Exploring Connectivity of the Brain's White Matter with Dynamic Queries

Anthony Sherbondy, David Akers, Rachel Mackenzie, Robert Dougherty, and Brian Wandell

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 or ellipsoid-shaped regions to selectively display pathways that pass through specific anatomical areas. These regions can be used in coordination with a simple and flexible query language which allows for arbitrary combinations of these queries using Boolean logic operators. A representation of the cortical surface is provided for specifying queries of pathways that may be relevant to gray matter structures and for displaying activation information obtained from functional magnetic resonance imaging. 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.

Index Terms—Computer graphics interaction techniques, computer graphics applications, 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, is comprised of 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 [17], [20], [9]. 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 noninvasive 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. (See Basser et al. [5] for a discussion.)

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. This could include postmortem dissections, animal studies, or Functional Magnetic Resonance Imaging (fMRI). fMRI suggests possible connections between regions of the brain's gray matter surface based on correlations in activity [7]; it is particularly compelling since it is non-invasive, like DTI.

Our contribution (extending the work of Akers et al. [3]) is a new interaction technique to assist in the exploration

- A. Sherbondy is with the Department of Electrical Engineering, James H. Clark Center, 318 Campus Dr., Room S324, Stanford University, Stanford, CA 94305. E-mail: Sherbond@stanford.edu.
- D. Akers is with the Department of Computer Science, Gates Computer Science Building, 353 Serra Mall, Room 3B-396, Stanford University, Stanford, CA 94305. E-mail: dakers@stanford.edu.
- R. Mackenzie, R. Dougherty, and B. Wandell are with the Department of Psychology, Jordan Hall, 450 Serra Mall, Stanford University, Stanford, CA 94305. E-mail: {rachelm, bobd, wandell}@stanford.edu.

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

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 or ellipsoid-shaped regions of interest within the brain, making it possible to selectively display pathways that pass through specific anatomical regions. The researcher can also use fMRI data to search for pathways that might connect several functionally defined regions of the brain's gray matter surface. 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. The 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. [35] for a summary). Direct volume rendering techniques [15], [16], [34] 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. Streamlines Tracking Techniques (STT) trace pathways by following the principle direction of diffusion [21], [7], [5]. Mori et al. [21] developed the Fiber Assignment through Continuous Tracking (FACT) algorithm, a variable-step STT method that can change directions at the boundary of each voxel. Conturo et al. [7] used a constant step size, while Basser et al. [5] suggested dynamically adjusting the step size to account for pathway curvature. Lazar et al. [18] described the tensor-deflection

algorithm (TEND) based on previous work by Weinstein et al. [33]. TEND may provide more accuracy in reconstructing certain anatomical features. Poupon et al. [24] developed a regularization technique for improved tracking and suggested ways to model branching of nerve fiber bundles. Zhukov and Barr [37] 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 [4], but they have been shown to be capable of recovering basic anatomical structures [37]. Zhang et al. [36] 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. [8] 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. [36] prefilter streamtubes based on length, average linear anisotropy, and distance separating neighboring streamtubes. Conturo et al. [7] use volumetric regions of interest to select pathways that connect anatomically or functionally defined regions. Wakana et al. [32] have combined region-of-interest filters with AND, OR, and NOT operations to isolate particular neurological pathways. All three groups filter streamtubes as a preprocessing 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. [36] have been displaying streamtubes in a CAVE environment to explore the possibility that virtual reality will help doctors to make diagnoses. In addition, interactive volume rendering techniques have also been employed for the purposes of DTI visualization [15], [16], [34].

Our interactive software is based on the principles of direct manipulation [28], [14] and dynamic queries [2]. An important motivation for our technique has been the development of recent methods for visual query and analysis. Hochheiser and Shneiderman [13] 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 2D spatio-temporal method to the 3D 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 $b = 800$ 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 a nominal $2x2x3$ mm voxel size (63 ms TE; 12 s TR; 1 NEX; 90° flip angle; 260 mm FOV; $128x128$ matrix size; ±110 kHz bandwidth; partial k-space acquisition).

We also collected high-resolution T1-weighted anatomical images using an eight minute sagittal 3D-SPGR sequence ($1x1x1$ mm voxel size). The DTI data were coregistered to the T1 data by automatically aligning the B0 to T1 using a mutual information algorithm [1]. We confirmed that this coregistration 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 following anatomical landmarks were defined in the T1 images: the anterior commissure (AC), the posterior commissure (PC), and the mid-sagittal plane. With these landmarks, we utilized a rigid-body transform to convert both DTI and T1 data 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 [22], taking care to rotate the tensors to preserve their orientation with respect to the anatomy. The registration process took about 20 minutes.

When combined with DTI, fMRI allows the neuroscientists to simultaneously view connectivity and activation information about specific brain regions [7]. Functional MR data were acquired with a spiral pulse sequence with 21-30 obliquely oriented slices acquired every 2.4 seconds (30 ms TE; 1.2 s TR; 2 interleaves; 70° flip angle; $2x2x3$ mm effective voxel size). Each individual functional scan lasted about four minutes and the subject was given a brief break between scans. Visual field maps were measured using rotating wedge and expanding ring stimuli. A thorough description of the collection of the fMRI data and the registration of the data with the high-resolution anatomy data is described by Dougherty et al. [10]. We map the fMRI data to our cortical surface mesh, as described in Section 3.3.

We demonstrate the coordination of the fMRI data with pathway queries in an example in Section 5.4.

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 [23]. 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 3.4), and in our interactive application to aid in navigation (Section 4.2). 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. [35] 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. [7].) 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. [36] and Vilanova et al. [31].

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 [18]:

- **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 [18]. 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 [37]. 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 5 minutes on an Intel Pentium 1.6 GHz PC.

3.3   Representing the Cortical Surface

We also precomputed two mesh representations of the cortical surface to allow a neuroscientist to answer questions about pathway connectivity to the cortex. First, the white matter voxels were classified semimanually by a neuroscientist in a separate custom software platform [30]. This process can take approximately 30 minutes for the expert to complete, depending on the success of the automated portion of the segmentation.

Next, we used the Marching-Cubes algorithm [19] to obtain a triangle mesh for the interface. The resulting cortical surface mesh is complex to navigate on with its many folds and creases. To remedy this, we smoothed the mesh using a windowed sinc filter [29]. Smoothing the cortical surface was first described by Sereno et al. [27], who used the smoothed representation to display fMRI data. In our system, we also used the smoothed cortical surface to help the researcher to perform surface constrained queries, as described in Section 4.2.3. The smoothed surface and the original surface are stored as unstructured meshes of triangle strips. Both surface meshes utilize the same list of vertex indices; the smoothed surface is different from the original surface only by vertex location. For our test subject, the cortical surface representations each contain about 230,000 vertices. The surface triangulation and smoothing required approximately 2 minutes on a 1.6 GHz Pentium laptop PC.

3.4   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

connections. 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. Fig. 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 or ellipsoid-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 4.2 and 4.3. Please also see the video footage at http://graphics.stanford.edu/papers/dti-query-extended for examples of its use.

4.1   Navigating the Scene

Before querying the data, an investigator must be able to navigate the volumetric data space represented in the scene window. 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.

4.2   Specifying Dynamic Queries

There are three ways to specify queries using our interface, as described below.

4.2.1   Querying by Pathway Statistics

Often a query sequence begins with the selection of a set of desired pathways based on the statistical criteria described in Section 3.4. 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.
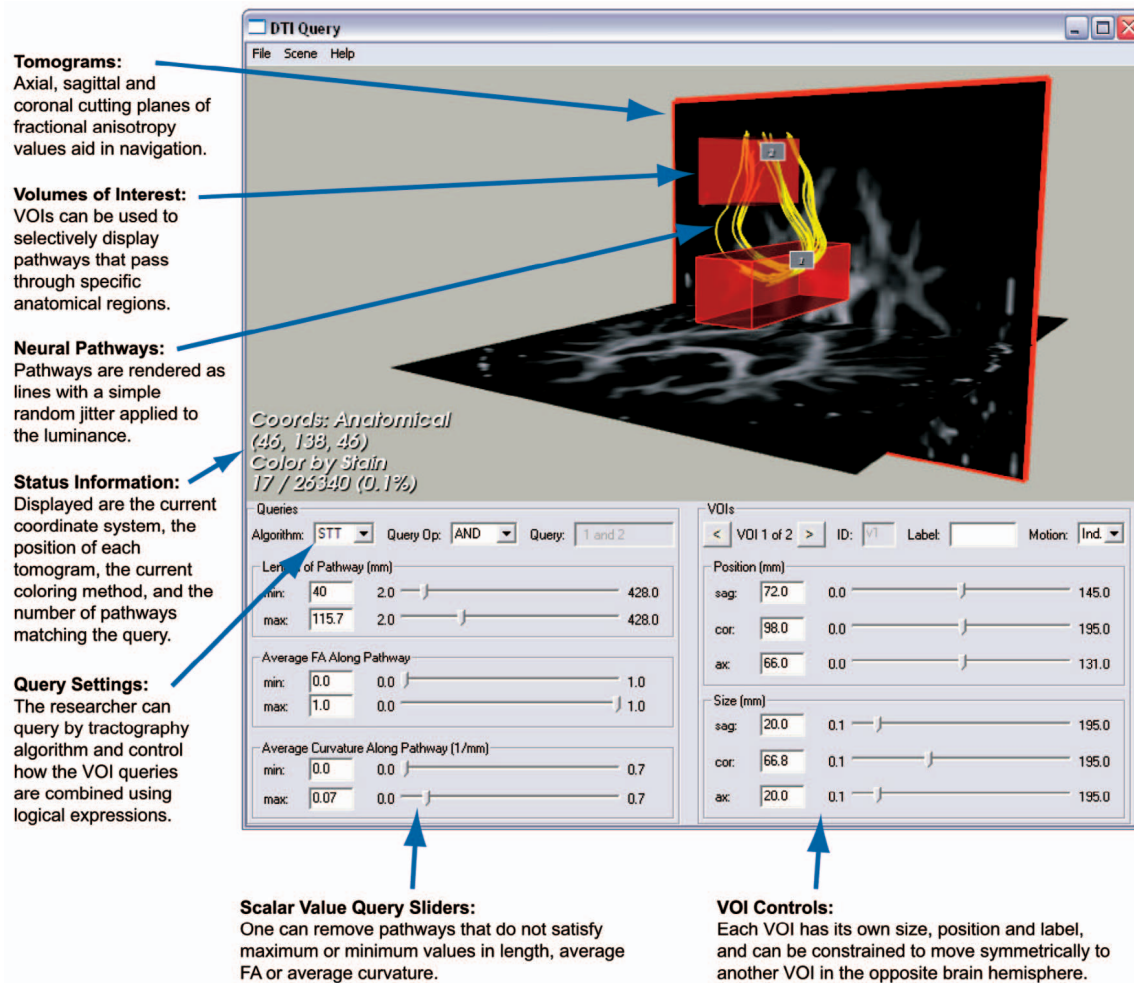
Fig. 2. The user interface to our pathway exploration software. The interface consists of three components: The VOI panel (bottom right) allows the investigator to specify box or ellipsoid-shaped regions for use in queries. The query panel (bottom left) provides mechanisms to query the pathways based 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.

4.2.2 Querying by Volumes of Interest

A key feature of our application involves the use of VOIs to display pathways that pass through specific anatomical regions (see Fig. 3). 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 an arbitrary Boolean logic expression. The VOI editing panel (Fig. 2, lower right) allows for the exact specification of VOI dimensions and position. As a VOI is modified using the slider bars or text widgets, the query is reevaluated 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 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.

4.2.3 Querying by Surface-Constrained Volumes of Interest

As shown in Fig. 4, our application also allows the neuroscientist to constrain the motion of a VOI to a representation of the cortical surface. After adding the surface representation to the scene window, the neuroscientist can simply drag a VOI along this surface using the mouse.

As described in Section 3.3, we also precomputed a smoothed version of the original surface. By dragging the VOI along the smoothed surface, the neuroscientist can easily reach regions of the brain surface that would otherwise be obscured from view by the many peaks (gyri) or valleys (sulci) of the brain. However, this approach is limited by low anisotropy near the cortical surface, which causes difficulties for path tracing. Moreover, the cortical surface is not perfectly aligned to the DTI data because of registration errors caused by the geometric distortions introduced during DTI data acquisition (as described in Section 3.1). This alignment problem is exacerbated by the mesh smoothing we perform.

Fig. 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 percent 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.

To display the VOI on the smoothed surface, the VOI is intersected with the original surface mesh and the resulting triangles are warped onto the smoothed representation using the one-to-one vertex correspondence. The warped surface patch is shaded green, as shown in Fig. 4. To accelerate the required picking operations on either surface mesh, we compute and store octree representations of the meshes.

To provide additional information to the neuroscientist, a color map can be applied to either surface representation. Coloring based on curvature helps to convey the shape of the cortical surface, particularly for the smoothed representation where the sulci and gyri have been smoothed. In Fig. 4, the surface of the smoothed and bumpy meshes both display light shades of gray for peaks and darker shades of gray for valleys. Optionally, the neuroscientist can overlay a color map of fMRI data. Coloring based on fMRI data allows the neuroscientist to drive VOI navigation with brain activation information. This was used extensively in Section 5.4 to guide VOI placement.


Fig. 4. Constraining VOI motion to the cortical surface. (a) The cortical surface with an ellipsoid-shaped VOI placed on a section of the back of the brain. (b) A smoothed representation of the cortical surface with a projection of the same VOI. Both surfaces present a higher intensity gray level for gyri and a lower intensity gray for the sulci. The smoothed surface allows the neuroscientist to manipulate a VOI on the cortical surface without the bumps and folds of the original surface obscuring the view to the VOI.

4.3 Pathway Rendering

While others have used streamtubes to represent pathways [36], 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 Fig. 1.

4.4 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) [26] 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 [12]. All VOIs and pathways are represented as sets of triangles that can be efficiently tested for intersection. The box or ellipsoid-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.

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. [12], 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 our current ellipsoid and box VOIs, using RAPID will allow us to implement more complex (e.g., nonconvex) VOI shapes in the future without a major change in performance.

On average, each pathway consumes approximately 20KB of memory, including the OBBTree structure and the points used for rendering. Accordingly, we use 510MB of memory to represent all $26,000$ pathways. The cortical surface meshes and all their accompanying data structures currently occupy about 160MB of memory. This memory is consumed by the original surface mesh, the smoothed surface mesh, and the octree acceleration data structure.

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. We will then show results of our system in a more complex neurobiological investigation of the pathways between the left and right hemispheres of the retinotopy map on the cortical surface of the brain. All four examples were produced by a novice user of our system who is a neuroscientist specializing in brain imaging.

5.1 Validating Known Pathways

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 Fig. 5, these pathways connect the occipital and frontal lobes in each hemisphere of the brain.[1] To locate the pathways, first our neuroscientist test subject placed a single VOI covering both occipital lobes, revealing neural pathways produced by our system which may be 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.

5.2 Exploring Unidentified Pathways

Our system also enables exploration of novel pathways that could motivate future research. Fig. 6 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

1. There is still some debate over the extent of the ILF; see Catani et al. [6] for a recent discussion.

(a)

(b)

(c)

(d)

Fig. 5. 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 pathways produced by our system that may be involved in human visual processing. (b) Placing a second VOI in the right frontal lobe and using an AND operation shows only the pathways that traverse between the occipital lobe and a position on the right frontal lobe, which are stained blue. (c) Moving the second VOI to a symmetric position on the left captures the pathways traversing between the occipital lobe and 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.


Fig. 6. Exploration of candidate neural pathways. The query results show many pathways (shown in green) reflecting the well-known connection between the left and right occipital lobes, passing through the splenium. Intriguingly, the query also suggests a possible connection (shown in purple) 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.

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. These pathways then return back to the occipital lobe. 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.

5.3 Comparing Tractography Algorithms

As a third example, our system can be used to visually compare the pathways estimated by different tractography algorithms. Fig. 7 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

Fig. 7. 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.

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 10 minutes.

5.4 Investigating Interhemispherical Connections between Visual Areas

Neuroscientists can utilize this system to explore connections between the left and right hemispheres within visual areas of the human brain. Fig. 8 shows pathways connecting the calcarine fissure on the right hemisphere and visual areas on the left hemisphere. Guided by fMRI activation data (as described in Section 4.2.3), the neuroscientist placed a VOI within the calcarine fissure on the right hemisphere and another VOI on an interesting visual area on the left hemisphere. The neuroscientist then slowly moved the left hemisphere VOI to a location slightly below the previous position. Both VOIs were then iteratively adjusted to visualize pathway shape and location. By performing incremental VOI adjustments on the cortical surface, the neuroscientist was able to explore how pathways may traverse from the calcarine fissure to dorsal or ventral regions of the visual areas on the opposite hemisphere. Such queries allow for neuroscientists to discover pathways that they can isolate for further study. The examination of the entire occipital lobe required approximately 10 minutes.

6 DISCUSSION

This section describes the potential applications of our system, current limitations, and future directions for research.

6.1 Applications

Our colleagues in neuroscience 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 [20], [9], [17], 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. [25] 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 generate large amounts of verifiable, high-resolution data that could be browsed efficiently with our system.

6.2 Limitations

The speed we have achieved during dynamic queries relies upon extensive precomputation and this does have certain disadvantages. For example, using our system, a neuroscientist cannot interactively manipulate the parameters used to generate the pathways and surfaces that we display. This is a standard trade-off between speed and flexibility, and similar trade-offs have been described in other literatures (e.g., precomputed lighting calculations for computer graphics).

The connectivity results we have obtained with our system have been limited by the resolution of the data, the geometric distortions introduced by the data acquisition process, 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 $2x2x3$ 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.

6.3 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. Pathways from multiple subjects could be used to compare pathological cases (e.g., multiple sclerosis) to normal ones or simply to understand normal

Fig. 8. Investigation of connections between hemispheres in the retinotopic map. Guided by the fMRI activation data, a neuroscientist investigates the proposed pathways that connect the calcarine fissure on the right hemisphere to visual areas on the left hemisphere. (a) An fMRI scan is mapped to the cortical surface, in order to guide the neuroscientist to the appropriate visual areas. The fMRI color map can be toggled on and off as VOIs are placed. (b) Placing a VOI in the right hemisphere within the calcarine fissure (orange arrow) and a second VOI in the left hemisphere within the visual areas (blue arrow). (c) The cortical surface representation is removed so that the paths are visible within the context of the tomograms. (d) The VOI on the left hemisphere (blue arrow) is moved slightly ventrally along the cortical surface. With dynamic queries, the neuroscientist can repeat many small adjustments to the VOI in the left hemisphere or the VOI in the calcarine fissure on the right hemisphere to understand the types of connections within the visual areas.

population variance. It could also be used to study developmental disorders as they unfold over time within a single subject. While these problems can benefit from a dynamic query approach, they will also require the development of methods to coregister the various data sets.

We 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.

We plan to make the source code to our system freely available at http://graphics.stanford.edu/projects/dti.

7 CONCLUSION

In summary, we have presented a novel interaction technique and a software system designed to assist in the exploration of white matter connectivity in the brain. The key to our system's utility is its ability to respond to queries at interactive rates. This allows neuroscientists to optimize the critical loop of hypothesis generation and evaluation.

Finding known anatomical pathways with our system has demonstrated the ability to resolve large-scale anatomical

structures with DTI. Using our software as a hypothesis-generation tool for previously undiscovered pathways has been intriguing, but risky due to current limitations in acquisition technology and path tracing algorithms. Comparing pathways suggested by multiple path tracing algorithms has pointed out the uncertainty in path tracing, giving us reason to distrust the pathway estimates suggested by individual algorithms.

As the field proceeds forward with new acquisition technologies and more sophisticated path tracing algorithms, we should continue to benefit from the dynamic query approach. Indeed, this system remains an integral part of our own long-term plan for answering questions about brain connectivity.

ACKNOWLEDGMENTS

REFERENCES

[1] Wellcome Dept. of Imaging Neuroscience, SPM2: Statistical Parametric Mapping, http://www.fil.ion.ucl.ac.uk/spm/software/spm2/, 2004.

[2] C. Ahlberg, C. Williamson, and B. Shneiderman, "Dynamic Queries for Information Exploration: An Implementation and Evaluation," *Proc. SIGCHI '92*, pp. 619-626, 1992.

[3] D. Akers, A. Sherbondy, R. Mackenzie, R. Dougherty, and B. Wandell, "Exploration of the Brain's White Matter Pathways with Dynamic Queries," *Proc. IEEE Visualization '04*, pp. 377-384, 2004.

[4] R. Bammer, B. Acar, and M.E. Moseley, "In Vivo MR Tractography Using Diffusion Imaging," *European J. Radiology*, vol. 45, pp. 223-234, 2002.

[5] P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi, "In Vivo Fiber Tractography Using DT-MRI Data," *Magnetic Resonance in Medicine*, vol. 44, pp. 625-632, 2000.

[6] M. Catani, D.K. Jones, R. Donato, and D.H. ffytche, "Occipito-Temporal Connections in the Human Brain," *Brain*, vol. 126, pp. 2093-2107, 2003.

[7] 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," *Proc. Nat'l Academy of Sciences*, pp. 10422-10427, 1999.

[8] M.J. da Silva, S. Zhang, C. Demiralp, and D.H. Laidlaw, "Visualizing Diffusion Tensor Volume Differences," *Proc. IEEE Visualization '01, Work in Progress*, 2001.

[9] 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.

[10] R.F. Dougherty, V.M. Koch, A.A. Brewer, B. Fischer, J. Modersitzki, and B.A. Wandell, "Visual Field Representations and Locations of Visual Areas V1/2/3 in Human Visual Cortex," *J. Vision*, vol. 3, pp. 586-598, 2003.

[11] M.S. Gazzaniga, R.B. Ivry, and G.R. Mangun, *Cognitive Neuroscience*. New York: W.W. Norton and Company, 2002.

[12] S. Gottschalk, M.C. Lin, and D. Manocha, "OBBTree: A Hierarchical Structure for Rapid Interference Detection," *Proc. SIGGRAPH 96*, pp. 171-180, 1996.

[13] H. Hochheiser and B. Shneiderman, "Visual Queries for Finding Patterns in Time Series Data," technical report, Univ. of Maryland, 2002.

[14] E. Hutchins, J. Hollan, and D. Norman, "Direct Manipulation Interfaces," *User Centered System Design*, D. Norman and S. Draper, eds., pp. 87-124, Erlbaum, 1986.

[15] G. Kindlmann and D. Weinstein, "Hue-Balls and Lit-Tensors for Direct Volume Rendering of Diffusion Tensor Fields," *Proc. IEEE Visualization '99*, pp. 183-189, 1999.

[16] G. Kindlmann, D. Weinstein, and D. Hart, "Strategies for Direct Volume Rendering of Diffusion Tensor Fields," *IEEE Trans. Visualization and Computer Graphics*, vol. 6, pp. 124-138, 2000.

[17] 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*, vol. 25, pp. 493-500, 2000.

[18] M. Lazar, D.M. Weinstein, J.S. Tsuruda, K.M. Hasan, K. Arfanakis, M.E. Meyerand, B. Badie, H.A. Rowley, V. Haughton, A. Field, and A.L. Alexander, "White Matter Tractography Using Diffusion Tensor Deflection," *Human Brain Mapping*, vol. 18, pp. 306-321, 2003.

[19] W.E. Lorensen and H.E. Cline, "Marching Cubes: A High Resolution 3D Surface Construction Algorithm," *Proc. SIGGRAPH '87*, pp. 163-169, 1987.

[20] 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," *J. Cognitive Neuroscience*, pp. 629-636, 2002.

[21] 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," *Annals of Neurology*, vol. 45, pp. 265-269, 1999.

[22] 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. Magnetic Resonance*, vol. 154, no. 1, pp. 85-100, 2002.

[23] C. Pierpaoli and P.J. Basser, "Toward a Quantitative Assessment of Diffusion Anisotropy," *Magnetic Resonance Magazine*, pp. 893-906, 1996.

[24] C. Poupon, J. Mangin, C.A. Clark, V. Frouin, J. Regis, D. LeBihan, and I. Bloch, "Towards Inference of Human Brain Connectivity from MR Diffusion Tensor Data," *Medical Image Analysis*, vol. 5, pp. 1-15, 2001.

[25] 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*, vol. 34, pp. 685-700, 2002.

[26] W. Schroeder, K. Marin, and B. Lorenson, *The Visualization Toolkit.* Prentice Hall, 1997.

[27] M.I. Sereno, A.M. Dale, J.B. Reppas, K.K. Kwong, J.W. Belliveau, T.J. Brady, B.R. Rosen, and R.B. Tootell, "Borders of Multiple Human Visual Areas in Humans Revealed by Functional MRI," *Science*, vol. 268, pp. 889-893, 1995.

[28] B. Shneiderman, "Direct Manipulation: A Step beyond Programming Languages," *Computer*, vol. 16, no. 8, pp. 57-69, Aug. 1983.

[29] G. Taubin, T. Zhang, and G. Golub, "Optimal Surface Smoothing as Filter Design," *Proc. Fourth European Conf. Computer Vision*, pp. 283-292, 1996.

[30] P.C. Teo, G. Sapiro, and B.A. Wandell, "Creating Connected Representations of Cortical Gray Matter for Functional MRI Visualization," *IEEE Trans. Medical Imaging*, vol. 16, no. 6, pp. 852-863, 1997.

[31] A. Vilanova i Bartroli, G. Berenschot, and C. van Pul, "DTI Visualization with Streamsurfaces and Evenly-Spaced Volume Seeding," *Proc. VisSym '04 Symp. Visualization*, pp. 173-182, 2004.

[32] 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*, pp. 77-87, 2004.

[33] D.M. Weinstein, G.L. Kindlmann, and E.C. Lundberg, "Tensorlines: Advection-Diffusion Based Propagation through Diffusion Tensor Fields," *Proc. IEEE Visualization '99*, pp. 249-254, 1999.

[34] A. Wenger, D.F. Keefe, S. Zhang, and D.H. Laidlaw, "Interactive Volume Rendering of Thin Thread Structures within Multivalued Scientific Data Sets," *IEEE Trans. Visualization and Computer Graphics*, vol. 10, pp. 100-107, 2004.

[35] C.F. Westin, M. SE, M. Nabavi, A. Jolesz, and F. Kikinis, "Processing and Visualization for Diffusion Tensor MRI," *Medical Image Analysis*, vol. 6, pp. 93-108, 2002.

[36] S. Zhang, C. Demiralp, and D. Laidlaw, "Visualizing Diffusion Tensor MR Images Using Streamtubes and Streamsurfaces," *IEEE Trans. Visualization and Computer Graphics*, vol. 9, pp. 454-462, 2003.

[37] L. Zhukov and A.H. Barr, "Oriented Tensor Reconstruction: Tracing Neural Pathways from Diffusion Tensor MRI," *Proc. IEEE Visualization '02*, pp. 387-394, 2002.

Anthony Sherbondy received the BS degree in electrical engineering in 1998 from Pennsylvania State University. He also received the MS degree in electrical engineering from Pennsylvania State University in 2001. Currently, he is a PhD candidate in the Department of Electrical Engineering at Stanford University. He is interested in visualization of biological and medical images. He conducted this research under the support of the US National Institutes of Health NIGMS Scientist Training Fellowship.

David Akers received the ScB degree in computer science and mathematics in 1999 from Brown University. He is currently a PhD student in the Department of Computer Science at Stanford University. He specializes in scientific visualization and educational software development, with a particular focus on interactive techniques. He is a recipient of the US National Science Foundation Graduate Research Fellowship.

Rachel Mackenzie received the BS degree in symbolic systems with a minor in biology in June 2004. Also in June 2004, Rachel received the MA degree in cognitive psychology, working with Professor Barbara Tversky. Her current research is on the narrative structure of diagrams and the use of diagrammatic symbols.

Robert Dougherty received the BA degree from Rutgers University in 1991 and the PhD degree in experimental psychology from the University of California at Santa Cruz in 1996. He specializes in the measurement of perception (visual and auditory psychophysics) and the measurement of brains using functional and anatomical MRI. He also worked as a postdoctral researcher in the Department of Ophthalmology at the University of British Columbia and BC's Children's Hospital Visual Neuroscience Lab.

Brian Wandell is the Stein Family Professor at Stanford University. His research includes image systems engineering and visual neuroscience. In engineering, he cofounded (with J. Goodman) the Stanford Center for Image Systems Engineering. In neuroscience, his work includes functional MRI, diffusion tensor imaging, and computational modeling to understand the activity and development of the visual portions of the brain. He is the author of *Foundations of Vision*, a textbook on vision science. He won the 1986 Troland Research Award from the US National Academy of Sciences for his work in color vision; he was made a fellow of the Optical Society of America in 1990; in 1997, he was made a McKnight Senior Investigator and given the Edridge-Green Medal in Ophthalmology for his work in visual neuroscience. In 2000, he was awarded the Macbeth Prize from the Inter-Society Color Council. He was elected to the US National Academy of Sciences in 2003.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.