.
- **Tractography Algorithm:** For later querying, the system tags each pathway with the algorithm used to generate it (STT or TEND). Querying by the algorithm allows the user of our application to compare the results of several tractography algorithms, as described in Section 5.

4 THE DYNAMIC QUERY APPLICATION

This section describes the user interface to the interactive application we have developed. The main purpose of our application is exploratory data visualization: we want to make it easier for neuroscientists to understand the neural pathways suggested by MR tractography algorithms. Figure 2 shows a labeled screen-shot of the application. With our direct-manipulation interface, it is possible to identify and display pathways that satisfy statistical constraints, or that pass through specific volumes of interest (VOIs). The interface consists of three components: The VOI panel (bottom right) allows the investigator to specify box-shaped regions for use in queries. The query panel (bottom left) provides mechanisms to query the pathways based on intersections with VOIs and statistical properties. The scene window (top) displays the currently selected pathways and assists in navigating the volumetric data space.

We explain the use of this interface in Sections 4.1 and 4.2. Please also see the included video footage for examples of its use.

4.1 Specifying Dynamic Queries

Before querying the data, an investigator must be able to navigate the volumetric data space represented in the scene window. To begin with, the investigator can change the camera position and orientation using a standard trackball/mouse interface. As a further aid to navigation, the scene window provides three moveable cutting planes (tomograms), which display planar reformations of FA data. Features visible in FA are commonly used by neuroscientists to navigate the brain’s white matter structures.

Often a query sequence begins with the selection of a set of desired pathways based on the statistical criteria described in Section 3.3. A set of slider bars in the query panel allows for the interactive specification of a range (min, max) of acceptable values. As the investigator drags any slider bar, the matching pathways are found and displayed in the scene window.

A key feature of our application involves the use of VOIs to display pathways that pass through specific anatomical regions. Once specified, VOIs can be used to form queries by setting the VOI query expression in the query panel. VOIs can be combined using simple AND and OR operations, or by typing in an arbitrary Boolean logic expression. The VOI editing panel (Figure 2, lower right) allows for the exact specification of VOI dimensions and position. As the VOI is modified using the slider bars or text widgets, the query is re-evaluated immediately and the scene window is updated with new pathway information. A VOI can be controlled more directly in the scene window, by using the mouse to click and drag the VOI. The investigator simply selects a tomogram and then

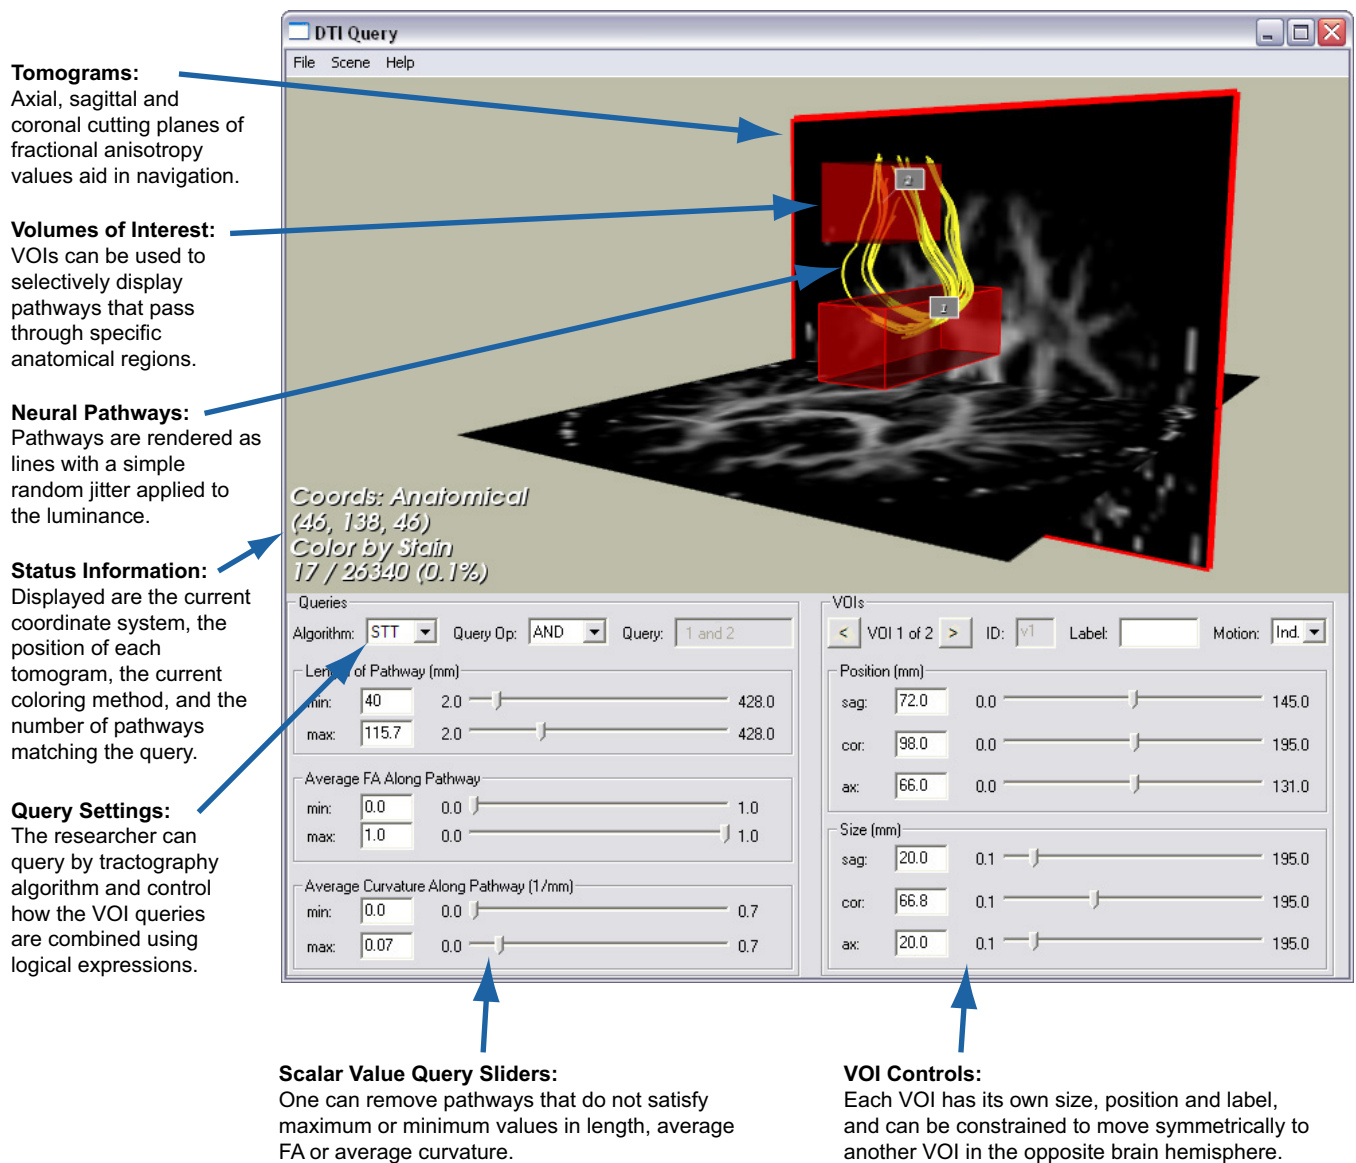


Figure 2: **The user interface to our pathway exploration software.** The interface consists of three components: The VOI panel (bottom right) allows the investigator to specify box-shaped regions for use in queries. The query panel (bottom left) provides mechanisms to query the pathways based on statistical properties and intersections with VOIs. The scene window (top) displays the currently selected pathways and assists in navigating the volumetric data space. Neuroscientists use the VOI and query control panels to identify specific neural pathways, which are then displayed in the scene window above.

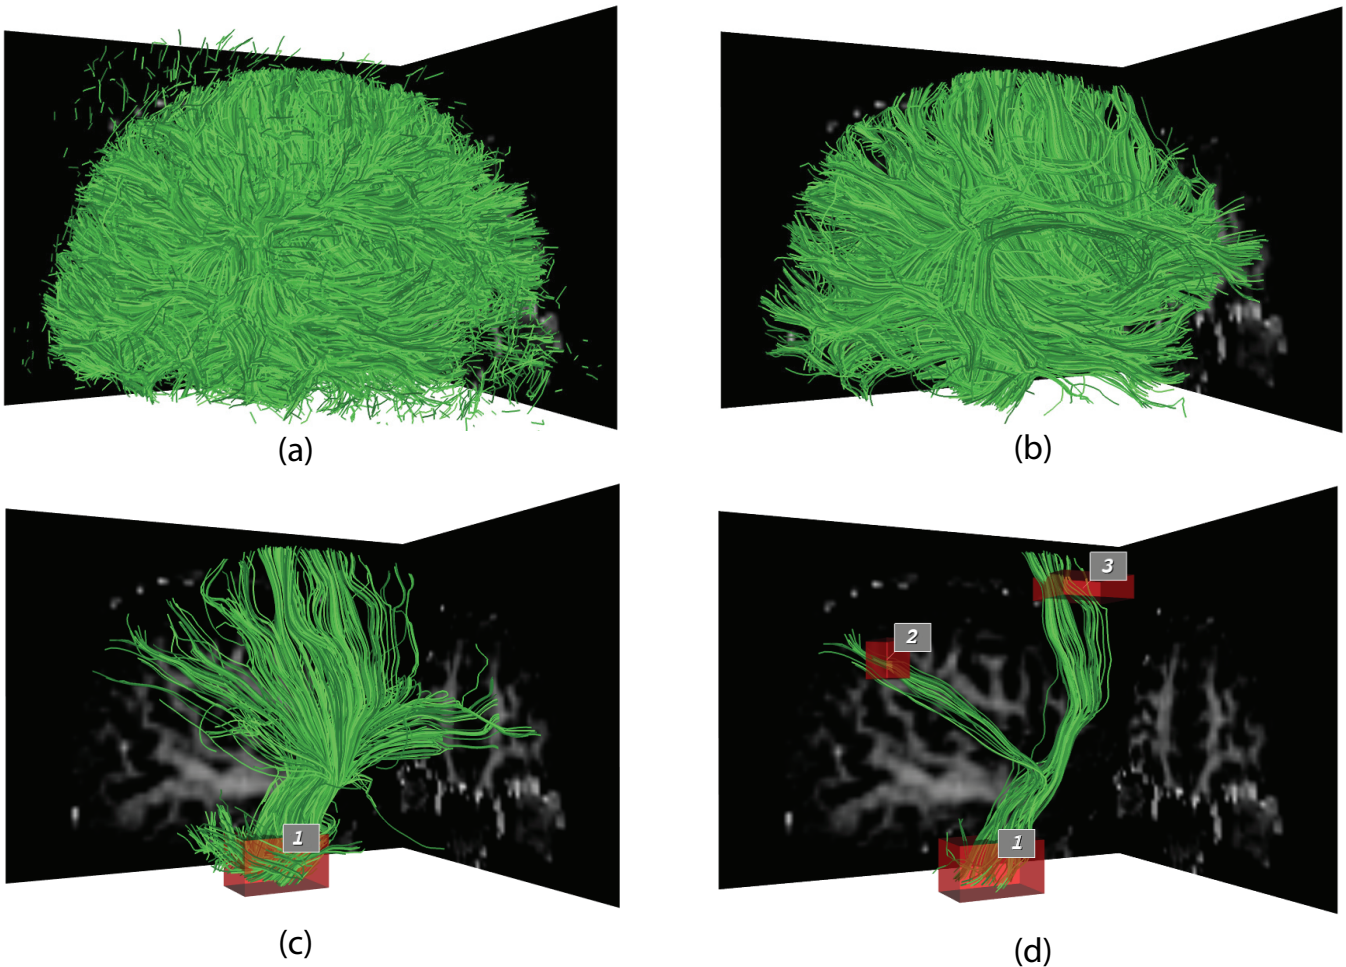


Figure 3: **Sequence of dynamic queries identifying the spatial organization of fiber pathways.** a) All 13,000 pathways computed using the STT algorithm. Patterns are difficult to discern because of visual clutter. b) Using a length filter, we show only the pathways that are greater than 4 cm in length (about 30% of the total number of pathways). c) By placing VOI 1 in the scene, we show only the pathways that pass through the internal capsule (bottom). d) By placing VOIs 2 and 3, we obtain a picture showing connections between 1 and either 2 or 3.

drags the VOI to any position on the plane. One can also link two VOIs to move them symmetrically in opposite brain hemispheres. This was made possible by aligning the data to AC-PC space, which defines the plane halfway between the hemispheres.

4.2 Pathway Rendering

While others have used streamtubes to represent pathways [27], we have chosen simply to use lines. Lines can be drawn much faster than streamtubes, and they adequately represent connectivity information of interest to neuroscientists (at the cost of losing local information about the underlying tensor data along the path). To visually distinguish the pathways, we use a simple ‘color jittering’ technique. In HSV color space, we compute a random luminance offset for each pathway. The luminance of each pathway is determined at startup and is held constant to avoid shimmering artifacts.

Differences in hue are used to establish logical groupings of pathways, using a process we call ‘virtual staining’. Here, the investigator can choose a hue and use it to color all of the pathways currently displayed in the scene window. This allows investigators to identify specific pathways and then visualize them within their surrounding context: as the query is modified, the original pathways remain stained. Virtual staining was used extensively in generating the results shown in Section 5, and in generating Figure 1.

4.3 Implementation

This section describes the implementation details of our interactive application. The program was written entirely in C++, and was designed to work on any modern inexpensive PC without any special hardware requirements. The program makes use of the visualization toolkit (VTK) [21] for 3D scene generation and interaction.

Since data exploration naturally involves making iterative adjustments to queries, our main goal has been to make the system immediately responsive when the investigator changes a query. One key to this interactivity is the preprocessing described in the previous section, but this preprocessing is not enough by itself to make our system interactive. At runtime, we also need to be able to interactively compute intersections between VOIs and pathways. To facilitate fast intersection tests, our program stores each pathway’s geometry as a hierarchical oriented bounding box (or OBBTree). For this we used the freely-available RAPID software from the University of North Carolina [8]. All VOIs and pathways are represented as sets of triangles that can be efficiently tested for intersection. The box-shaped VOIs are trivial to triangulate, and the pathways are triangulated as very small area (long and thin) triangles. Since the RAPID software only reports object intersections between triangles, our application also tests the endpoints of each pathway to determine whether the pathway is fully contained within the VOI.

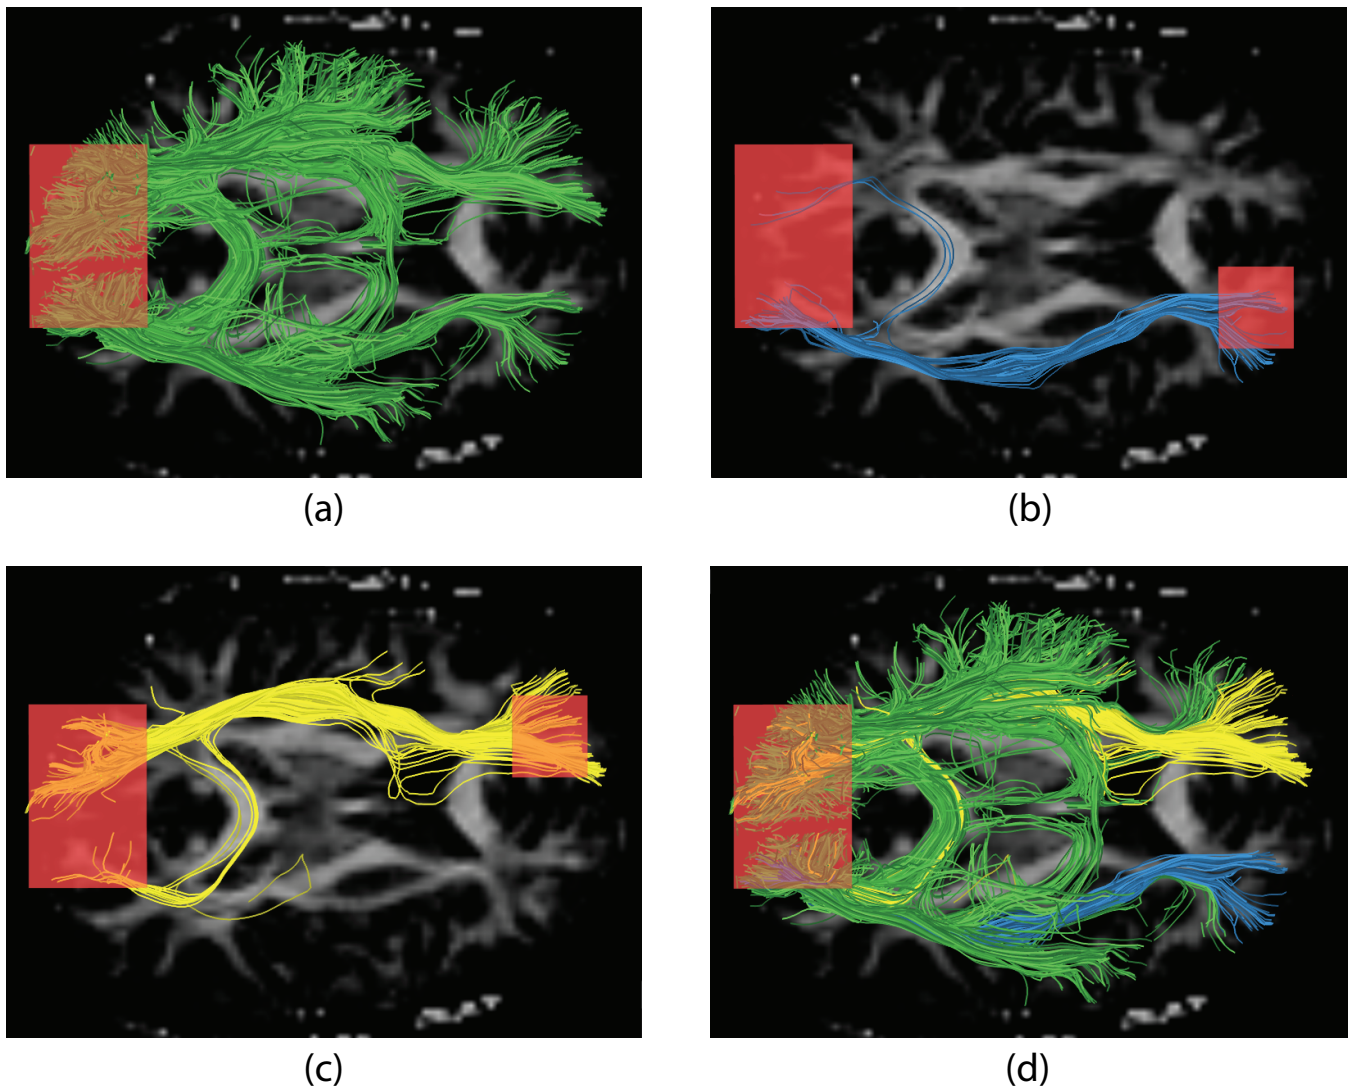


Figure 4: **Validation of known white matter pathways.** The left and right inferior longitudinal fasciculi are known anatomical pathways that connect the occipital and frontal lobes in each hemisphere. In four simple steps, a researcher used the system to produce a visual representation of these known pathways. a) Placing a VOI covering the occipital lobes shows all pathways involved in human visual processing. b) Placing a second VOI in the right frontal lobe and using an AND operation shows only the pathways in the right hemisphere, which are stained blue. c) Moving the second VOI to a symmetric position on the left captures the pathways in the left hemisphere, which are stained yellow. d) Removing the second VOI from the query shows the results of b and c within the original set of pathways. The pathways shown were all produced using the STT algorithm and were also limited to lengths greater than 4 *cm* to reduce visual clutter.

Queries based on precomputed pathway properties are very fast since the precomputed values need only be compared against the current range of the query. The performance of RAPID is described by Gottschalk et al. [8], but their execution times are based on an older SGI Reality Engine. In our own benchmarks on a 1.6 *GHz* Pentium laptop PC, we are able to intersect a VOI with between 80,000 and 220,000 pathways per second, depending on the size of the VOI. (Larger VOIs require more bounding-box tests, since they intersect with more of the pathways.) This allows us to maintain a frame rate of 3-8 fps while manipulating the VOIs. While not the most efficient solution for intersection with box-shaped VOIs, using RAPID will allow us to implement more complex (e.g. non-convex) VOI shapes in the future, without a major change in performance.

On average, each pathway consumes approximately 20 *KB* of memory, including the OBBTree structure and the points used for rendering. Accordingly, we use 510 *MB* of memory to represent all 26,000 pathways.

5 RESULTS

In this section we demonstrate some of the capabilities of our system. First, we acquired a DTI data set collected from a single normal subject (described in Section 3.1). Using this data set as input to our system, we identified three types of queries that are especially useful to neuroscientists. In particular, we will show how our system has been used to validate known white matter pathways, to explore previously unidentified pathways, and to visually compare tractography algorithms. All three examples were produced by a novice user of our system who is a neuroscientist specializing in brain imaging.

Using our dynamic query system, the neuroscientist easily identified two known neural pathways in the test subject data, the left and right Inferior Longitudinal Fasciculi (ILF). Shown in Figure 4, these pathways connect the occipital and frontal lobes in each




Figure 5: **Exploration of candidate neural pathways.** The query results show many pathways reflecting the well-known connection between the left and right occipital lobes, passing through the splenium. Intriguingly, the query also suggests a possible connection through the anterior commissure. This suggested pathway may incite further validation research. Our system helps to form such hypotheses by allowing researchers to interactively pose and answer specific questions about connectivity. The pathways shown were all produced using the STT algorithm and were also limited to lengths greater than 4 cm to reduce visual clutter.




Figure 6: **Visual comparison of tractography algorithms.** The VOI in this query is placed within the corpus callosum. The pathways were computed using the STT (blue) and TEND (yellow) algorithms. Using virtual staining, a neuroscientist was able to easily inspect the differences by overlaying the pathways. While both TEND and STT show callosal projections to superior regions, the TEND pathways also include callosal projections to both temporal lobes. This example shows the extent to which the two algorithms can differ. The pathways shown were limited to lengths greater than 4 cm to reduce visual clutter.

hemisphere of the brain.¹ To locate the pathways, first our neuroscientist test subject placed a single VOI covering both occipital lobes, revealing all neural pathways involved in visual processing. Next, he placed an additional VOI in the right frontal lobe above the eye and used an AND operation to show pathways connecting the two brain regions. The many pathways passing between these VOIs comprise the right ILF. Interested to see whether these neural pathways were located symmetrically on both sides of the brain, the researcher moved the second VOI to a symmetric position in the left hemisphere, identifying the left ILF. Finally, using virtual staining, the neuroscientist separately colored the pathways from each hemisphere so that they could be visually compared. This exploration was performed in about five minutes.

Our system also enables exploration of novel pathways that could motivate future research. Figure 5 shows all the pathways generated by our system which pass between the occipital lobes (responsible for visual processing). To isolate these pathways, the neuroscientist placed VOIs on each of the occipital lobes and displayed the conjunction of the VOI queries. The majority of the connections follow a known neural pathway, crossing the posterior corpus callosum and terminating at a symmetric location in the opposite hemisphere. Interestingly, some of them travel forward to cross at what appears to be the anterior commissure, a small bundle of fibers connecting the two hemispheres beneath the corpus callosum. Further research is necessary to determine whether this anterior pathway is real. In situations like this, our system can help form hypotheses about novel pathways by allowing researchers to interactively pose and answer specific questions about connectivity. This exploration was performed in about five minutes.

Finally, our system can be used to visually compare the pathways estimated by different tractography algorithms. Figure 6 shows pathways generated by the STT and TEND algorithms (described in Section 3.2). By virtually staining pathways passing through a region in the corpus callosum, the neuroscientist was able to visualize an important difference between the two algorithms. The STT algorithm, following the direction of greatest diffusion at each point, generates only 'U-shaped' pathways. The TEND algorithm additionally produces pathways passing from the corpus callosum to each temporal lobe. As this example illustrates, the pathways generated by these algorithms can differ greatly. Such visualizations are useful to the neuroscientist who is uncertain about the reliability of estimates across algorithms, and to the expert in tractography who wants to understand the consequences of algorithmic assumptions. This exploration was performed in about ten minutes.

6 DISCUSSION

The key to our system's utility is its ability to render data and respond to queries at interactive rates. Our colleagues in psychology stressed the significance of this program as an exploratory tool; quickly browsing through connections in the brain could be invaluable in identifying areas of interest for future study. The system could also assist scientists investigating the neurological bases of disorders, as has been done with other methods for analyzing DTI data [15, 7, 13], or provide a diagnostic tool for such disorders. It could be employed as an educational aid for students learning about neuroanatomy, as it allows for interactive viewing of the primary anatomical pathways.

Our system may also be useful for exploring data from sources other than DTI. Saleem et al. [20] have developed a method of tracing axonal connections across synapses in live monkeys. $MnCl_2$ is injected into the monkey brain and transported along neuronal tracts where it can later be detected with MRI. This approach may

¹There is still some debate over the existence of the ILF; see Catani et al. [4] for a recent discussion.

generate large amounts of verifiable, high-resolution data that could be browsed efficiently with our system.

The results we have obtained with our system have been limited by the resolution of the data and the quality of the tractography algorithms. It is important to realize that these are not limitations of our interactive technique, but rather they are limitations of the current acquisition technology and algorithms. Current DTI data are highly restricted in spatial resolution. The $2 \times 2 \times 3$ mm resolution used in this scan represents the current state of the art. The voxel dimensions are roughly two orders of magnitude larger than the cross-sectional width of white matter axons (between 10 and 50 microns). This makes ambiguities in tracing inevitable when estimating neural pathways from the diffusion tensor field. However, despite limitations in data quality our system remains viable and useful for exploring MR tractography data and suggesting possible hypotheses about connectivity.

7 FUTURE WORK

The methodology of precomputation and dynamic queries should yield several interesting enhancements in the future. In particular, it will be useful to expand our system to handle multiple data sets. For instance, one could include pathways from multiple subjects for the purposes of comparing pathological cases to normal ones, or simply understanding population variance. In addition, our technique should be integrated with functional magnetic resonance imaging so that the user can simultaneously view connectivity and activation information about specific brain regions [5]. While these methods can benefit from a dynamic query approach, they will also require the development of methods to co-register the various data sets.

We also believe that useful improvements could be made in the visual representation of the pathways. Currently the pathways are drawn simply as lines; however, it might be advantageous to aggregate pathways into groups, or to simplify their paths for easier interpretation. Such abstractions could also contain visual cues that measure either local DTI properties or statistical information regarding the certainty of the path estimates.

8 ACKNOWLEDGMENTS

We would like to thank Roland Bammer and Michael Moseley of the Stanford University Lucas Center, for providing useful discussion and developing the pulse sequences used to collect our data.

REFERENCES

- [1] C. Ahlberg, C. Williamson, and B. Shneiderman. Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of SIGCHI 92*, pages 619–626. ACM Press, 1992.
- [2] R. Bammer, B. Acar, and M. E. Moseley. In vivo MR tractography using diffusion imaging. *European Journal of Radiology* 45, pages 223–234, 2002.
- [3] P. J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
- [4] M. Catani, D.K. Jones, R. Donato, and D.H. ffytche. Occipitotemporal connections in the human brain. *Brain*, 126:2093–2107, 2003.
- [5] T. E. Conturo, N. F. Lori, T. S. Cull, E. Akbudak, A. Z. Snyder, J. S. Shimony, R. C. McKinstry, H. Burton, and M. E. Raichle. Tracking neuronal fiber pathways in the living human brain. In *Proceedings of the National Academy of Sciences*, pages 10422–10427, 1999.
- [6] M.J. da Silva, S. Zhang, C. Demiralp, and D.H. Laidlaw. Visualizing diffusion tensor volume differences. In *IEEE Visualization 01 Proceedings, Work in Progress*, 2001.
- [7] G.K. Deutsch, R.F. Dougherty, R. Bammer, W.T. Siok, J.D.E. Gabrieli, and B. Wandell. Children’s reading performance is correlated with white matter structure measured by diffusion tensor imaging. *Cortex*, In press.
- [8] S. Gottschalk, M. C. Lin, and D. Manocha. OBBTree: A hierarchical structure for rapid interference detection. In *Proceedings of SIGGRAPH 96*, pages 171–180, 1996.
- [9] H. Hochheiser and B. Shneiderman. Visual queries for finding patterns in time series data. Technical report, University of Maryland, 2002.
- [10] E. Hutchins, J. Hollan, and D. Norman. Direct manipulation interfaces. In D. Norman and S. Draper, editors, *User Centered System Design*, pages 87–124. Erlbaum, 1986.
- [11] G. Kindlmann and D. Weinstein. Hue-balls and lit-tensors for direct volume rendering of diffusion tensor fields. In *IEEE Visualization 99 Proceedings*, pages 183–189. IEEE Computer Society Press, 1999.
- [12] G. Kindlmann, D. Weinstein, and D. Hart. Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 2000, pages 124–138, 2000.
- [13] T. Klingberg, M. Hedehus, E. Temple, T. Salz, J.D. Gabrieli, M.E. Moseley, and R.A. Poldrack. Microstructure of temporo-parietal white matter as a basis for reading ability: Evidence from diffusion tensor magnetic resonance imaging. *Neuron*, 25:493–500, 2000.
- [14] M. Lazar, D. M. Weinstein, J. S. Tsuruda, K. M. Hasan, K. Arfanakis, M. E. Meyerand, B. Badie, H. A. Rowley, V. Houghton, A. Field, and A. L. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18:306–321, 2003.
- [15] N. Molko, L. Cohen, J.F. Mangin, F. Chochon, S. Lehericy, D. L. Bihan, and S. Dehaene. Visualizing the neural bases of a disconnection syndrome with diffusion tensor imaging. *Journal of Cognitive Neuroscience*, pages 629–636, 2002.
- [16] S. Mori, B.J. Crain, V.P. Chacko, and P.C. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol*, 45:265–269, 1999.
- [17] S. Pajevic, A. Aldroubi, and P.J. Basser. A continuous tensor field approximation of discrete DT-MRI data for extracting microstructural and architectural features of tissue. *J Magn Reson*, 154(1):85–100, 2002.
- [18] C. Pierpaoli and P.J. Basser. Toward a quantitative assessment of diffusion anisotropy. *Magnetic Resonance Magazine*, pages 893–906, 1996.
- [19] C. Poupon, J. Mangin, C.A. Clark, V. Frouin, J. Regis, D. Le Bihan, and I. Bloch. Towards inference of human brain connectivity from MR diffusion tensor data. *Medical Image Analysis*, 5:1–15, 2001.
- [20] K. S. Saleem, J. Pauls, M. Augath, T. Trinath, B. Prause, T. Hashikawa, and N. K. Logothetis. Magnetic resonance imaging of neuronal connections in the macaque monkey. *Neuron*, 34:685–700, 2002.
- [21] W. Schroeder, K. Marin, and B. Lorensen. *The Visualization Toolkit*. Prentice Hall, New Jersey, USA, 1997.
- [22] B. Shneiderman. Direct manipulation: A step beyond programming languages. *IEEE Computer* 16(8), pages 57–69, 1983.
- [23] A. Vilanova i Bartroli, G. Berenschoot, and C. van Pul. DTI visualization with streamsurfaces and evenly-spaced volume seeding. In *VisSym '04 Symposium on Visualization*, pages 173–182, 2004.
- [24] S. Wakana, H. Jiang, L.M. Nagae-Poetscher, P.C. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, pages 77–87, 2004.
- [25] D. M. Weinstein, G. L. Kindlmann, and E. C. Lundberg. Tensorlines: Advection-diffusion based propagation through diffusion tensor fields. In *IEEE Visualization 99 Proceedings*, pages 249–254, San Francisco, 1999.
- [26] C.F. Westin, M. SE, M. Nabavi, A. Jolesz, and F. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6:93–108, 2002.
- [27] S. Zhang, C. Demiralp, and D. Laidlaw. Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE Transactions on Visualization and Computer Graphics* 2003, pages 454–462, 2003.
- [28] L. Zhukov and A. H. Barr. Oriented tensor reconstruction: Tracing neural pathways from diffusion tensor MRI. In *IEEE Visualization 02 Proceedings*, pages 387–394. IEEE Computer Society, 2002.