

Exploration of the Brain's White Matter Structure through Visual Abstraction and Multi-Scale Local Fiber Tract Contraction

Maarten H. Everts, Eric Begue, Henk Bekker,

Jos B.T.M. Roerdink, *Senior Member, IEEE*, and Tobias Isenberg, *Senior Member, IEEE*

To appear in an IEEE VGTC sponsored conference.

Abstract—We present a visualization technique for brain fiber tracts from DTI data that provides insight into the structure of white matter through visual abstraction. We achieve this abstraction by analyzing the local similarity of tract segment directions at different scales using a stepwise increase of the search range. Next, locally similar tract segments are moved toward each other in an iterative process, resulting in a local contraction of tracts perpendicular to the local tract direction at a given scale. This not only leads to the abstraction of the global structure of the white matter as represented by the tracts, but also creates volumetric voids. This increase of empty space decreases the mutual occlusion of tracts and, consequently, results in a better understanding of the brain's three-dimensional fiber tract structure. Our implementation supports an interactive and continuous transition between the original and the abstracted representations via various scale levels of similarity. We also support the selection of groups of tracts, which are highlighted and rendered with the abstracted visualization as context.

Index Terms—Diffusion Tensor Imaging (DTI), fiber tracts, visual abstraction, multi-scale representation, illustrative visualization.



1 INTRODUCTION

THE task of uncovering the functionality of the human brain in all its detail has engaged researchers for centuries. Fundamental to gaining such an understanding is to comprehend the brain's complex anatomy which comprises ganglia, blood vessels, grey matter, as well as white matter. This white matter that contains connections between nerve cells in the gray matter and that consists of approximately 10^5 km of myelinated axons is the focus of our work. Often, neuroscientists do not only want to understand its anatomy in general but rather based on patient-specific data. While modern imaging techniques provide non-invasive ways to obtain data about the white matter's anatomy, it remains difficult to understand the structure because of its complexity.

We propose a new technique for showing the brain's white matter structure using fiber tracts extracted for the whole brain. Our goal is to represent all extracted fiber tracts simultaneously and at the same time visually abstract them so that global patterns become clear for a given level of scale. Our approach is based on diffusion tensor imaging (DTI) as one of the non-invasive medical imaging technologies [2]. DTI records information about anisotropic physical structures in that it reflects the self-diffusion

properties of water in fibrous material such as the brain's white matter. The resulting tensor field that represents the anisotropic diffusion is used in a variety of analysis methods, of which fiber tracking [31] is a popular approach. The general idea of fiber tracking is, starting from a seed point, to generate a curve following the main diffusion direction in the tensor field. Even though the resolution of DTI is not nearly high enough for tracing individual axons, fiber tracking can provide useful information about the connectivity of brain regions and the structure of axon bundles. Such white matter fiber bundles are typically either tubular-shaped, such as the *cingulum*, or sheet-like, such as the body of the *corpus callosum*, the *corona radiata*, and the rostral part of the corticospinal tract [41], [42].

Placing seeds to initialize fiber tracking is typically done in a particular region-of-interest or throughout the entire brain. We use the latter method which produces a huge collection of dense fiber tracts (typically 10^4 – 10^5) that pass through the brain in complex ways (Fig. 1(a)). When visualized in a naïve way, most of the global structure of the tracts including tube, sheet, or fan formations, is occluded due to the high density of the tracts. If we look at the tracts in detail, however, we can observe that many neighboring tracts follow almost the same path for their entire length, others follow the same path only for subsections and then diverge to run in parallel with other tracts. We use this observation that many tracts are locally virtually parallel as the basis of our approach.

Specifically, we start by subdividing all tracts into small segments of equal size. For every pair of consecutive vertices of nearby tracts we then determine the local tract similarity based on the direction of incident segments. This processing

- M.H. Everts is with TNO, the Netherlands. E-mail: maarten.everts@tno.nl.
- E. Begue and H. Bekker are with the University of Groningen, the Netherlands. E-mail: {ericbeg@gmail.com, h.bekker@rug.nl}.
- J.B.T.M. Roerdink is with the University of Groningen and with the Neuroimaging Center of the University Medical Center Groningen, the Netherlands. E-mail: j.b.t.m.roerdink@rug.nl.
- T. Isenberg is with Inria, France. E-mail: tobias.isenberg@inria.fr.




Fig. 1. DTI fiber tract abstraction based on similarity and proximity (tractography was performed for the complete brain, but all images show only one hemisphere of the brain). Tracts are colored according to direction [33]; top-down: green; left-right: red; perpendicular to view-plane: blue. Central structure in red: *corpus callosum*; green structure directly above the *corpus callosum*: *cingulum*. (a) uncontracted; (b) contracted at a scale of 5 mm; (c) a subset of the contracted structure highlighted.

is done for vertices with a separation less than the targeted scale. For every vertex pair with a similarity score above a certain threshold we add an edge to a similarity graph that connects these vertices and, thus, the fiber tracts. Based on this graph we carry out an iterative lateral contraction of tract segments by relocating tract vertices to be closer to similar ones. The total number of tracts remains unchanged during the contraction. This process results in a lateral contraction of fiber tract structures and thus a visually abstracted representation (Fig. 1(b)).

While the reduction of visual information is useful in itself, it is also accompanied by the creation of *volumetric voids*; space previously occupied by many parallel tract segments becomes mostly empty through the contraction process (Fig. 1(b)). These added voids throughout the brain decrease occlusion of the fiber tract patterns by the fibers themselves so that it becomes possible to see deeper into the brain and better interpret its white matter structure. Full-brain tractography methods are used to get an initial and unbiased overview of the connectivity structure. Often, it is a starting point of a more detailed analysis which involves selecting and analyzing meaningful clusters of the tracts. One of the main benefits of our technique is that it removes noise from the large fiber tract datasets, as well as clarifies the structure, and thus allows for easier fiber bundle selection.

To further enhance the understanding we allow users to interactively and *continuously move* between the original and the contracted state of the fiber tracts which makes it easier to relate the two states. Furthermore, we allow users to select subsets of fibers to explore local patterns in the contracted state (Fig. 1(c)). We emphasize, however, that before doing such a local analysis an initial whole-brain fiber computation is required. Restricting the fiber contraction to a local subset of fibers would give the contraction process too much freedom, leading to large displacements of contracted fiber bundles from their original anatomical locations, especially at the boundaries of the selected region, which is undesirable.

Of course, despite the inherent tract deformation that comes with our method, the introduced abstraction needs to be anatomically meaningful. For this purpose we analyze the amount of displacement our abstraction scheme introduces

to the location of the fiber tracts and compare this to the original anatomical locations. The results show that, for up to a search radius of ≈ 2 mm, the contracted results are located within the regions occupied by the original fiber tracts with a high accuracy. However, higher search radii that would normally lead to anatomical incorrectness can still be useful. Domain experts pointed out that bundles can be *selected* in the contracted state (with some anatomical distortion), but then *analyzed* in the uncontracted state. The continuous movement between the original and the contracted state that our method supports can help the user to quickly build up a mental image of the relation between the different states.

We emphasize that our approach differs from the process of clustering or grouping DTI fiber tracts (sometimes called fiber bundle identification [45]) into anatomically correct fiber bundles (*corpus callosum*, *cingulum*, etc.). This process is typically accompanied by a color labeling but no displacement of the fibers tracts occurs. In contrast, our goal is to create a simplified visualization of large sets of fiber tracts; we thus do not aim to identify anatomical fiber bundles. Our method also bears some similarity to edge bundling, which determines both a grouping of the edges and the paths of the edges. The difference is that we displace only local fiber tract segments, while in edge bundling a global displacement of the paths occurs. To avoid confusion, we avoid the term “bundling” when referring to our method and instead speak about fiber tract contraction.

The remainder of this paper is organized as follows. Related work is discussed in Section 2. The creation of contracted fiber tract configurations is explained in Section 3 and their visualizations in Section 4. Next, we present our results in Section 5 and analyze the introduced fiber tract displacement in Section 6. We discuss the merits and limitations of our technique in Section 7 and conclude the paper in Section 8.

2 RELATED WORK

There exists a wide variety of techniques for generating abstractions and visualizations that provide an understanding of the brain’s white matter structure. The vast majority of these approaches rely on DTI data because these capture the local anisotropy of the myelinated axons in the brain.

One may start from the fractional anisotropy (FA) field derived from DTI data [3]. The tract-based spatial statistics approach by Smith *et al.* [37] produces a 3D volumetric skeleton from a mean FA volume as an abstract representation of white matter. A related representation—anisotropy crease surfaces—was proposed by Kindlmann *et al.* [23], also based on the FA values. In a complementary approach, Schultz *et al.* [36] analyze the topology of the tensor field based on probabilistic fiber tracking, thus illustrating the structurally important information. Another technique by Kindlmann and Westin [24] represents the tensor field by glyphs and uses glyph packing to better illustrate the underlying structure.

In contrast to these techniques and similar to our own method, many approaches that visualize white matter first employ deterministic tractography [31] and then visualize the extracted fiber tracts, possibly processing them further. The visualization of fiber tracts (e.g., [13], [18], [34], [35], [38]) is typically done for subsets of tracts because rendering a large number of fiber tracts not only leads to occlusion and clutter but also obscures the structurally important aspects of the data. Methods to determine which tracts to show include seeding at regions-of-interest and interactive tract selection [1], [5]. With our technique we can both include all extracted fiber tracts in the visualization such that the overall structure is revealed and at the same time make use of tract selection for emphasis.

The overall goal of our technique is to reveal the structure of the brain's white matter by introducing abstraction. A similar effect can also be achieved by careful tract seeding—without the need to abstract the locations of fiber tracts. Such approaches were presented, for instance, by Merhof *et al.* [29], Vilanova *et al.* [40], and Zhang *et al.* [46], using streamlines, streamtubes, or streamsurfaces. While these techniques produce good visual results and need less geometry, they always only show representations of specific instances of fiber tracts—as opposed to abstracted tracts that represent many different extracted paths in the volume.

An important approach for understanding white matter structure is the use of grouping or clustering [7], [12], [25], [39]. Here, DTI fiber tracts are classified into anatomically correct fiber bundles (*corpus callosum*, *cingulum*, etc.), often accompanied by a color labeling [45]. Segmentation of sheetlike fiber tracts into anatomically meaningful bundles was considered by Maddah *et al.* [27], [28]. O'Donnell *et al.* [32] proposed the use of a white matter atlas to learn a model of the common white matter structures present in a group of subjects. Interaction with the clustered data can be added to support exploration and navigation through tractography data [8], [20]. Evaluation of fiber clustering methods was considered by Moberts *et al.* [30], who developed a framework to validate different clustering methods and different (global) fiber similarity measures by comparison with a manual classification which was used as a ground truth.

Clustering relates to our own approach. In both cases the similarity of tracts is important. Our approach, however, differs in that, first, clustering approaches do not introduce

a displacement of the fiber tracts which our technique does. Second, our approach does not aim to identify anatomical fiber bundles, in the sense of classifying tracts into anatomical structures; our goal is to obtain an abstraction of large sets of fiber tracts through simplification by lateral contraction. Third, while clustering determines the similarity based on the whole length of the tract, we establish similarity of fiber tract segments at a local level. This local similarity aspect relates to the visualization of fiber tract coherence by Hlawitschka *et al.* [15] who calculate a coherence measure based on the deviation of fiber tracts in small neighborhoods. The resulting scalar field is not only visualized directly, but also used as a transparency mask to enhance fiber tract visualization. In contrast, we aim to show the overall structure while keeping the connectivity intact. This structure can also be captured in terms of an abstract network [14] which nicely illustrates the connectivity of regions but loses the relation to the actual anatomy to some degree.

To study functional brain connectivity, Böttger *et al.* [6] apply edge bundling based on resting-state fMRI data. A method to cluster functional MRI data based on functional unit maps was proposed by Crippa *et al.* [10]. For functional connections, however, the issue of anatomical faithfulness is less pressing, since there is no one-to-one relation between anatomical and functional connections in the brain.

In another related paper, Yu *et al.* [44] introduce hierarchical streamline bundles to simplify and visualize 3D flow fields defined on regular grids. Seed placement and streamline generation occur according to flow saliency. Spatially neighboring and geometrically similar streamlines are grouped to construct a hierarchy of streamline bundles that partitions the flow field. Differences with our approach are: (i) only a grouping, not a displacement, of streamlines is determined (like in fiber clustering); (ii) a global, not a local, metric is used to compute streamline similarity; and (iii) seeding is selective (saliency based), while ours is uniform in the volume.

Finally, our approach of fiber tract contraction is related to techniques used in 2D graph drawing, including force-directed layouts [11], and, in particular, the graph edge bundling by Holten *et al.* for hierarchical [16] and general graphs [17]. The latter is especially relevant due to its use of an iterative force-directed approach. An important difference with our approach is that for the edges in the 2D graph the path is a feature denoting the connection of two nodes whose spatial location is completely flexible, whereas for our fiber tracts the path location is relevant and important as an anatomical feature, i.e., the deformation of tracts has to be limited. A common property of both approaches is that they create spaces between the original edges by displacing them, so as to improve understandability.

3 FIBER TRACT ANALYSIS & CONTRACTION

Our goal is to visualize DTI fiber tracts in their entirety such that the depictions assist viewers in understanding the structure of white matter, at different levels of scale. Our technique is based on fiber contraction that takes the local

conditions into account but also considers the location of the original tracts to provide anatomic relevance. We achieve these goals through a two-stage process. First, an analysis stage inspects the fiber tracts to analyze the local similarity of tract segments with respect to proximity and direction. Next, an iterative process uses this similarity information to move fiber tract segments towards similarly oriented neighboring segments, resulting in a contracted state. In interactive exploration, scientists can transition between the uncontracted and the contracted state via several scale levels.

3.1 Analysis of Local Direction Similarity

We consider two tracts to be locally similar if the directions of the considered tract segments are approximately parallel and not located far from each other. We capture this similarity information in a graph whose edges connect fiber tracts where they are locally similar. Hence, the nodes of the similarity graph are vertices of the fiber tracts, while the graph's edges represent similarity relations between these vertices and, thus, between tract segments. The goal of the analysis stage is to create the similarity graph, which will later be used to inform the contraction stage.

To faithfully describe the local similarity for a set of fiber tract segments with the similarity graph we rely on the polygonal fiber tracts being evenly tessellated, i. e., all tract segments should have the same length. Because the fiber tracts resulting from tractography typically do not have this property, we first resample each tract such that each pair of neighboring vertices on the tract has the same distance d_{sample} .

Based on the set of resampled tracts we capture their local similarity. While such similarity can be established for a more general case of points on two planes [22], we use a more direct heuristic¹ because we only deal with vertices on linear structures. Similar to Corouge *et al.* [9] and Ding *et al.* [12], we process each pair of fiber tracts A and B to find, for each vertex on A , its nearest neighbor vertex on B , and vice versa. This step yields a list L of vertex pairs of the form $(v, nn(v))$, where v is a vertex on tract A or B and $nn(v)$ is v 's nearest neighbor on the other tract. At this stage, the number of pairs in L is equal to the sum of the number of vertices in A and B . However, L may also contain superfluous entries, i. e., a pair (p, q) can be removed if (q, p) also exists in L .

Next, we filter the entries of L based on three conditions, a distance condition, an angle condition, and a combinatorial condition, before creating edges in the similarity graph. We create an edge if and only if a vertex pair (p, q) in L meets all of the following conditions:

Condition 1: the distance between p and q is smaller than a given length d_{max} ;

Condition 2: at least one of the segments incident to p is approximately parallel to one of the segments incident to q , where “approximately parallel” means that the angle




Fig. 2. Illustration of how tract vertices are connected by edges. Only lines with pure dashes connect vertices that form an edge in the similarity graph, the other vertex pairs are discarded because of Condition 1 (dots) and Condition 3 (dash-dot) in Section 3.1.

between the segments is smaller than a predefined angle θ_{par} ;

Condition 3: the nearest-neighbor relation of p and q is approximately mutual, that is, it holds that $|\text{index}(nn(q)) - \text{index}(p)| \leq 1$ and $|\text{index}(nn(p)) - \text{index}(q)| \leq 1$, where the function $\text{index}(v)$ returns the index of v on the tract to which it belongs.

This last condition requires some additional motivation. Fig. 2 shows a common configuration of tracts. Assume that tracts A and B are approximately 3 mm apart, $d_{\text{sample}} = 1$ mm and $d_{\text{max}} = 6$ mm. If only Condition 1 is used then the pairs (A_1, B_1) , (A_1, B_2) , (A_1, B_3) would not define an edge because their distances are all greater than d_{max} , but the pairs (A_1, B_4) , (A_1, B_5) , (A_1, B_6) , (A_1, B_7) would each lead to an edge in the similarity graph because they satisfy Conditions 1 and 2. Such an edge is undesirable because this would mean that during the actual contraction process vertices A_1 and B_4 would move toward each other, A_1 towards B_5 , etc. However, we only want A_1 and B_6 to move towards each other and A_1 and B_7 to move towards each other. Condition 3, therefore, prevents (A_1, B_4) and (A_1, B_5) from being turned into edges and to cause unwanted edge contraction later on. The value of 1 in the condition is somewhat arbitrary. It has to be greater than 0 but not too large, and a value of 1 proved to work well in our experiments.

Each newly added edge is also given a real-valued attribute that captures the distance between the pair's vertices. This value is later used in the iterative contraction process.

The parameters θ_{par} and d_{max} can be interactively changed by the user, depending on the noise level of the data.

3.2 Iterative Contraction of Fiber Tracts

In the second stage, the information in the similarity graph is used to iteratively move vertices toward each other. This has the general effect that tract segments move towards nearby tract segments of similar direction. The process is realized iteratively because a vertex in the similarity graph is, in general, connected to multiple other vertices. During each iteration, we accumulate the influence of all incident graph edges for each vertex in a combined displacement vector.

One iteration consists of the following steps. First, the displacement vector of each vertex is initialized to zero. Similar to force-directed layouts [11, Ch. 10] we then calculate the displacement for each graph edge (p, q) that will move p and q towards each other such that they will

1. We experimented with variations of this heuristic but what we describe here proves to strike a good balance between simplicity and effectiveness.




Fig. 3. Illustration of noise (dents, bulges, steps) in neighboring tracts. We remove this by applying a Gaussian filter to the displacements.




Fig. 4. Illustration of diverging tracts (top) and double-linking in the similarity graph (bottom).

meet halfway, i. e., at $(\mathbf{r}_p + \mathbf{r}_q)/2$, with \mathbf{r}_p and \mathbf{r}_q being the current position vectors of vertices p and q , respectively. If p and q would not be connected to any other vertices they would end up at the same position, i. e., both halfway between \mathbf{r}_p and \mathbf{r}_q . However, that will not be the case in general because p and q are also connected to other vertices.

Because such a simple force-directed approach can potentially lead to run-away situations, we divide the displacement vector by the number of edges connected to the processed vertex to compute the average, similar to what is done in the barycenter force-directed method [11, Ch. 10.2]. In addition, we remove noise along a fiber tract (as illustrated in Fig. 3) by applying a small Gaussian kernel to the displacements along each tract. This smoothing also addresses the issue of otherwise sudden transitions where merging tract segments come into the influence range of other tract segments (Fig. 4, top), and the issue of the occasional double-linking (Fig. 4, bottom). In a final adjustment, we ensure that vertices can only be moved perpendicularly to a tract's local direction, as obtained from the average orientation of the two tract edges incident to the processed vertex, and project the accumulated and filtered displacement vectors onto the orthogonal plane. After all displacement vectors have been derived, they are applied to their respective vertices and the next iteration starts. Several iterations are necessary to achieve a contracted state; we typically use in the order of 40 iterations.

The contraction process deforms the tracts themselves: fiber tracts are not simply subjected to rigid transformations but instead are locally modified. These confined deformations force tracts to locally follow the same paths. This process creates voids between sets of fiber tracts, one of

the main goals of the method. However, the amount of deformation is kept to a minimum by moving vertices mainly perpendicularly to the tract direction and by applying Gaussian smoothing.

3.3 Scale-Dependent Abstraction

At this stage, the fiber tract analysis and contraction incorporate a notion of scale due to the use of d_{max} to limit the neighborhood of what is being considered to be similar. Different values for d_{max} result in different notions of scale: only close structures are contracted for small values of d_{max} , while larger structures are merged for larger values of d_{max} . During contraction, however, we only process graph edges whose length is smaller than a given $d_{max,i} \in [0, d_{max}]$, with i indicating scale. By performing the contraction for increasing values of $d_{max,i}$ we can create contracted fiber tract representations for growing scales. Because we use the same fiber tract tessellations for each of these computations we can relate the position of a fiber tract vertex in the uncontracted state to the position of the same vertex in a contracted state for any of the computed scales.

To better understand the impact of d_{max} on the abstraction that is introduced through the contraction we have explored d_{max} values of up to the equivalent of 7 mm and report on this evaluation in Section 6.

4 VISUALIZATION AND INTERACTION

The result of the contraction process is that for each tract not only its original vertex positions have been stored but also the positions of its vertices at the different contraction scales. Based on this data we provide a visualization that enables the user to interactively and seamlessly move from one contraction scale to the next to explore white matter structure. In addition, we provide means to employ filtering, a 2D lens tool, and tract selection.

4.1 Visualization at Multiple Scales

In the user interface of our visualization users can interactively control the scale of contraction. Even though we only compute a discrete number of scales (typically in 1 mm increments, so for $d_{max} = 5$ mm we compute 5 discrete contracted representations), we provide a continuous transition by linearly interpolating the tract vertex positions between consecutive scales. This continuous transition prevents sudden jumps of tracts while users are interacting with the scale slider and helps them to follow the tract locations across the series of scales.

4.2 Rendering

For rendering we store the tract vertex positions for the different scales on the graphics card and perform linear interpolation between consecutive scales in a vertex shader. We can thus render the tract representation and move between scales at interactive speeds. The rendering is based on the depth-dependent halo visualization for dense line data [13], with one small adaptation. Instead of duplicating

a line's vertices to create a view-oriented triangle strip, we use a geometry shader. The reason is that the creation of the view-oriented triangle strips depends on the local direction of the line at each vertex, which was pre-computed in the original approach. In our case, however, pre-computation is unfeasible because the interpolation between scales results in the local direction of the lines depending on the chosen scale. This problem is solved by calculating the direction and creating the triangle-strip on the fly in a geometry shader. This has the additional advantage that it requires much less data to be transferred to the graphics card.

The complete process employs a vertex shader that interpolates vertex positions depending on the chosen scale, then transfers the positions to the geometry shader, which turns the line strip into a view-oriented triangle strip. The fragments that result from rasterization are processed in a fragment shader that assigns white for halos, black for lines, and performs the depth manipulation as explained by Everts *et al.* [13]. The lines can also be colored based on the direction as is common in DTI-tract visualization [33] (e.g., Fig. 1). (Note that color is not used here to label anatomical fiber bundles, as in fiber bundle identification [45].) The additional colors assist viewers in distinguishing directions in dense regions after contraction: in strongly contracted regions the difference in depth is minimal so that no halo is generated which would otherwise support depth perception.

4.3 Filtering for Additional Abstraction

The contraction of tracts based on local similarity results in volumetric voids which reduce occlusion artifacts (e.g., Fig. 1(b)). Occlusion can be decreased even more by interactively filtering out smaller structures that do not have as many neighboring tracts with similar direction (Fig. 10). The filtering attribute is the number of similarity graph edges that are connected to a vertex of a fiber tract. By passing the vertex degree as an attribute to each vertex, the filtering can be done in the fragment shader. This has the added effect that the interpolation done in the rasterization stage enables filtering over a continuous range.

4.4 Focus+Context with a 2D Lens

We also allow users to explore the effects of scale locally with a two-dimensional virtual lens [4]. If this lens is placed over the scene in screen-space (Fig. 11) it affects the scale value that is used to derive the positions of the vertices beneath it, while other vertices remain at their normal position as determined by the global scale setting. We first check for each rendered vertex' screen position whether it is inside the lens and, based on its location inside the lens, compute an attenuation factor (which is 1 in a center region and 0 at the perimeter). Next, we reposition affected vertices to the location determined by the lens scale setting and the vertex attenuation factor, which ensures a gradual change from global to local scale. We implemented the lens in the same vertex shader that is also responsible for the scale-dependent interpolation of vertices to ensure fast processing. As a result, users can move the lens over the




Fig. 5. The selected subset of tracts is shown in its uncontracted state, while the remainder is shown contracted. In addition, a large amount of filtering was applied to the context, showing only vertices with a degree in the similarity graph of more than 210. Notice the sheet-like nature of the *corpus callosum* while the *cingulum* is in its original state (compare it with its contracted shape, e.g., in Fig. 1(c)).

scene in screen-space to locally explore the effect of the contraction.

4.5 Tract Selection

As an additional means of exploring white matter structure and its connectivity we support a simple interactive selection of groups of fiber tracts. Selected tracts are highlighted in the contracted state by rendering the selection with regular colors while the remaining tracts are rendered with less vivid colors (Fig. 1(c) and Fig. 5). Specifically, we let the tract direction affect the color in a range between 20% and 40% of each color channel. This has the effect that the fiber tracts in the context are considerably darker than the ones in focus. The context's subdued color scheme enforces the emphasis of the selection which still uses vivid colors. Alternatively, we also permit users to only show selected tracts or to render selected tracts at their original locations as shown in Fig. 5.

The interactive selection of fiber tract subsets makes it easier for viewers to follow the path and position of the selected tracts and relate these to their contracted state. We use a sphere as the interaction object that can be placed in 3D space; tracts that pass through it in the *uncontracted state* are selected [1], [5]. For the positioning of the sphere we use an approach similar to that found in many 3D modeling programs (such as Blender). Users can click on the sphere and, while the mouse button is down, three lines are drawn through the center of the sphere (Fig. 6), each representing one coordinate axis. Upon moving the mouse, the sphere is moved along the straight line from the original to the current mouse position. With repeated clicks the sphere can be positioned anywhere in the dataset. For improved depth perception of the sphere placement we extend the three coordinate axis lines to where they intersect the bounding box of the whole dataset (see Fig. 6).




Fig. 6. Fiber bundle selection using spheres and coordinate axes. The blue sphere (partially hidden) represents the current selection, while the red one is being moved and will specify a new region of interest.

5 RESULTS

The fiber tract visualizations shown in this paper are based on DT-MRI data that was acquired on a 3T MRI system (Philips Intera) from a healthy volunteer whose consent was obtained prior to scanning. Diffusion Tensor Imaging was performed using a diffusion weighted spin-echo, echo-planar imaging technique. The DTI parameters were as follows: 240 mm \times 240 mm field of view; 128 \times 128 matrix size; 51 slices; 1.85 \times 1.85 \times 2 mm³ imaging resolution; 5,485 ms repetition time; 74 ms echo time. In total, 61 volumes were acquired per subject, one without diffusion weighting ($b = 0$ s/mm²) and 60 volumes with diffusion weighting ($b = 800$ s/mm²) along 60 noncollinear directions. To correct for susceptibility artefacts, two acquisitions were used: one with fat-shift direction in the posterior direction (APP) and one in the anterior direction (APA). The fiber tracts were generated by the Diffusion Toolkit program, using the FACT algorithm [31]. The contraction method, however, does not depend on the specific method to generate tracts. As an alternative, for example, one could use an integration with a fixed step-size as described, e. g., by Zhang *et al.* [46].

Fig. 1(a) shows the uncontracted fiber tracts of one brain hemisphere, the color values indicating the tracts' local direction. While it is possible to identify certain structures such as the *corpus callosum* (the flat bundle of neural fibers connecting the left and right cerebral hemispheres), much of the structure of the fiber tracts below and above this region is hidden by other tracts. Fig. 1(b) shows the same view, contracted to a scale of 5 mm. Now, the sheet-like structure of the *corpus callosum* becomes evident and more detail is visible for the upper regions. Similarly, the *cingulum* (a white matter fiber bundle located above the *corpus callosum* and interconnecting limbic structures) is now visualized as a pronounced structure and tracts emerging from it to the brain's outer regions are visible. Also, the structure in the brain's central region is shown much more clearly; the merging and splitting of fiber tract sets is well depicted. Finally, a selection on the *corpus callosum* was made in Fig. 1(c) which shows the fiber tracts in contracted form emerging from this location. Here we show the whole length of the tracts in focus, while the context of unselected tracts is partly removed by a cutting plane.

Fig. 7 further emphasizes that fiber tracts in several major




Fig. 7. The sheetness of the underlying tract data in both (a) the original and (b) the contracted state is emphasized by using a cutting plane. To further highlight this characteristic, parts of the fiber tracts close to the cutting planes are rendered in black.




Fig. 8. Two randomly chosen fiber tracts, shown in red, with the remaining fibers in black and white: (a) uncontracted; (b) contracted at a scale of 2 mm. In the uncontracted situation, the red fiber tracts are hardly visible while the contracted situation gives a good impression of their global trajectory.

brain white matter structures (body of the *corpus callosum*, rostral part of the corticospinal tract) are organized in sheets rather than tubular bundles [28], [41], [42]. The parts of the fiber tracts close to an active cutting plane are rendered in




Fig. 9. A selection of fiber tracts in the *corpus callosum* (a) before and (b) after contraction. In (c) and (d) the same selection of fiber tracts is shown from a different viewpoint. Notice that the contraction is done with respect to all tracts, not only the selection.

black to clearly show where the tracts intersect the planes. While the uncontracted state in Fig. 7(a) still shows large regions with intersections, the visualization of the contracted state in Fig. 7(b) clearly illustrates the sheet-like structure of many white matter regions. In addition, the added space that results from the iterative contraction allows us to look further into the brain in Fig. 7(b), and thus to see smaller structures such as the ones that emerge from the *cingulum*.

As we will argue in Section 6, beyond $d_{max} \approx 2$ mm anatomical distortion becomes significant. Therefore we investigate the effect of volumetric voids for $d_{max} = 2$ mm. To that end we compare two images (Fig. 8), one in the uncontracted state and the other contracted at a scale of 2 mm. Two fairly randomly chosen tracts are shown in red. In the uncontracted state these tracts are almost invisible, while in the contracted state the behaviour of the tracts is much clearer. This shows that, even at a fairly low contraction scale of 2 mm, enough ‘volumetric voids’ are created which facilitate a deeper look into the white matter structure.

Fig. 9 shows a selection of tracts in the *corpus callosum* both (a) before and (b) after contraction with $d_{max} = 5$ mm. Note that the selection of tracts only affects the rendering, not the contraction itself which is still done for all fiber tracts. Therefore, the selection provides a nice example for the sheet-like structure of the *corpus callosum* in the neighborhood of the selected tracts, apparent in particular in the rotated view in Fig. 9(c).

As stated above, the volumetric voids created by our contraction approach assist in reducing occlusion. The application of filtering based on the degree of the vertices in the similarity graph decreases the occlusion even more, see Fig. 10. As the average vertex degree is ≈ 200 (approx. $4 \cdot 10^8$ edges, $2 \cdot 10^6$ vertices) a significant reduction is achieved by filtering with a threshold as shown in Fig. 10. Fig. 10(b) and (c) show filtering with a vertex degree threshold of 160 and 560, respectively. Notice how only the major structures of the white matter remain in Fig. 10(c).

The 2D lens interaction described in Section 4.4 is shown in Fig. 11. One mode shown in Fig. 11(a) shows the tracts under the lens in the original positions while the remainder of the tracts is in its contracted state. The other mode shown in Fig. 11(b) is the reverse: now the context is in its uncontracted state and the lens allows the user to see the contracted state of the tracts inside the lens.




Fig. 11. Lens interaction with the contracted fiber tracts. The lens can have two modes, one (a) where the lens reveals the original state with the rest of the tracts in the contracted state, and one (b) that does the reverse.

6 EVALUATION AND VALIDATION

We need to understand the kinds of changes that are being introduced by the contraction process. Specifically, it is important to determine to what extent brain structures are being deformed, how much—for a given contraction scale—tracts are moved on average, and how well brain structures in the contracted state still coincide with their corresponding non-deformed fiber tracts. We thus analyzed the introduced contraction changes quantitatively as well as visually, and present the results of the analysis in this section.

6.1 Visual Evaluation of the Contraction Process

During the abstraction process, the contraction of the fiber tract paths is primarily due to a lateral movement of the vertices on each of the fiber tracts. Insight into the contraction process can thus be gained by visualizing the effects of this lateral motion. Hence, we specifically evaluate the contraction process in a 2D space that best portrays this lateral motion. For this purpose we determine and visualize the intersections of the contracted or uncontracted fiber tracts with *skull-like surfaces* at different depths in the brain, allowing us to easily perceive how the fibers behave under the contraction process. Thus we avoid the problems that would be introduced by comparing directly against 3D structures such as fiber bundles identified by clustering methods, or FA skeletons [37].

First, the approximate surface S of the skull is extracted using the upper part of a near-zero iso-surface of the FA data. The resulting triangulated surface is smoothed and the number of triangles is reduced to about 2000. Next, we determine three scaled-down instances S_i ($i = 1, 2, 3$) of S at different depths of the brain (at 4 mm, 28 mm, and 58 mm (± 3 mm), respectively; see Fig. 12). We use a point in the middle of the brain (near the brain stem; but the exact location is not too critical) as the center of the scaling to achieve both a reduction in size and the desired displacement. We then texture-map the respective local FA value in gray-scale onto S_i as a reference and then proceed to indicate the intersection points of the fiber tracts. For every tract t , the point of intersection of t and S_i is rendered as a red dot on S_i , as shown in Fig. 13.

When transitioning seamlessly from the uncontracted to the fully contracted state, the intersection points start




Fig. 10. Filtering on vertex degree in the similarity graph reduces occlusion more. Here we show tracts (a) unfiltered, (b) with a filter threshold of 160, and (c) with a filter threshold of 560.




Fig. 12. Three skull-like surfaces S_i used as reference surfaces to visualize the effects of contraction.




Fig. 13. Fiber tracts are intersected with a surface of interest S_i . The surface is textured with the local FA values and the intersection points are visualized.

moving towards each other, forming clusters of points on S_i , as shown in Fig. 14. Generally, each tract intersects a given surface S_i only once. Only sporadically does a tract which used to intersect the surface in the uncontracted state no longer intersect it in a more contracted state, or the other way around. Such cases, however, do not pose a problem due to the large number of tracts used in our case (roughly 77 400 fiber tracts).

Fig. 14 shows four stages of the contraction process for S_1 for $d_{max} = 0, 1, 3$, and 7 mm, respectively. Fig. 14(a) (i. e., $d_{max} = 0$ mm) shows the uncontracted stage in which the intersection points cover the white FA areas. Next, for $d_{max} = 1$ mm (Fig. 14(b)) the intersection points start to cluster and coincide very well with the white FA areas. At the same time, voids between contracted fiber tracts emerge. At $d_{max} = 3$ mm (Fig. 14(c)), the contraction becomes stronger but this happens at the cost of distortion, i. e., the shape of many sets of contracted fiber tracts deviates somewhat from the shape of the corresponding FA areas. Finally, at $d_{max} = 7$ mm (Fig. 14(d)) the contraction is so strong that contracted fiber tracts end up far away from their




Fig. 14. The contraction process visualized on a skull-like surface for $d_{max} = 0, 1, 3$, and 7 mm, resp. The 1 mm contraction results correspond well to the FA field, while the contraction with 3 mm and higher values results in increasingly distorted sets of fiber tracts.




Fig. 15. Histograms of the vertex displacement (to the uncontracted positions) for four contraction distances, based on a given d_{max} ((a) 1 mm, (b) 2 mm, (c) 3 mm, and (d) 7 mm) and for every tract-vertex combination.

corresponding FA areas and the shape of the contracted fiber sets deviates strongly from their corresponding FA areas.

We can conclude that for the first approx. 30% of the contraction (i. e., $0 \leq d_{max} \leq 2$ mm) the position and shape of the generated fiber tract sets correspond well with FA

TABLE 1
Displacement statistics for Fig. 15's contraction data.

d_{max} (mm)	mean (mm)	var (mm)	max (mm)
1	0.2072	0.0244	1.9573
2	1.0082	0.3088	4.2964
3	2.0089	1.1120	8.0925
7	4.5832	5.6305	19.2375

position and shape. As we increase d_{max} beyond 2 mm, fiber tract sets move increasingly to unrealistic positions and become distorted. While we only reported results for the skull-like surface S_1 , our experiments show that the same conclusions may be drawn for the contraction visualizations on S_2 and S_3 .

6.2 Quantitative Analysis

The analysis described in the previous section relies on visual inspection and, thus, is still subjective to some degree. To better understand the effects of the contraction we also analyzed the introduced displacement quantitatively.

For a number of stages of the contraction process (i. e., for $d_{max} = 1, 2, 3, 7$ mm) and for every vertex v of every fiber tract, we calculated the distance δ_v between the vertex position in the contracted and the uncontracted states. In Fig. 15 we show the histograms of these distances, grouped by 20 bins of equal width. The position of the last bin was determined by the largest observed value of δ_v for the d_{max} under consideration. Since tracts mainly move in the lateral direction, we can consider the displacement of a vertex on a fiber tract to be a measure for this lateral tract displacement. The distances in the histograms, therefore, can be interpreted as a signature of the local lateral tract displacement.

Because the shape of the histograms in Fig. 15(b)–(d) is roughly the same, we can deduce that the relative displacement distribution does not depend strongly on d_{max} , apart from a linear scaling of the displacement. The peak of each distribution is at $\approx d_{max}/2$, and the probability that a point on a fiber tract is moved by more than d_{max} is less than $\approx 20\%$. Moreover, Table 1 shows that the maximum displacement is less than $3 \cdot d_{max}$ in all cases.

6.3 Expert Feedback

We presented our fiber tract contraction approach to a neuroscientist (6 years of professional experience). His initial reaction was enthusiastic, he could clearly distinguish, for example, the *cingulum* bundle which he found clearly separated from the *corpus callosum*. However, he found the lateral parts more difficult to interpret and could not really distinguish specific bundles. To a large extent this is not surprising as we did not design our method to identify and label the anatomical fiber bundles. He did, however, appreciate the contracted visualization of a selection of tracts, as shown in Fig. 9(b), stating it nicely conveyed the sheet-like nature of several white matter structures. Furthermore, he was interested in seeing our visualization together with FA slices and would like to compare it to the FA skeleton by Smith *et al.* [37].

6.4 Focus Group with Neuroscientists

In addition to this informal feedback, we also conducted a focus group (e. g., [26]) with three neuroscientists (also, a research intern was present who did not contribute to the discussion) to discuss the benefits and problems of our technique and the utility of full-brain fiber tracking in general. One was a neuroscience researcher (R (male); 40 years old; 19 years of professional experience), while the other two were neuroscience research engineers (E1 (female) and E2 (male); ages 28 and 32 years, resp.; 4 and 7 years of professional experience, resp.). R reported to work on brain connectivity data on a weekly basis, E1 several times daily, and E2 on a monthly basis. R also works with full-brain tractography data on a weekly basis, while E1 and E2 work with it several times a year. One of the authors met with them at their research institute, first explained our approach using images and videos, and then lead a semi-structured interview on the benefits and limitations/problems of our approach. The participants were specifically instructed to voice both positive and negative comments. The session took approx. 90 minutes and was video-recorded to be able to verify and complete the written notes.

First, we addressed the usefulness of full-brain tractography in today's neurological research. The participants stated that full-brain tractography (both deterministic and probabilistic) is actively being used in connectomics research and is a fundamental technique; "it is often best practice" to get an overview of the connectivity structure (R: "It's often the first step."; E2: "This is less biased than putting your seeds somewhere and then doing tracking."). However, full-brain tractography is typically not used much for "pure visualization" purposes but as a starting point of a detailed analysis; in particular if the task is to select and analyze "meaningful clusters of the tracts" (R). So while the experts felt that our approach would probably not be applicable in much of *clinical practice* by itself (even though full-brain tractography is sometimes used in surgery as was pointed out), it would be quite useful in neurological research.

In fact, the focus group stated that our technique is *applicable to both deterministic and probabilistic tractography* since both approaches rely on tracing fiber tracts. Only some frequently employed software such as FSL² [19] does not provide views of individual fibers. Moreover, they pointed out that our technique would likely be particularly suitable to the fiber tracts generated in probabilistic fiber tracking because it is likely to show the crossing fibers well.³

Our participants also emphasized that a main benefit of our approach is its removal of noise from the large fiber tract datasets which currently prevents the use of full-brain tractography. This *removal of noise is essential* to understand the data—it would also allow researchers to check their datasets for errors which are invisible in the original data.

The noise removal and the resulting abstraction, the experts stated, thus open exciting possibilities for easy

2. <http://fsl.fmrib.ox.ac.uk/>

3. The choice of deterministic vs. probabilistic tracking is independent from that of single-fiber vs. crossing-fiber models: both crossing-fiber deterministic tracking and single-fiber probabilistic tracking exist.

fiber bundle selection (E1: “We don’t see spaghetti, we see bundles!”). After bundles are thus selected in the contracted state, they could then be analyzed in the uncontracted state. This selection in the contracted state can even be done at d_{max} ranges that would normally lead to anatomical incorrectness: at these contraction stages more noise is removed and the larger-scale bundle structure becomes even more apparent, making the selection of certain structures easier. One exciting next step would thus be to automatically extract a bundle representation from the contracted fiber tracts, i. e., an analytic three-dimensional model of bundle sheets that split and possibly merge. Such a representation would be very useful for bundle selection, in a similar way that Jianu *et al.* [21] use abstracted 2D connectivity maps.

Based on such an analytic representation, the focus group pointed out that another very promising application domain for our method would be the *registration of different patient’s datasets*. Here, high (and anatomically inaccurate) d_{max} ranges would not be a problem since larger-scale fiber bundle structures are of importance. Such an approach would be similar to MRI dataset registration techniques that rely on the (also anatomically incorrect) mapping of an individual’s brain surface on a spherical unit surface [43] for brain registration, as currently done with tools such as FreeSurfer⁴.

To address the anatomical incorrectness issue, the domain experts suggested to use the FA field in the contraction process, as opposed to only the locations of neighboring fiber tracts—such an extended approach could ensure anatomical correctness for d_{max} values of more than 2 mm. Alternatively, they suggested to explore approaches that can extract the distortion field with respect to the anatomically correct fiber tract locations and then to warp the FA field or extracted anatomical structures (e. g., tumors) accordingly. This would provide a frame of reference even for highly abstracted representations of the fiber tracts.

The focus group also emphasized (on their own accord, without being prompted by us) that it would be extremely interesting for them to apply our technique to specific local regions, as opposed to only the whole brain. They stated that, in particular, deep-brain tractography of regions such as the brain stem (for which data is available, e. g., from 7 T scans) would be interesting to be examined with our method because their structure is not yet understood.

Finally, the researchers stated that a software tool that would provide our fiber tract contraction technique as well as the functionality suggested by them would be highly useful and would be much appreciated in neuroscience research.

7 DISCUSSION

Next we discuss the evaluation and our visualization choices, report on performance, review the specific parameters that we used, and mention a number of limitations.

7.1 Evaluation/Validation Results

Both the quantitative evaluation and the subjective observations of the contraction process indicate that the abstraction

introduced by the contraction is successful. While it is difficult to formulate objective criteria for the quality of contraction without specifying the precise purpose or application domain, our results have shown that it is realistic to use the contracted representation for up to $d_{max} \approx 2$ mm for purposes that require good anatomical precision. This is a positive result since already at this small value range for d_{max} we can achieve a significant level of abstraction (see Fig. 14(b)). Beyond that value, the contracted representation can still be used, e. g., to study the connections of the brain in a more schematic manner, and can also serve as context for small groups of uncontracted tracts. Specifically the smooth interactive transition between an uncontracted and a contracted representation as enabled by our technique can enhance insight in the brain connectivity structure.

7.2 Visualization Choices

Our visualization relies on view-oriented triangle strips combined with the depth-depended-halos rendering technique [13]. Thus, individual fibers are depicted as surfaces, even though these surfaces do not exist in reality. One could argue that this defeats the purpose of revealing the spatial structure of fiber tracts, and that it would be better to use shading to give insight into the spatial structure. We think, however, that this would be a bad choice: The tracts are very close to each other and shading would require to use tubes which would heavily overlap. Moreover, we allow the user to interactively transition between the uncontracted and contracted states, thus requiring a visualization technique that also works well for depicting a dense set of individual lines. The use of depth-depended-halos combined with direction-dependent color depicts the spatial layering well, still provides cues on the direction of the fibers as can be seen, e. g., in Fig. 10(b), and facilitates an interactive transition between contracted and uncontracted stages.

7.3 Performance and Parameters

Both the graph generation and the edge contraction algorithms were written in C while the visualization program was written in Python, making extensive use of shaders, vertex arrays, and vertex buffer objects. The latter part, therefore, is interactive, but the graph generation and contraction stages are done in a pre-processing step.

The first part of this preprocessing, the graph generation process, is easy to parallelize since each pair of tracts can be compared independently. We chose a simple thread-based approach for parallelization and achieved a near-linear speedup with an increasing number of CPU cores. Nevertheless, due to the amount of data and the computational complexity of our approach, both the graph generation and iterative contraction take a considerable amount of time.

The generation of the similarity graph requires that each tract is compared to each other tract, so the number of tract comparisons is quadratic with respect to the number of tracts. For each pair, finding the nearest neighbor for each vertex requires a quadratic amount of operations with respect to the number of vertices in both tracts. This amounts to

4. <http://surfer.nmr.mgh.harvard.edu/>




Fig. 16. Similarity graph edge # as a function of d_{max} .

a computational complexity of $O(N^2M^2)$, where N is the number of tracts and M the average number of vertices per tract. We improved the computational complexity and, consequently, the run time of our implementation in three ways: (a) by using a grid search for the distance Condition 1 in Section 3.1; (b) by avoiding multiple angle calculations for segment pairs for Condition 2; and (c) by using, for Condition 3, a linear sweep algorithm running along two tracts simultaneously. We ran the naïve $O(N^2M^2)$ algorithm as well as the improved algorithm on a machine on which up to 4 Intel® Xeon® X7350 processors could be used, each with 4 cores, running at 2.93 GHz and using 128 GiB of memory. The test data set consisted of 77 389 tracts with a total of 1 944 570 vertices, and d_{sample} was set to 1 mm. On this setup and using 4 threads, the naïve graph generation implementation took 76 minutes, while the improved implementation took 15 minutes.

An iteration of the contraction process is linear in the number of edges in the similarity graph, but the number of edges in the graph is typically large. In the given example, a similarity graph with approx. $345 \cdot 10^6$ edges was generated, using a direction similarity threshold θ_{par} of 11.48° and a maximum search range d_{max} of 5 mm. With this graph we created 5 scale levels, each with d_{max} increased by 1 mm. This computation took 4 hours and 45 minutes on a single core of the aforementioned machine.

Fig. 16 shows a graph of the number of edges in the graph as a function of d_{max} . It can be observed that over a range of approximately 1–6 mm the size of the graph grows more or less linearly. This is probably caused by two opposing effects of increasing d_{max} . On average, the number of edges in the graph increases quadratically with d_{max} . But this effect is canceled out by the fact that edges with a large distance are less likely to be parallel.

7.4 Limitations

One important limitation of our technique is the computational complexity of the graph generation and iterative contraction, in terms of running time and memory usage, as reported in Section 7.3. While the improved algorithm for the similarity graph generation reduced the running time to around 15 minutes for our test data, the computation time of the iterative contraction process is still in the order of hours. Options to improve this situation include reducing the number of edges in the similarity graph, not calculating the displacement vector in each iteration, experimenting with the parameters for the contraction (e.g. scaling of the




Fig. 17. Potential topology changes during contraction.

displacement vector and the number of iterations), and restructuring the algorithm to allow for parallelization and better cache coherency.

Also, topological properties may change as a result of contraction. For example, consider two ribbon-shaped fiber groups A and B (see Fig. 17) which intersect each other perpendicularly (such cases would not occur in FACT-based DTI fiber tracking but should be considered for fiber tracking algorithms that can produce intersections). Assume that the leftmost tract b of B intersects some tract a in A in point p , i.e., a and b have the point p in common. During contraction, the tracts do not affect each other because of the perpendicular arrangement, but the width of each of the ribbons decreases while their overall spatial positions remain more or less the same. So it is reasonable to assume that the contracted ribbons A' and B' do not intersect anymore, that is, the tracts a' and b' do not have a point in common. This shows that contraction may change the property ‘ a and b have a point in common.’ In the same way, properties like ‘between’ and ‘in front of’ may change under contraction. When answering topological questions of this kind it is advisable to make a number of transitions between contracted and original state and determine what the situation is in the original state.

8 CONCLUSION

We have presented a visualization method for DTI fiber tracts that can help users to understand white matter structure through visual abstraction. It differs from previous work in that we are able to process and visualize all fiber tracts extracted for a brain, deriving an abstracted representation through neighborhood similarity analysis based on local tract orientation and iterative contraction. The resulting representation of the fiber tracts enables a better insight into the structure of the white matter of the brain. Increased voids between the contracted tracts not only reduce the occlusion normally arising from visualizing a large number of fibers but also reveal the sheet-like structure of several major white matter structures [28], [41], [42]. Even for abstractions in which the anatomical location of the fiber tracts is distorted too much, the abstract representations can be used as a context for displaying a subsection of the original, uncontracted fibers to be examined in detail. While the pre-processing requires a considerable amount of time, it only needs to be done once for a given dataset after which it is possible to explore the resulting data interactively. Moreover, the runtime of the algorithm depends on d_{max} : a lower value

means faster processing, thus using the suggested 2 mm instead of 7 mm saves a significant amount of time. Because we compute the contraction for a number of stages for increasing neighborhood sizes, users can make a seamless transition between the different scales of abstraction. In addition, we support filtering, tract selection, and a lens as interactive tools to explore the effects of the contraction.

Our technique can be applied, in particular, in neurological research and the exploration of new datasets. For example, selected subsets of fibers with their correct anatomy can be examined in the abstracted context of all fibers. Moreover, feedback from domain experts in the neurosciences suggests that our approach can be useful for error checking of datasets and for effective fiber bundle selection and registration.

Future work includes improving the running time of the pre-processing steps, as mentioned in Section 7.4. This would also make it easier to further explore the effect of parameter settings on the contraction. Furthermore, our method needs more validation with the help of neuroscientists. Finally, we are planning to combine our contraction approach with fiber tract clustering techniques, as reviewed in Section 2, which could yield interesting results.

ACKNOWLEDGMENTS

We thank Cris Lanting and Pim van Dijk from the NeuroImaging Center in Groningen, the Netherlands, for the brain dataset used in this paper. We also thank Klaas Mussche and Jorn van de Beek for designing and testing the improved methods for similarity graph generation. We gratefully acknowledge Alessandro Crippa, Leonardo Cerliani, and Arjen van Hulzen for assistance and feedback. We also thank the anonymous reviewers for their useful comments. This research was funded by the Netherlands Organisation for Scientific Research (NWO), “VIEW” program, project no. 643.100.501. Tobias Isenberg’s work was supported, in part, by a DIGITEO chair of excellence in France.

REFERENCES

- [1] D. Akers, A. Sherbondy, R. Mackenzie, R. Dougherty, and B. Wandell. Exploration of the brain’s white matter pathways with dynamic queries. In *Proc. IEEE Visualization*, pp. 377–384. IEEE Computer Society, Los Alamitos, 2004. doi> 10.1109/VISUAL2004.30
- [2] P. J. Basser, J. Mattiello, and D. LeBihan. Estimation of the effective self-diffusion tensor from the NMR spin-echo. *Journal of Magnetic Resonance*, 103(3):247–254, 1994.
- [3] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance B*, 111(3):209–219, June 1996. doi> 10.1006/jmrb.1996.0086
- [4] E. A. Bier, M. C. Stone, K. Pier, W. Buxton, and T. D. DeRose. Toolglass and magic lenses: The see-through interface. In *Proc. SIGGRAPH*, pp. 73–80. ACM, New York, 1993. doi> 10.1145/166117.166126
- [5] J. Blaas, C. P. Botha, B. Peters, F. M. Vos, and F. H. Post. Fast and reproducible fiber bundle selection in DTI visualization. In *Proc. IEEE Visualization*, pp. 59–64. IEEE Computer Society, Los Alamitos, 2005. doi> 10.1109/VISUAL2005.1532778
- [6] J. Böttger, A. Schäfer, G. Lohmann, A. Villringer, and D. S. Margulies. Three-dimensional mean-shift edge bundling for the visualization of functional connectivity in the brain. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):471–480, Mar. 2014. doi> 10.1109/TVCG.2013.114
- [7] A. Brun, H. Knutsson, H.-J. Park, M. E. Shenton, and C.-F. Westin. Clustering fiber tracts using normalized cuts. In *Proc. MICCAI*, part I, volume 3216 of *LNCIS*, pp. 368–375. Springer, Berlin, 2004. doi> 10.1007/978-3-540-30135-6_45
- [8] W. Chen, Z. Ding, S. Zhang, A. MacKay-Brandt, S. Correia, H. Qu, J. A. Crow, D. F. Tate, Z. Yan, and Q. Peng. A novel interface for interactive exploration of DTI fibers. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1433–1440, Nov./Dec. 2009. doi> 10.1109/TVCG.2009.112
- [9] I. Corouge, G. Gerig, and S. Gouttard. Towards a shape model of white matter fiber bundles using diffusion tensor MRI. In *Proc. ISBI*, pp. 344–347. IEEE, Los Alamitos, 2004. doi> 10.1109/ISBI.2004.1398545
- [10] A. Crippa and J. B. T. M. Roerdink. Data-driven visualization of functional brain regions from resting state fMRI data. In *Proc. VMV*, pp. 247–254. Eurographics Association, Goslar, Germany, 2011. doi> 10.2312/PE/VMV/VMV11/247-254
- [11] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall, Upper Saddle River, NJ, USA, 1999.
- [12] Z. Ding, J. Gore, and A. Anderson. Classification and quantification of neuronal fiber pathways using diffusion tensor MRI. *Magnetic Resonance in Medicine*, 49(4):716–721, Apr. 2003. doi> 10.1002/mrm.10415
- [13] M. H. Everts, H. Bekker, J. B. T. M. Roerdink, and T. Isenberg. Depth-dependent halos: Illustrative rendering of dense line data. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1299–1306, Nov./Dec. 2009. doi> 10.1109/TVCG.2009.138
- [14] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns. Mapping the structural core of human cerebral cortex. *PLoS Biology*, 6(7):1479–1493, July 2008. doi> 10.1371/journal.pbio.0060159
- [15] M. Hlawitschka, C. Garth, X. Tricoche, G. Kindlmann, G. Scheuermann, K. I. Joy, and B. Hamann. Direct visualization of fiber information by coherence. *International Journal of Computer Assisted Radiology and Surgery*, 5(2):125–131, Apr. 2010. doi> 10.1007/s11548-009-0302-5
- [16] D. Holtén. Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):741–748, Sept./Oct. 2006. doi> 10.1109/TVCG.2006.147
- [17] D. Holtén and J. J. van Wijk. Force-directed edge bundling for graph visualization. *Computer Graphics Forum*, 28(3):983–990, June 2009. doi> 10.1111/j.1467-8659.2009.01450.x
- [18] T. Isenberg. A survey of illustrative visualization techniques for DTI-based fiber tracking. In T. Schulz and I. Hotz, eds., *Visualization and Processing of Tensors and Higher Order Descriptors for Multi-Valued Data*, chap. 12. Springer, Berlin/Heidelberg, 2015. To appear.
- [19] M. Jenkinson, C. F. Beckmann, T. E. Behrens, M. W. Woolrich, and S. M. Smith. FSL. *NeuroImage*, 62(2):782–790, Aug. 2012. doi> 10.1016/j.neuroimage.2011.09.015
- [20] R. Jianu, C. Demiralp, and D. Laidlaw. Exploring 3D DTI fiber tracts with linked 2D representations. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1449–1456, Nov./Dec. 2009. doi> 10.1109/TVCG.2009.141
- [21] R. Jianu, C. Demiralp, and D. H. Laidlaw. Exploring brain connectivity with two-dimensional neural maps. *IEEE Transactions on Visualization and Computer Graphics*, 18(6):978–987, June 2012. doi> 10.1109/TVCG.2011.82
- [22] S. E. Jones, B. R. Buchbinder, and I. Aharon. Three-dimensional mapping of cortical thickness using Laplace’s equation. *Human Brain Mapping*, 11(8):12–32, Aug. 2000. doi> 10.1002/1097-0193(200009)11:1<12::AID-HBM20>3.0.CO;2-K
- [23] G. Kindlmann, X. Tricoche, and C.-F. Westin. Delineating white matter structure in diffusion tensor MRI with anisotropy creases. *Medical Image Analysis*, 11(5):492–502, Oct. 2007. doi> 10.1016/j.media.2007.07.005
- [24] G. Kindlmann and C.-F. Westin. Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1329–1336, Sept./Oct. 2006. doi> 10.1109/TVCG.2006.134
- [25] J. Klein, F. Ritter, H. K. Hahn, J. Rexilius, and H.-O. Peitgen. Brain structure visualization using spectral fiber clustering. In *Research Posters of ACM SIGGRAPH*, p. 168:1. ACM, New York, 2006. doi> 10.1145/1179622.1179816
- [26] J. Lazar, J. H. Feng, and H. Hochheiser. Interviews and focus groups. In *Research Methods in Human-Computer Interaction*, chap. 8, pp. 177–216. John Wiley & Sons Ltd, 2010.

- [27] M. Maddah, W. E. L. Grimson, S. K. Warfield, and W. M. Wells. A unified framework for clustering and quantitative analysis of white matter fiber tracts. *Medical Image Analysis*, 12(2):191–202, Apr. 2008. doi> 10.1016/j.media.2007.10.003
- [28] M. Maddah, J. V. Miller, E. V. Sullivan, A. Pfefferbaum, and T. Rohlfing. Sheet-like white matter fiber tracts: Representation, clustering, and quantitative analysis. In *Proc. MICCAI*, part II, volume 6892 of *LNCS*, pp. 191–199. Springer, Berlin, 2011. doi> 10.1007/978-3-642-23629-7_24
- [29] D. Merhof, M. Sonntag, F. Enders, V. P. Hastreiter, R. Fahlbusch, C. Nimsky, and G. Greiner. Visualization of diffusion tensor data using evenly spaced streamlines. In *Proc. Vision, Modeling and Visualization*, pp. 79–86. IOS Press, Amsterdam, 2005.
- [30] B. Moberts, A. Vilanova, and J. J. van Wijk. Evaluation of fiber clustering methods for diffusion tensor imaging. In *Proc. IEEE Visualization*, pp. 65–72. IEEE Computer Society, Los Alamitos, 2005. doi> 10.1109/VS.2005.29
- [31] S. Mori and P. C. M. van Zijl. Fiber tracking: Principles and strategies – a technical review. *NMR in Biomedicine*, 15(7–8):468–480, Nov./Dec. 2002. doi> 10.1002/nbm.781
- [32] L. O'Donnell and C.-F. Westin. Automatic tractography segmentation using a high-dimensional white matter atlas. *IEEE Transactions on Medical Imaging*, 26(11):1562–1575, Nov. 2007. doi> 10.1109/TMI.2007.906785
- [33] S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42(3):526–40, Sept. 1999. doi> 10.1002/(SICI)1522-2594(199909)42:3<526::AID-MRM15>3.0.CO;2J
- [34] V. Petrovic, J. Fallon, and F. Kuester. Visualizing whole-brain DTI tractography with GPU-based tuboids and LoD management. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1488–1495, Nov./Dec. 2007. doi> 10.1109/TVCG.2007.70622
- [35] T. Schultz, N. Sauber, A. Anwender, H. Theisel, and H.-P. Seidel. Virtual Klingler dissection: Putting fibers into context. *Computer Graphics Forum*, 27(3):1063–1070, May 2008. doi> 10.1111/j.1467-8659.2008.01243.x
- [36] T. Schultz, H. Theisel, and H.-P. Seidel. Topological visualization of brain diffusion MRI data. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1496–1503, Nov./Dec. 2007. doi> 10.1109/TVCG.2007.70602
- [37] S. M. Smith, M. Jenkinson, H. Johansen-Berg, D. Rueckert, T. E. Nichols, C. E. Mackay, K. E. Watkins, O. Ciccarelli, M. Z. Cader, P. M. Matthews, and T. E. Behrens. Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. *Neuroimage*, 31(4):1487–1505, July 2006. doi> 10.1016/j.neuroimage.2006.02.024
- [38] P. Svetachov, M. H. Everts, and T. Isenberg. DTI in context: Illustrating brain fiber tracts in situ. *Computer Graphics Forum*, 29(3):1024–1032, June 2010. doi> 10.1111/j.1467-8659.2009.01692.x
- [39] A. Tsai, C.-F. Westin, A. O. Hero, and A. S. Willsky. Fiber tract clustering on manifolds with dual rooted-graphs. In *Proc. CVPR*. IEEE Computer Society, Los Alamitos, 2007. doi> 10.1109/CVPR.2007.383096
- [40] A. Vilanova, G. Berenschot, and C. van Pul. DTI visualization with stream surfaces and evenly-spaced volume seeding. In *EG/IEEE TCVG Symposium on Visualization*, pp. 173–182. Eurographics Association, Goslar, Germany, 2004.
- [41] S. Wakana, H. Jiang, L. M. Nagae-Poetscher, P. C. M. van Zijl, and S. Mori. Fiber tract-based atlas of human white matter anatomy. *Radiology*, 230(1):77–87, Jan. 2004. doi> 10.1148/radiol.2301021640
- [42] V. J. Wedeen, D. L. Rosene, R. Wang, G. Dai, F. Mortazavi, P. Hagmann, J. H. Kaas, and W.-Y. I. Tseng. The geometric structure of the brain fiber pathways. *Science*, 335(6076):1628–1634, Mar. 2012. doi> 10.1126/science.1215280
- [43] B. Yeo, M. Sabuncu, T. Vercauteren, N. Ayache, B. Fischl, and P. Golland. Spherical demons: Fast diffeomorphic landmark-free surface registration. *IEEE Transactions on Medical Imaging*, 29(3):650–668, Mar. 2010. doi> 10.1109/TMI.2009.2030797
- [44] H. Yu, C. Wang, C.-K. Shene, and J. H. Chen. Hierarchical streamline bundles. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1353–1367, Aug. 2012. doi> 10.1109/TVCG.2011.155
- [45] S. Zhang, S. Correia, and D. H. Laidlaw. Identifying white-matter fiber bundles in DTI data using an automated proximity-based fiber-clustering method. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1044–1053, Sept./Oct. 2008. doi> 10.1109/TVCG.2008.52
- [46] S. Zhang, C. Demiralp, and D. H. Laidlaw. Visualizing diffusion tensor MR images using streamtubes and streamsurfaces. *IEEE*

Transactions on Visualization and Computer Graphics, 9(4):454–462, Oct.–Dec. 2003. doi> 10.1109/TVCG.2003.1260740



Maarten Everts received his M.Sc. (2006) and Ph.D. (2011) degrees in computing science at the University of Groningen. Currently he is employed by TNO, The Netherlands. His research interests include computer graphics, visualization, and computer security.



Eric Begue received his M.Sc. degree in computing science at the University of Groningen in 2011. Currently he is a game programmer at ISOTX, Groningen Area, the Netherlands.



Henk Bekker received his Ph.D. degree in 1996 at the University of Groningen in the area of molecular dynamics simulations. Currently, he is an assistant professor at the Johann Bernoulli Institute for Mathematics and Computer Science of the University of Groningen. His research interests include modelling and simulation, computational geometry, and visualization.



Jos B. T. M. Roerdink received a Ph.D. in theoretical physics (1983) from the University of Utrecht, the Netherlands. After a postdoctoral position (1983-1985) at the University of California, San Diego, he joined the Centre for Mathematics and Computer Science in Amsterdam, working on image processing and tomographic reconstruction. He was appointed associate professor (1992) and full professor (2003), respectively, at the Johann Bernoulli Institute for Mathematics and Computer Science of the University of Groningen, where he currently holds a chair in Scientific Visualization and Computer Graphics. He also has a joint appointment with the University Medical Center Groningen.



Tobias Isenberg is a senior research scientist with the AVIZ project team at Inria-Saclay, France. He received his doctoral degree from the University of Magdeburg, Germany, in 2004. Previously he held positions as postdoctoral fellow at the University of Calgary, Canada, and as assistant professor at the University of Groningen, the Netherlands. His research interests comprise topics in scientific visualization, in illustrative and non-photorealistic rendering, and interactive visualization techniques and he is particularly interested how abstraction can benefit illustrative visualization.