

Segmentation of DT-MRI Anisotropy Isosurfaces

Thomas Schultz¹, Holger Theisel², and Hans-Peter Seidel¹

¹ MPI Informatik, Saarbrücken, Germany — {schultz, hpseidel}@mpi-inf.mpg.de

² BieGraph Group, Bielefeld University, Germany — theisel@techfak.uni-bielefeld.de

Abstract

While isosurfaces of anisotropy measures for data from diffusion tensor magnetic resonance imaging (DT-MRI) are known to depict major anatomical structures, the anisotropy metric reduces the rich tensor data to a simple scalar field. In this work, we suggest that the part of the data which has been ignored by the metric can be used to segment anisotropy isosurfaces into anatomically meaningful regions. For the implementation, we propose an edge-based watershed method that adapts and extends a method from curvature-based mesh segmentation [MW99]. Finally, we use the segmentation results to enhance visualization of the data.

Categories and Subject Descriptors (according to ACM CCS): I.4.6 [Image Processing and Computer Vision]: Region growing, partitioning

1. Introduction

Diffusion-tensor magnetic resonance imaging (DT-MRI) is a medical imaging modality that measures the self-diffusion of water molecules. This allows conclusions about the microstructure of the tissue that restricts molecular movement [BP96]. In each voxel, the method determines a symmetric 3×3 matrix (a *diffusion tensor*) that models the observed distribution of apparent diffusivities.

Most frequently, DT-MRI is applied to the human brain. The high complexity of such diffusion tensor fields has led to the development of a number of complementary visualization techniques, each conveying some portion of the data (cf. book chapters [VZKL06, ZKL05]).

One of these techniques are isosurfaces of scalar anisotropy measures, which give a large-scale overview of the data. It has been shown that these surfaces allow one to identify important anatomic structures within the brain [ZMB⁺03] and they have subsequently been made available in packages for DT-MRI visualization (e.g., [Sci06]). In the present work, we suggest a method for automatic segmentation of these isosurfaces into regions that correspond to anatomic units revealed by the data. Subsequently, we use the results to enhance visualization of the data.

Examining anisotropy isosurfaces is made difficult by the fact that they are complex and convoluted, so structures deep within the brain are occluded by structures further outside

from most viewpoints. User-defined clipping planes are suitable for cutting along any of the coordinate planes; however, they do not offer a natural and convenient tool to eliminate outer parts of the surface.

Reducing the visibility of outer structures by changing the isovalue will not in general give acceptable results, as it can also change the appearance of inner structures significantly. Segmenting the surface easily allows the user to view only the parts she is interested in.

Moreover, coloring the extracted isosurface mesh according to the regions helps to visually identify anatomically relevant parts at first glance and choosing region colors based on an average tensor representative reveals if a region is uniform in orientation and tensor shape.

In the following section, we will discuss existing methods for mesh segmentation and relate them to our own work. Section 3 gives a review of the anisotropy measures from which the isosurfaces are extracted. Section 4 formalizes our segmentation criterion, which is then used by the algorithm described in Section 5. Results are presented in Section 6 before the paper concludes in Section 7.

2. Related Work

One fundamental idea in our work is to use the magnitude of directional derivatives in the tensor field as a criterion to segment anisotropy isosurfaces. While previous work

has confirmed that the magnitude of the tensor field gradient indicates semantically significant boundaries in the data [PAB02, OGW04, Kin06] and can be used to steer geodesic active contours [FWB03], such differential information has so far not been used to segment isosurface meshes. Moreover, we consider directional derivatives rather than the full gradient magnitude, which is critical for our results.

Watershed-type segmentation methods are steered by a scalar “height” field. Since the derivative magnitude can be used to define such a scalar field on the isosurface, they offer themselves as a natural candidate for the implementation of our idea. The watershed principle has previously been used for surface segmentation by Mangan and Whitaker [MW99] and several authors after them (cf. [CG06] and references therein). However, these works have a completely different background: They aim at segmenting general surface meshes according to their geometric properties, while our method depends on the underlying tensor field and finds regions which are not defined by surface geometry.

Both Rettmann et al. [RHHP02] and Vivodtzev et al. [VLB*03] have segmented isosurfaces from cranial MRI. Their approaches are based on geodesic depth and curvature, respectively, which are suitable to segment the gyri and sulci of the cortical surface, but do not allow segmentation of the major white matter structures depicted by anisotropy isosurfaces. Our method draws on information about fiber orientation, which is specific to DT-MRI and not present in the conventional MRI data these two works deal with.

All watershed-based mesh segmentation methods we are aware of define a height field on the mesh vertices to steer the segmentation process. For our application, we found it more advantageous to define edge weights, as discussed in Section 4. While the work by Page et al. [PKA03] belongs to the group of geometry-based methods, it comes closest to our approach in that it uses a *directional* height field to control the final step of the segmentation. However, at that stage, all regions have already been found and 70–90% of the vertices have been labeled based on vertex-specific information, so beside its distinct segmentation criterion, our work can be regarded as a contribution to mesh segmentation in general.

As an alternative to anisotropy isosurfaces, direct volume rendering of DT-MRI data has been suggested [KWH00]. While that approach is not limited to depicting isosurfaces of anisotropy metrics, it has been used to produce results that are comparable to renderings of isosurfaces [VZKL06] and share their problems regarding occlusions. In this work, we concentrate on isosurfaces because our segmentation method requires explicit geometry.

3. Measures of Anisotropy

Tensor fields on \mathbf{R}^3 in general do not have isosurfaces. Instead, the surfaces we consider in this paper are isosur-

faces of a scalar field that results from the application of an anisotropy metric.

Anisotropy is a measure of the degree to which the apparent diffusivity in a voxel is directionally dependent. Alternatively, anisotropy can be thought of as the degree to which the diffusion ellipsoid deviates from a sphere. In the human brain, high anisotropy indicates coherently organized nerve fibers, so anisotropy isosurfaces outline the contours of major white matter tracts.

In the context of DT-MRI, fractional anisotropy (FA) [BP96] is the most widely used measure of anisotropy. Let \mathbf{D} denote a diffusion tensor and $\bar{\mathbf{D}} := \mathbf{D} - \frac{1}{3}\text{tr}(\mathbf{D})\mathbf{I}$ be its anisotropic part, where tr is the matrix trace. Then, the FA is defined as the relative magnitude of $\bar{\mathbf{D}}$ with respect to the Frobenius norm $\|\mathbf{D}\| := \sqrt{\text{tr}(\mathbf{D}\mathbf{D}^T)}$, rescaled to lie in $[0, 1]$:

$$\text{FA} := \sqrt{\frac{3}{2}} \frac{\|\bar{\mathbf{D}}\|}{\|\mathbf{D}\|} \quad (1)$$

However, the FA ignores the *mode* of anisotropy, i.e., it does not tell us whether the diffusion ellipsoid has a more linear or a more planar shape. The geometrical measures suggested by Westin et al. [WPG*97] are a frequent choice when this additional information is required. They are interpreted as coordinates c_l , c_p and c_s of the tensor relative to the extremal cases of linear, planar, or spherical shape, respectively. Their most common definition in terms of the sorted eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and the eigenvalue sum $\bar{\lambda} := \lambda_1 + \lambda_2 + \lambda_3$ of \mathbf{D} is:

$$c_l := \frac{\lambda_1 - \lambda_2}{\bar{\lambda}} \quad c_p := \frac{2(\lambda_2 - \lambda_3)}{\bar{\lambda}} \quad c_s := \frac{3\lambda_3}{\bar{\lambda}} \quad (2)$$

The isosurfaces described in this paper are obtained by evaluating FA or, respectively, c_l for each tensor in the dataset and running the standard Marching Cubes algorithm [LC87] on the resulting scalar field.

4. Definition of Edge Weights

Our approach exploits the fact that high local contrast in the data indicates an anatomical boundary [PAB02], where local contrast is measured as the magnitude of a directional derivative. In tensor fields, directional derivatives are again tensors, and their magnitude is given by the Frobenius norm.

According to this, two nodes on the isosurface are likely to belong to different regions if the integrated magnitude of the directional derivative along their common edge is high. We will refer to this value as the edge weight w

$$w := \int_{\text{edge}} \left\| \frac{\partial \mathbf{D}(\mathbf{r}(t))}{\partial \mathbf{e}} \right\| ds \quad (3)$$

where \mathbf{e} is the direction of the edge, which is parametrized as $\mathbf{r}(t)$. In practice, we approximate the integral by evaluating the magnitude of the directional derivative at the center of the edge, and multiplying the result with the edge length.

Tensor field derivatives are obtained by convolution with the derivative of a C^2 piecewise-cubic reconstruction kernel.

Since Equation (3) takes the integral of the derivative magnitude along the edge rather than its average, it induces a bias towards clustering geometrically close vertices together. This stabilizes the segmentation in cases where the Marching Cubes algorithm generated extremely short edges.

In cerebrospinal fluid (CSF), the apparent diffusion coefficients are much higher than within tissue, leading to large diffusion tensor traces. Varying amounts of partial voluming with CSF at ventricle boundaries lead to spurious derivative magnitudes that dominate those found in the white matter structures we are primarily interested in, which renders Equation (3) unusable in such locations. However, tensor trace has been reported to be approximately constant over functional tissue [PJB*96], so it is safe to filter out this disturbance, without running the risk of losing relevant information elsewhere. Like previous authors [PAB02], we achieve this by considering the derivative of the trace-free deviatoric $\bar{\mathbf{D}} := \mathbf{D} - \frac{1}{3}\text{tr}(\mathbf{D})\mathbf{I}$. By linearity of matrix trace, it can be obtained by subtracting from the derivative tensor \mathbf{D}' its isotropic part $\bar{\mathbf{D}}' = \mathbf{D}' - \frac{1}{3}\text{tr}(\mathbf{D}')\mathbf{I}$.

Taking directional derivatives along FA or c_l isosurfaces implicitly excludes variations in these measures, so our segmentation exploits the part of the DT-MRI data which has been ignored by the anisotropy metric. Since we explicitly eliminate the influence of tensor trace, the only remaining degrees of freedom are those anisotropy changes which are not captured by our metric (i.e., changes in anisotropy mode for FA or, respectively, in planar anisotropy for c_l) and, most of all, changes in orientation. This observation makes it plausible that the segmentation criterion presented here will yield anatomically meaningful results, since sharp changes in fiber orientation are likely to indicate an anatomical boundary.

As in scalar fields, directional derivatives in tensor fields can be determined from the gradient via the chain rule. The gradient $\nabla\mathbf{D}$ of a tensor field is a third-order tensor, which can be thought of as a three-vector of its partial derivatives $\mathbf{D}^{(x)}$, $\mathbf{D}^{(y)}$, and $\mathbf{D}^{(z)}$. A directional derivative is obtained by taking the inner vector-tensor product with a unit-length vector $\mathbf{e} = (e_x, e_y, e_z)^T$ that defines the desired direction. Combined with the filtering of tensor trace, the formula for the traceless directional derivative magnitude M reads

$$M = \left\| \sum_{i \in \{x,y,z\}} e_i \left[\mathbf{D}^{(i)} - \frac{1}{3}\text{tr}(\mathbf{D}^{(i)})\mathbf{I} \right] \right\| \quad (4)$$

We first tried to apply a variant of the algorithm in [MW99] to our problem, which requires the definition of a height field on the mesh vertices. According to the same basic segmentation criterion, such a field can be defined from the derivatives in any two orthogonal directions perpendicular to the surface normal. However, there are vertices at which the derivative magnitude varies considerably for dif-

ferent directions within the tangent plane, so we may want such vertices to be included in the same region as *specific* neighbors. For this reason, the novel segmentation algorithm described in the following section performs better than existing watershed-based methods that rely on vertex heights.

5. Performing the Segmentation

Like most watershed-based algorithms, our segmentation method finds initial regions based on local minima, followed by a merging process to reduce the oversegmentation which typically occurs in the first stage. Additionally, our algorithm requires a third stage in which it removes undesired “noise” regions that have not been matched by the merging criterion.

5.1. Finding Initial Regions

The idea behind the initial step is that each vertex most likely lies in the same region as the vertex to which it is connected by the edge of least weight w . From that neighbor, we recursively traverse the mesh until we either meet a vertex that already has a label, which can then be copied to all vertices on our path, or until we enter a vertex through its minimal edge, in which case a new region is formed.

For clarity, the explanation above adopts the *top-down* formulation which is frequently met in the watershed literature [MW99, PKA03]. Our actual implementation uses the bottom-up equivalent, starting at edges that are minimal for both their vertices and recursively labeling all nodes that are reached via their minimal edge. This choice is a mere matter of taste, as the results of both formulations will be the same.

This step of our algorithm is very similar to the original method by Mangan and Whitaker [MW99], except that it uses edge weights to decide on the transitions. On a typical isosurface mesh with 44k vertices (half of which is shown in Figure 3), this step produces almost 13k regions.

5.2. Merging Close Regions

Watershed methods frequently use closely related criteria for the initial oversegmentation and the merging. For example, previous authors [MW99, RHHP02] have simply merged regions based on the lowest point on their common boundary. In our context, high edge weights along a large part of a boundary are a good indicator that the adjacent regions are distinct, even if they share one or two cheaper edges. Consequently, we decide to integrate the directional derivative magnitudes over the full boundary.

Figure 1 clarifies the terminology of the following section: On a triangular mesh, a *boundary line* is formed by the edges that connect the nodes at the boundary of a region. The triangle strip between the boundary lines of any two adjacent regions is what we call their *boundary*. Any triangle for which exactly two of its vertices belong to the same region is




Figure 1: In our terminology, the area between any two adjacent regions (shades of gray) is called their boundary (white). Boundary lines are shown dashed, the mid-boundary line is strong black. A single boundary triangle has been hatched.

called a *boundary triangle*. The line segment that connects the midpoints of its two edges within the boundary is part of the *mid-boundary line*.

The integrated weight W along a boundary is given by

$$W := \iint_{\text{boundary}} \left[\sum_{i=1}^2 \left\| \frac{\partial \mathbf{D}(\mathbf{b}(s,t))}{\partial \mathbf{v}_i} \right\|^2 \right]^{\frac{1}{2}} dA \quad (5)$$

where $\mathbf{b}(s,t)$ parametrizes the surface between the boundary lines and the \mathbf{v}_i are any two unit vectors that span the surface element dA . In practice, we again filter the derivatives, as described by Equation (4).

We approximate the integral in (5) for each boundary triangle by evaluating the directional derivative magnitudes at the midpoint of each of its edges, averaging the results and multiplying them with the area of the triangle. Among the sampling schemes we have tried, this one was the most reliable, presumably because it places sampling points along the mid-boundary line, at the center of the boundary. Also, it can be efficiently implemented by storing the tensor field gradients which have been calculated in the first stage as edge attributes and only re-evaluating the inner vector-tensor products, as well as the Frobenius norm of the result.

The length of boundary lines is not an indicator of region saliency: In particular, distinct regions can meet along short boundary lines. Consequently, we define the effective weight $W' := W/l$, where l is the length of the mid-boundary line.

Junctions are triangles for which each vertex belongs to a different region, so they are places in which three boundaries meet. We distribute the information from junctions to the three adjacent boundaries by adding one third of the triangle's weight, area and mid-boundary length to each of them. While the integral in (5) can be evaluated as before, the length of the mid-boundary line is now estimated by averaging over both possible configurations (cf. Figure 2).

Since the boundaries we integrate over are changed by merges, the order of merge operations plays a significant role




Figure 2: In a junction, all three vertices belong to different regions (central triangle). Thus, there are two possible continuations of each mid-boundary line, shown here dotted for the one coming from below.

for the final result. We keep all boundaries in a priority queue and iteratively merge the lightest one, until a user-defined threshold on W' is reached.

The merging step is the core of our method, reducing the 13k initial regions in the above example to 623 more meaningful ones. In terms of mesh segmentation in general, it is also the main novelty of our approach. We expect that it can be transferred to other segmentation goals, as long as a reasonable definition of edge weights can be given.

5.3. Cleaning up

While watershed approaches that are based on a height field have an inherent notion of a basin depth, which relates boundary height to the minimal height within the region and is frequently used as a merging criterion [MW99, RHHP02], our described algorithm exclusively considers absolute edge weights along the boundary of a region. Consequently, for regions that are fully contained in areas of high gradient magnitude, none of the boundaries are ever selected in the merging step, and are left as noise at the end of that stage.

We eliminate this remaining oversegmentation in a final step, by forcing a merge of all regions for which the number of internal edges is lower than the number of edges that form the boundary. This condition successfully treats both cases of noise regions encountered in practice, i.e., very small regions and thin and elongated regions along the boundary of two larger ones, without introducing any new parameters. In all cases, large and legitimate regions were preserved.

The merging partner for the noise regions was again selected using the lowest effective weight W' . However, this choice is not critical, since it only affects minor regions. In our example, the final step removed 536 out of the 623 regions left after merging.

5.4. Notes on the Implementation

To implement the described segmentation method in an efficient manner, we traverse the mesh only once, for the initial

labeling. The subsequent merging and cleaning stages are performed by keeping a simple data structure for each region and boundary, respectively. For each region, we store the number of internal and boundary edges, a list of all boundary indices, and a list of all regions that have been merged into this region. For each boundary, we store the associated region and edge indices, as well as the accumulated weights W , and mid-boundary line length l .

Merging itself is performed using a priority queue that holds all boundary indices and their effective weights W' . We maintain an array that tells us for each region of the initial oversegmentation, to which of the remaining regions it is merged. If a third region is adjacent to both regions we are about to merge, we also need to merge the corresponding boundaries, which requires the priority queue to allow efficient updates of non-top elements. In addition, some care has to be taken as a merge may turn a former junction into a boundary triangle and the boundary information has to be updated accordingly. This case can be checked using the edge list of the boundary that caused the merge.

Our prototype implementation uses the C++ standard template library (STL) for most required data structures, and the teem library for taking derivatives in the tensor field. For typical datasets, it goes through the whole segmentation pipeline, from computation of the derivatives to output of the final segmented mesh, in less than three seconds on a 2 GHz Athlon 64 processor. This timing holds both for the example discussed above (44k vertices), and for all other results presented in the following section.

6. Results

Anisotropy isosurfaces typically consist of one large, central surface and several smaller ones around it. To avoid visual clutter, we have performed a connected component analysis and only retained the main component in all our examples. Thus, the presented regions have all been found by our segmentation method.

Like previous authors [ZMB*03, VZKL06], we have experimented with different values to find an isolevel which clearly depicts the structures of interest. However, since both FA and c_l are normalized to the range $[0, 1]$, similar isovalues give comparable results on different datasets.

6.1. Segmentation Results and Region-based Clipping

Figure 3 presents a sample result from running our algorithm on an isosurface at FA=0.5. In Subfigure (b), we assign a random pseudocolor to each region in order to emphasize the boundaries found by our method. A comparison to images in a brain atlas [SSS06] suggests that our segmentation correctly captures a number of anatomic structures. For example, the *cingulum bundle* (Cing) is clearly separated from the *corpus callosum* (CC), and the *internal capsule* (IC) is recognized as a region of its own.

Note that the *corpus callosum* has not been separated from the *corona radiata* (CR), reflecting the fact that there is a smooth transition between both structures: In fact, the fibers from the *internal capsule* which fan out in the *corona radiata* are intermingled with fibers that pass the *corpus callosum*. Thus, a purely boundary-based segmentation method cannot distinguish them.

In Figure 4, we demonstrate how the segmentation results can be used to clip parts of the surface that may occlude parts we are more interested in. Subfigure (a) displays the largest connected component of an FA isosurface, while the user has selected the region that corresponds to the *corpus callosum* and the *corona radiata* with the mouse and clipped all other regions to get a better view on the ventral part of the CC, which is mostly hidden in (a).

6.2. Coloring Regions by Representative

Once we have obtained a segmentation, it is possible to compute one tensor per region that represents its average diffusion behaviour. We determine this representative as a weighted average of the data within the region, where the tensor at each vertex is weighted by the sum of areas of all adjacent triangles that belong to the same region.

Consequently, the region color can be chosen based on its representative. We apply the standard XYZ-RGB scheme in which the absolute values of the x , y , and z components from the principal eigenvector are mapped to the red, green and blue color channels, respectively. Since the principal eigenvector is only well-defined if the largest eigenvalue is sufficiently distinct from the second largest one, the resulting color must be modulated with a measure of linear anisotropy.

We decide to scale saturation with the square root of c_l . While using c_l itself is more common, averaging typically leads to more-or-less isotropic tensors, which result in undistinguishable region colors when using standard schemes. Even with this modification, our colors appear fainter than what the viewer may be used to from looking at DT-MRI data, but we chose to preserve this effect as a visual reminder of the fact that we are dealing with average tensors, whose direction is inherently less clearly defined than the direction of a single data point may be.

In order to ensure expressive region representatives which are well-suited for coloring, we modify the termination criterion for the segmentation: In addition to providing a threshold on the effective boundary weight W' , we now specify a lower threshold for the linear anisotropy c_l of the region representative that would result from a merge. If it is too low, we are no longer able to assign a clear color to the resulting region. Consequently, the selected boundary is removed from the queue without causing a merge. Otherwise, our algorithm is carried out as described in Section 5.

Since changes in tensor orientation dominate our segmentation criterion (cf. Section 4), it is unlikely that two adjacent

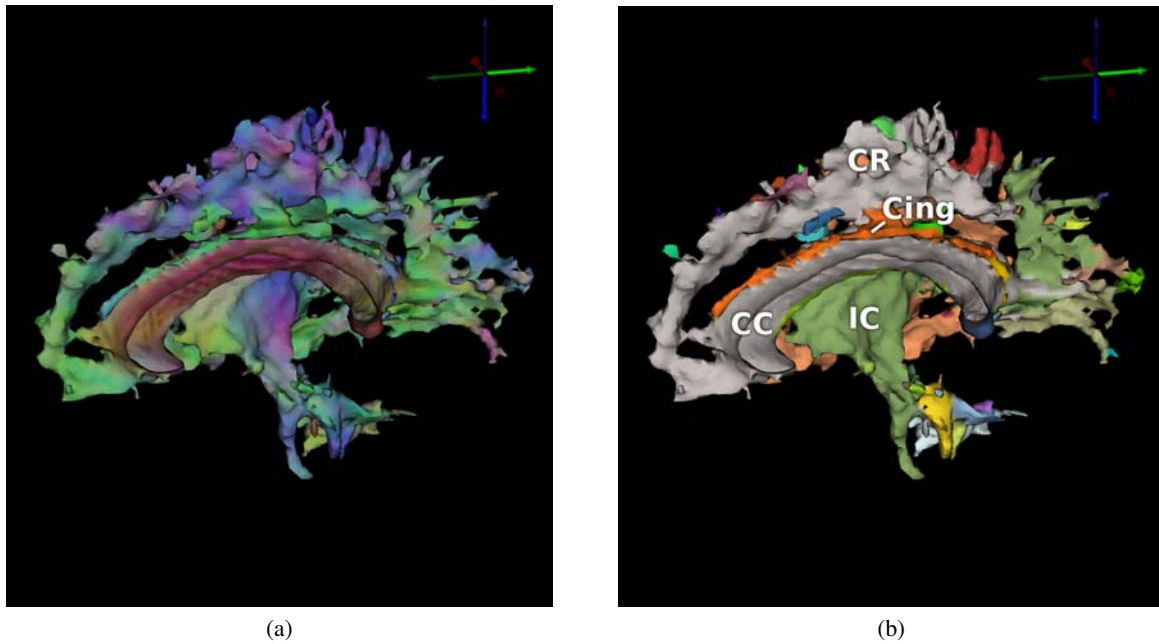


Figure 3: A part of an $FA=0.5$ isosurface, viewed from the midsagittal plane. (a) shows the standard XYZ-RGB color scheme, (b) presents an annotated segmentation result in random pseudocolors.

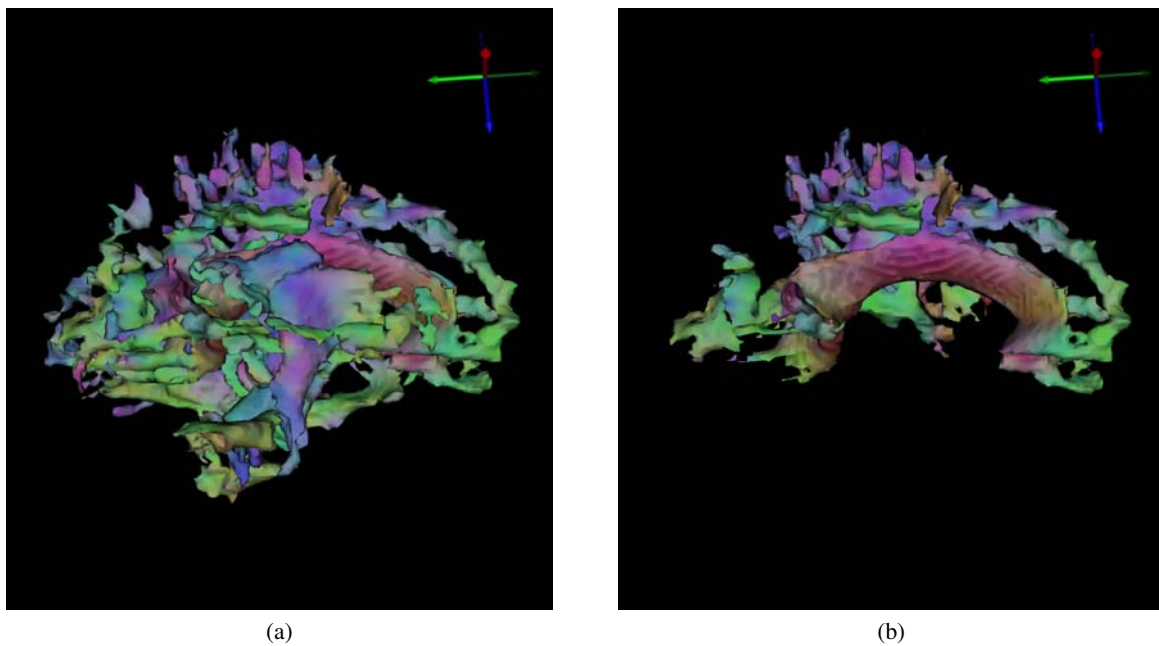


Figure 4: Segmentation results can be used to reveal parts of the surface which are otherwise occluded. While (a) shows the largest connected component of an FA isosurface, parts of it have been clipped in (b) to give a better view on the corpus callosum.




Figure 5: An isosurface at $c_1 = 0.26$, seen from the front/top/right. (a) shows the standard XYZ-RGB color scheme, while (b) presents the abstraction provided by a segmentation with region representative coloring. The annotations illustrate that our method has identified anatomically relevant regions.

regions will have representatives which are so similar that the resulting region colors cannot be distinguished visually.

In Figure 5 (b), we use a combination of stopping region merging with the c_1 criterion and coloring regions by their representative to abstract from variations within anatomic regions, allowing to tell apart the significant units at first sight and still indicating their overall orientation.

The result clearly reflects a number of anatomic structures: Again, a region has been identified that corresponds to the *internal capsule* (IC). The *cerebellar peduncle* (CP) and the *inferior fronto-occipital fasciculus* (IF) have also been segmented, but are mostly occluded in Figure 3. The *corpus callosum* (CC) is separated from the *cingulum bundle* (Cing), and the region-based c_1 criterion even allowed to distinguish it from the *corona radiata* (CR).

7. Conclusion and Future Work

In this work, we have suggested that the part of diffusion tensor data which is ignored by the anisotropy measure can be used to segment anisotropy isosurfaces in an anatomically meaningful manner, effectively adding information to the visualization and allowing the user to concentrate on parts of the surface that may be of particular interest.

We have demonstrated practical segmentation results using an efficient edge-based watershed approach, which extends methods that have previously been described in the

context of geometry-based mesh segmentation. The segmented regions have been shown to reflect a number of anatomically distinct structures and we have used the results to provide an abstracted view on the data that only shows relevant structures and their overall properties.

Even though anisotropy isosurfaces have been introduced for the visualization of DT-MRI data early on, they have so far not received much scientific attention, presumably because applying an anisotropy metric reduces the tensor data to a simple scalar field that can be processed using any of the well-known methods for isosurface extraction. However, note that both the FA and c_1 measures are nonlinear in the tensor values, so we introduce an error by assuming linear interpolation when extracting the isosurface. In high-curvature regions of the mesh, we have observed artifacts that stem from this fact. Upsampling the tensor field before applying the metric helps, but its advantages are minor on large parts of the mesh and it makes isosurface extraction too slow to be practicable. Thus, the development of an adaptive method for anisotropy isosurface extraction may be an interesting goal of future research.

Also, it may be interesting to see where the results of our method differ when applied to high angular resolution diffusion (HARD) MRI measurements, which do not employ the standard second-order diffusion tensor model. Different anisotropy metrics for such data have been proposed [Fra01, ÖVM05], so we expect that it would not be difficult

to extract isosurfaces from them and segment those by appropriately modifying our definition of edge weights.

Acknowledgments

We would like to thank Natascha Sauber for discussions at all stages of this project, as well as Rhaleb Zayer for a discussion on mesh segmentation in general. We are grateful to Alfred Anwander, who is with the Max Planck Institute for Human Cognitive and Brain Science in Leipzig, Germany, for providing the dataset used in this paper and helping to identify some of the segmented regions.

Our implementation builds on the BioTensor application [Sci06], which has been developed at the Scientific Computing and Imaging Institute at the University of Utah, and the team libraries <http://teem.sf.net/> by Gordon Kindlmann.

This research has partially been funded by the Max Planck Center for Visual Computing and Communication (MPC-VCC).

References

- [BP96] BASSER P. J., PIERPAOLI C.: Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor mri. *Journal of Magnetic Resonance B*, 111 (1996), 209–219. 1, 2
- [CG06] CHEN L., GEORGANAS N. D.: An efficient and robust algorithm for 3D mesh segmentation. *Multimedia Tools and Applications* 29, 2 (2006), 109–125. 2
- [Fra01] FRANK L. R.: Anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine* 45 (2001), 935–939. 7
- [FWB03] FEDDERN C., WEICKERT J., BURGETH B.: Level-set methods for tensor-valued images. In *Proc. Second IEEE Workshop on Geometric and Level Set Methods in Computer Vision* (Nice, France, October 2003), Faugeras O. D., Paragios N., (Eds.), pp. 65–72. 2
- [Kin06] KINDLMANN G.: Tensor invariants and their gradients. In *Visualization and Processing of Tensor Fields* (2006), Weickert J., Hagen H., (Eds.), Springer, pp. 215–224. 2
- [KWH00] KINDLMANN G., WEINSTEIN D., HART D.: Strategies for direct volume rendering of diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 6, 2 (April 2000), 124–138. 2
- [LC87] LORENSEN W. E., CLINE H. E.: Marching cubes: A high resolution 3D surface construction algorithm. In *SIGGRAPH '87: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1987), ACM Press, pp. 163–169. 2
- [MW99] MANGAN A. P., WHITAKER R. T.: Partitioning 3D surface meshes using watershed segmentation. *IEEE Transactions on Visualization and Computer Graphics* 5, 4 (1999), 308–321. 1, 2, 3, 4
- [OGW04] O'DONNELL L., GRIMSON W. E. L., WESTIN C.-F.: Interface detection in diffusion tensor MRI. In *Seventh International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI '04)* (2004), Barillot C., Haynor D., Hellier P., (Eds.), vol. 3216 of *LNCIS*, Springer, pp. 360–367. 2
- [ÖVM05] ÖZARSLAN E., VEMURI B. C., MARECI T. H.: Generalized scalar measures for diffusion MRI using trace, variance, and entropy. *Magnetic Resonance in Medicine* 53 (2005), 866–876. 7
- [PAB02] PAJEVIC S., ALDROUBI A., BASSER P. J.: A continuous tensor field approximation of discrete DT-MRI data for extracting microstructural and architectural features of tissue. *Journal of Magnetic Resonance* 154 (2002), 85–100. 2, 3
- [PJB*96] PIERPAOLI C., JEZZARD P., BASSER P. J., BARNETT A., DI CHIRO G.: Diffusion tensor MR imaging of the human brain. *Radiology* 201, 3 (1996), 637–648. 3
- [PKA03] PAGE D. L., KOSCHAN A. F., ABIDI M. A.: Perception-based 3D triangle mesh segmentation using fast marching watersheds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'03)* (2003), pp. 27–32. 2, 3
- [RHHP02] RETTMANN M. E., HAN X., HU C., PRINCE J. L.: Automated sulcal segmentation using watersheds on the cortical surface. *NeuroImage* 15, 2 (2002), 329–344. 2, 3, 4
- [Sci06] SCIENTIFIC COMPUTING AND IMAGING INSTITUTE (SCI): BioTensor: A scirun power app for processing and visualizing diffusion tensor images, 2006. 1, 8
- [SSS06] SCHÜNKE M., SCHULTE E., SCHUMACHER U.: *Prometheus: LernAtlas der Anatomie. Kopf und Neuroanatomie*. Thieme, 2006. 5
- [VLB*03] VIVODTZEV F., LINSEN L., BONNEAU G.-P., HAMANN B., JOY K. I., OLSHAUSEN B. A.: Hierarchical isosurface segmentation based on discrete curvature. In *Proceedings of the symposium on Data visualisation (VISSYM '03)* (2003), Bonneau G.-P., Hahmann S., Hansen C. D., (Eds.), pp. 249–258. 2
- [VZKL06] VILANOVA A., ZHANG S., KINDLMANN G., LAIDLAW D. H.: An introduction to visualization of diffusion tensor imaging and its applications. In *Visualization and Processing of Tensor Fields* (2006), Weickert J., Hagen H., (Eds.), Springer, pp. 121–153. 1, 2, 5
- [WPG*97] WESTIN C.-F., PELED S., GUDBJARTSSON H., KIKINIS R., JOLESZ F. A.: Geometrical diffusion measures for MRI from tensor basis analysis. In *International Society for Magnetic Resonance in Medicine '97* (Vancouver, Canada, April 1997), p. 1742. 2
- [ZKL05] ZHANG S., KINDLMANN G., LAIDLAW D. H.: Diffusion tensor MRI visualization. In *The Visualization Handbook*, Hansen C. D., Johnson C. R., (Eds.), Elsevier, 2005, pp. 327–340. 1
- [ZMB*03] ZHUKOV L., MUSETH K., BREEN D., WHITAKER R., BARR A. H.: Level set segmentation and modeling of DT-MRI human brain data. *Journal of Electronic Imaging* 12 (2003), 125–133. 1, 5