

Comparative Visualization for Diffusion Tensor Imaging Group Study at Multiple Levels of Detail

C. Zhang¹, T. Höllt^{1,2}, M.W.A. Caan³, E. Eisemann¹, A. Vilanova¹

¹Computer Graphics and Visualization, Delft University of Technology, Delft, The Netherlands

²Computational Biology Center, Leiden University Medical Center, Leiden, The Netherlands

³Brain Imaging Center, Academic Medical Center, University of Amsterdam, Amsterdam, The Netherlands

Abstract

*Diffusion Tensor Imaging (DTI) group studies often require the comparison of two groups of 3D diffusion tensor fields. The total number of datasets involved in the study and the multivariate nature of diffusion tensors together make this a challenging process. The traditional approach is to reduce the six-dimensional diffusion tensor to some scalar quantities, which can be analyzed with univariate statistical methods, and visualized with standard techniques such as slice views. However, this provides merely part of the whole story due to information reduction. If to take the full tensor information into account, only few methods are available, and they focus on the analysis of a single group, rather than the comparison of two groups. Simultaneously comparing two groups of diffusion tensor fields by simple juxtaposition or superposition is rather impractical. In this work, we extend previous work by Zhang et al. [ZCH*17] to visually compare two groups of diffusion tensor fields. To deal with the wealth of information, the comparison is carried out at multiple levels of detail. In the 3D spatial domain, we propose a details-on-demand glyph representation to support the visual comparison of the tensor ensemble summary information in a progressive manner. The spatial view guides analysts to select voxels of interest. Then at the detail level, the respective original tensor ensembles are compared in terms of tensor intrinsic properties, with special care taken to reduce visual clutter. We demonstrate the usefulness of our visual analysis system by comparing a control group and an HIV positive patient group.*

CCS Concepts

•Human-centered computing → Scientific visualization; Visualization design and evaluation methods;

1. Introduction

Diffusion Tensor Imaging (DTI) is an imaging modality that is able to characterize fibrous tissues such as brain white matter. DTI group studies aim at spatially locating diffusion-related differences between groups in order to discover potential markers for white matter structures, and are normally carried out in a voxel-wise manner. For example, several studies [SDT05, SJJ*06, DMA*09, VHNE*09] compare a control group and a patient group to search for differences caused by certain pathologies. Due to the complex and multifaceted nature of the tensor data, the analysis of a single group of diffusion tensor fields is already challenging. Inter-group comparison only adds more complexity by introducing the need for comparative solutions in the analysis. The standard approach relies on the univariate statistical analysis of tensor-derived scalar quantities such as fractional anisotropy (FA), rather than the multivariate diffusion tensor itself. Slice views are commonly used to display the analysis results, which are some specific highly aggregated scalar values (e.g., the t-statistic). This type of scalar-based comparative analysis provides merely one facet of the full tensor information, and leaves a large amount of information unused.

To keep the full tensor information, one has to resort to the multivariate analysis tools such as multivariate normal distribution [BP03, WWHT07]. However, the intrinsic structure of diffusion tensors is not taken into consideration, and it is also unclear how to effectively compare the generated fourth-order covariance tensors. Furthermore, only few visualizations are available for analyzing a single DTI group. For example, Abbasloo et al. [AWHS16] focus on visualizing the complex fourth-order covariance tensor [BP07] at multiple levels of detail while Zhang et al. [ZCH*17] summarize and visualize a single DTI group from the perspective of the intrinsic properties of diffusion tensors (i.e., tensor scale, shape, and orientation). However, to the best of our knowledge, there is no comparative visualization technique developed specifically for inter-group comparison.

Our work extends Zhang et al. [ZCH*17] to facilitate the visual comparison of two DTI groups at multiple levels of detail. Our comparative visualization framework is built based on the Shneiderman mantra “Overview first, zoom and filter, then details-on-demand” [Shn96]. The spatial overview is for comparing the multivariate statistical summary information of tensor ensembles, i.e., mean tensor & tensor variations [ZCH*17]. We com-

bine the checkerboard style pair-wise tensor difference (Tender) glyph [ZSL*16] and the tensor summary glyph [ZCH*17] to support the visual comparison in the spatial domain. More importantly, we integrate the details-on-demand scheme into the glyph visualization such that, as analysts zoom in, more statistical information will be progressively revealed. We refer to the visual comparison of statistical information in the spatial space as *summary comparison*. The purpose of the spatial overview is to guide analysts to locate voxels of interest for the further detailed comparison of the original tensors without aggregation. We extend the detail visualization of Zhang et al. [ZCH*17] by utilizing superposition and explicit difference encoding to support the comparison in terms of tensor intrinsic properties as the finest level, and refer to it as *feature comparison*. Furthermore, we avoid visual clutter by employing the kernel density estimation. The main contributions of this work can be summarized as:

- We present a level-of-detail glyph representation to support the visual comparison of tensor ensemble summary information at different zoom levels in the spatial domain;
- We facilitate the visual comparison of two groups of the diffusion tensors in terms of tensor intrinsic properties by superposition and explicit encoding.

2. Related Work

Our work mainly belongs to two visualization categories: comparative visualization and the visual analysis of DTI ensembles. We restrict the discussion here to literature relevant to these two areas.

2.1. Comparative Visualization

Comparative visualization refers to the process of understanding the differences between two elements of study by utilizing easy-to-interpret visual representations. Gleicher et al. [GAW*11] divide comparative visualization into three fundamental categories: juxtaposition, superposition, and explicit encoding. Juxtaposition refers to place the objects to be compared side-by-side. It is straightforward to implement, and it is the most commonly used technique in DTI group analysis [SDT05, VHNE*09]. Hotz et al. [HSNHH10] employ juxtaposition to compare tensor interpolation results. Schultz et al. [SSSSW13] juxtapose two sets of glyphs to compare fiber variability. Superposition takes objects into the same frame of reference for comparison but can easily cause occlusions and visual clutter. Jones et al. [JGA*02] overlay multiple tensor glyphs to evaluate the registration results while Bergmann et al. [BKLW06] overlay two tensor glyphs to depict the tensor estimation results. To suppress visual occlusion when superposing tensor glyphs, either image blending [AWHS16] or transparency [MVB*17] can be applied. Explicit encoding directly encodes the differences between objects into visual representations but it implies a suitable measure for difference quantification, which is not a trivial task especially for complex objects such as diffusion tensors. Kindlmann et al. [KSJEN*07] use slice views to visualize tensor shape- and orientation-specific differences. Zhang et al. [ZSL*16] combine explicit encoding and juxtaposition to support the voxel-wise comparison of two diffusion tensor fields.

2.2. Visual Analysis of DTI Ensemble

Basser and Pajevic [BP03, BP07] assume that the set of diffusion tensors follows a multivariate normal distribution, and propose the fourth-order covariance tensor to quantify the variation. Additionally, they use radial glyphs for the visualization. Kindlmann et al. [KEWW07] propose a decomposition framework, which is based on tensor gradient invariants and rotation tangents, to provide an intuitive understanding of the complex fourth-order covariance tensor. Abbasloo et al. [AWHS16] combine slice views, volume rendering, and superquadric glyphs [Kin04] to visualize the covariance tensor at multiple levels of detail. Zhang et al. [ZCH*17] propose to summarize diffusion tensor ensembles in terms of tensor intrinsic properties. They extend the conventional superquadric glyph [Kin04] to additionally encode aggregated variation information, and propose three detail views to show the actual distribution of the original tensors in tensor scale, shape, and orientation, respectively. Jiao et al. [JPGJ12] construct the so-called shape inclusion probability function to visualize an ensemble of diffusion orientation distribution functions (dODFs) by volume rendering.

Our work is the extension and combination of two previous works, i.e., the visual comparison of two diffusion tensor fields [ZSL*16] and the visual analysis of a single group of diffusion tensor fields [ZCH*17]. Our visual analysis approach facilitates the comparison of two groups of diffusion tensor fields in terms of tensor intrinsic properties at multiple levels of detail, ranging from tensor statistical summary to the original tensors. This type of comparative visualization technique, to the best of our knowledge, has not been reported previously.

3. Comparative Visualization Design

Neuroscientists normally apply voxel-wise analysis to locate diffusion related differences between two groups of DTI datasets. Therefore, the main objects to be compared in this work are two groups of diffusion tensors, the so-called *tensor ensembles*, at the same spatial location. Each tensor within the ensemble corresponds to one individual subject. We denote them as $\{\mathbf{D}_i\}_A$ and $\{\mathbf{D}_j\}_B$ with $i \in [0, m)$ and $j \in [0, n)$, where m and n are the total number of subjects in the corresponding groups A and B , respectively. We use two complementary and colorblind-safe colors, i.e., orange and blue [DBB*13, AWHS16], to identify different groups.

In the following sections, we first develop the visualization for *directly* comparing $\{\mathbf{D}_i\}_A$ and $\{\mathbf{D}_j\}_B$ but not considering the field information, and then design the visualization for comparing the ensemble summary information in the 3D spatial domain.

3.1. Direct Tensor Ensemble Comparison

Zhang et al. [ZCH*17] visualize a single tensor ensemble in terms of three tensor intrinsic properties (scale, shape, and orientation), each of which has a biologically meaningful interpretation. We extend the *detail views* [ZCH*17] to facilitate the visual comparison of two tensor ensembles, also from the perspective of the tensor intrinsic properties. Two synthetic tensor ensembles (see Figure 1a) are used for illustration in the following sections.

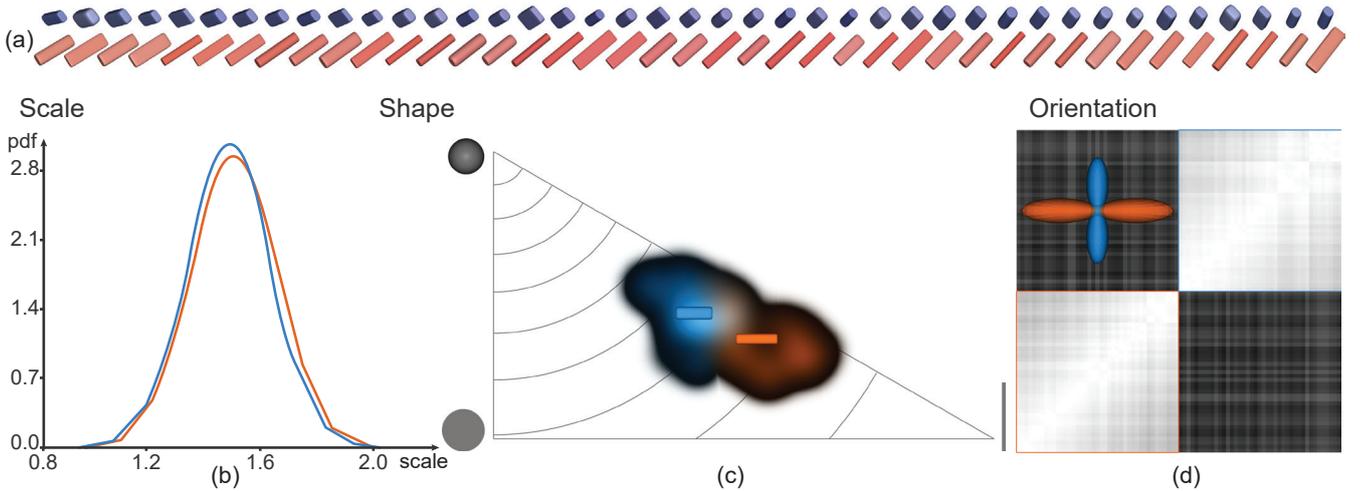


Figure 1: Direct visual comparison of two tensor ensembles. (a) each row represents a synthetic tensor ensemble, which is the subject of comparison. All the tensor glyphs are colored according to the major eigenvectors with the conventional spherical colormap [PP99]. The color is desaturated based on the tensor invariant c_1 [WMK*99]. The major eigenvectors of these two ensembles are perpendicular to each other. (b) shows the line plots of the probability density function of the one dimensional tensor scale for the two ensembles. (c) shows the blended 2D density plot in the triangular tensor shape space, overlaid with the mean tensor glyphs for each ensemble. (d) shows the pair-wise orientation dissimilarity matrix with the two dODF [ALS*10] glyphs. (b–d) blue and orange colors indicate the tensor ensemble at the upper and lower row in (a), respectively. Note that the glyphs in (a, c, d) are visualized with different viewing settings.

Tensor Intrinsic Properties. A symmetric positive-definite diffusion tensor \mathbf{D} can be eigen-decomposed into three ordered real eigenvalues ($\lambda_1 \geq \lambda_2 \geq \lambda_3 > 0$), and the corresponding orthonormal eigenvectors $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$. λ_i and \mathbf{e}_i ($i = 1, 2, 3$) are referred to as the major, medium, and minor eigenvalue and eigenvector, respectively. The tensor trace $|\mathbf{D}|$ is defined as the sum of the eigenvalues $|\mathbf{D}| = \sum_{i=1}^3 \lambda_i$. Finally, the ordered normalized eigenvalues ($\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3$) are given by normalizing the eigenvalues with respect to the trace $(\lambda_1, \lambda_2, \lambda_3) / |\mathbf{D}|$.

Tensor intrinsic properties include scale, shape, and orientation. Tensor scale is mathematically expressed as the trace $|\mathbf{D}|$, which represents the amount of diffusion. Tensor shape is two dimensional, and mathematically expressed as the set of normalized eigenvalues $(\tilde{\lambda}_1, \tilde{\lambda}_2, \tilde{\lambda}_3)$. Tensor shape reflects the underlying fiber configurations. For example, for regions with coherently aligned fiber tracts, the diffusion tensors are of linear shape. Tensor orientation is given by the set of eigenvectors $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$, which can be used to deduce the pathway of the fiber tracts. Interested readers are referred to previous work [EK06, ZSL*16] for more details.

Scale Comparison. Zhang et al. [ZCH*17] employ a 1D probability density function (PDF) to depict the distribution of tensor scale, which is visualized as a line plot. It is straightforward to compare two distributions of tensor scale. We use the *superposition* strategy to overlay two line plots within the same frame of reference, each of which is colored according to their group identifier color. For example, as shown in Figure 1b, two line plots largely overlap, indicating that two corresponding distributions are quite similar to each other.

Shape Comparison. Tensors with all possible shapes are located on a right triangle [ZCH*17], and hence we compare the shapes of two tensor ensembles in this triangular shape space. However, displaying all the tensor glyphs at once or drawing them as a scatter plot generates clutter or does not allow an easy understanding of the actual shapes, respectively. We tackle these problems by combining 2D probability density estimation [LH11] with glyph-based visualization. First, we draw the 2D density plot for each tensor ensemble separately. We then blend the two plots in image space as $c_{i,j} = p_{i,j}^A * c_A + p_{i,j}^B * c_B$, where (i, j) is the pixel coordinate, p is the density value, c_A and c_B are the corresponding group identifier colors (see Figure 1c). Abbasloo et al. [AWHS16] employ the same strategy to avoid occlusions for glyph-based visualization. Blending the colors as described will create an impression of neutral gray, where the density values are similar. If one of the values is larger, the corresponding group color will become obvious at this position, and thus allow the identification of differences in the distribution of the tensor shape between ensembles. However, the mental reconstruction of the actual tensor shapes is not easy. Therefore, we draw several representative glyphs to assist the perception of the blended density plot without introducing much occlusion. We provide two options. First, the user can select to simply show the mean tensor for each group, resulting in just two glyphs in total rendered on top of the density plot (see Figure 1c). However, the shape distribution can be multi-modal. Therefore, the second option is based on automatic peak finding in the density and subsequent clustering. To this end, we apply GPU-based mean shift clustering based on the generated density plot, similar to previous work by Höllt et al. [HPvU*16]. The mean shift algorithm [FH75] has several useful properties. The density clusters will closely match the visual

clusters that analysts would identify. Furthermore, since we compute the density plot for visualization, we can use the same density estimate as input for the mean shift algorithm. Finally the mean shift algorithm can also function as a peak finding algorithm on the density plot and the identified peaks can be used to place one representative glyph for each cluster. To obtain the clusters, based on the density estimate, we first compute the gradient of the density estimate by calculating the central differences for each pixel in image space. Then, starting with the initial position for each glyph, we ascend along the gradient direction until we reach a peak. All glyphs that ascend to the same peak are considered as belonging to the same cluster and the peak can be used for placing the representative glyph for that cluster. Finally, we reconstruct the shapes for the representative glyphs, based on the position of the peak in image space, and render the glyphs on top of the density plots. An example with tensor glyphs at cluster centers can be seen later in Figure 4d top. By combining the 2D density plot and the tensor glyphs, we effectively reduce the visual clutter and provide a straightforward comparison in the tensor shape distribution.

Orientation Comparison. It is challenging to define a space to explicitly convey the orientation distributions due to the inherent coupling of tensor shape and orientation. As suggested by Zhang et al. [ZCH*17], the pair-wise dissimilarity matrix of tensor orientations proves to be useful in identifying similar orientation behaviors. Therefore, we order the elements for each tensor ensemble, in the same way as Zhang et al. [ZCH*17], and then concatenate these two ensembles together to build the final dissimilarity matrix. The orientation matrix view is therefore composed of three parts (see Figure 1d). A white-to-black perceptually linear colormap is employed to color the matrix. The lower diagonal part depicts the orientation variation of one ensemble, while the upper diagonal part shows that of the other. The remaining two parts, which are symmetric with respect to the diagonal, depict the inter-group orientation differences. The diagonal parts allow exploring the orientation variation within the ensemble, and the off diagonal parts allow inspecting the inter-ensemble orientation differences, which belongs to the comparison category of *explicit encoding*.

In order to support the interpretation of tensor orientation variation of each ensemble, we convert each tensor into the so-called diffusion orientation distribution function (dODF) [ALS*10], and obtain an average dODF glyph for each ensemble [ZCH*17]. In principle the dODF glyph is a directional probability density function defined on a unit-sphere. We apply the *superposition* strategy to compare two dODF glyphs, and provide the option to add transparency. We place the dODF glyphs on top of one of the two symmetric off diagonal parts, which will not cause information loss. The dissimilarity matrix is two dimensional while the dODF glyphs are three dimensional objects. Therefore, we offer the option to freely rotate the dODF glyphs to have a comprehensive perception without affecting the viewing settings. Hereby, we not only show the orientation variations for a single ensemble, but also facilitate the visual comparison between two ensembles.

As shown in Figure 1d, both ensembles exhibit a small intra-ensemble orientation variation (i.e., brighter colors), while the inter-ensemble orientation differences are rather large (i.e., darker

colors). Additionally, the superposed dODF glyphs clearly shows the differences in dODF for every directions.

Linked Brushing. Each view is extended to support the inter-group comparison in one tensor intrinsic property. Therefore, linked brushing is used to connect all of the three views. To minimize visual clutter, we only highlight the selected tensor glyphs with red halos (e.g., see Figure 4d bottom). A filled line plot (e.g., see Figure 4e) is drawn to reflect the contributions of selected tensors to the overall scale distributions, and the corresponding diagonal entries of the matrix view (e.g., see Figure 4f) are also colored red. We use a table view to display the metadata (e.g., age) of each individual subject, which can also be used for linked selection.

3.2. Ensemble Summary Comparison

The extended comparative visualization, described in the previous section is useful for directly comparing the tensor ensembles for voxel(s) of interest. However, this visualization cannot provide spatial information. Therefore, we use a separate spatial visualization, described in this section, to show the ensemble differences in the 3D spatial domain. Before performing voxel-wise ensemble comparison information needs to be aggregated. Zhang et al. [ZCH*17] summarize the tensor ensemble $\{\mathbf{D}_i\}$ into a mean tensor $\bar{\mathbf{D}}$ and the variations in scale σ_{scale} , shape σ_{shape} , and orientation $\sigma_{\text{orientation}}$. Furthermore, they extend the superquadric tensor glyph [Kin04] to visualize the summary information.

We employ the same aggregation method, which handles the tensor intrinsic properties separately, and design a new glyph to visually compare the summary information in 3D. Our new design is a combination of the checkerboard style Tender glyph [ZSL*16], designed for comparing two diffusion tensors, and the tensor ensemble summary glyph [ZCH*17], designed for visualizing the summaries of a single tensor ensemble. Both types of glyphs are suitable for detailed inspection at close-up views. In this work, we incorporate the level-of-detail concept by designing different visual representations at different zoom levels. For a small field of view, it is more suitable to encode the full information for a limited number of voxels, while for a large field of view it is more suitable to encode a limited amount of information for a large amount of voxels to provide an overview. The zoom level is reversely proportional to the extent of information aggregation; the higher the zoom level, the less aggregation is needed. Therefore, in the following, we describe the levels from the most detailed to the coarsest by gradually removing information.

Aggregation Level 1 (Zoom Level 3) : Mean tensors and tensor variations.

At the finest zoom level, the aggregate information to be compared is the mean tensor and three scalar-valued tensor variations. We apply the checkerboard style Tender glyph [ZSL*16] to compare the mean tensors. The original design uses the color channel to compare tensor scales. However, we decide to reserve the color channel for later usage at the coarsest zoom level. Therefore, we have to resort to other visual channels to encode the tensor scale information. We find that the original sector representation, which is shown both above and below the glyphs, for the orientation difference is redundant. Hence, we divide the sector into two parts,

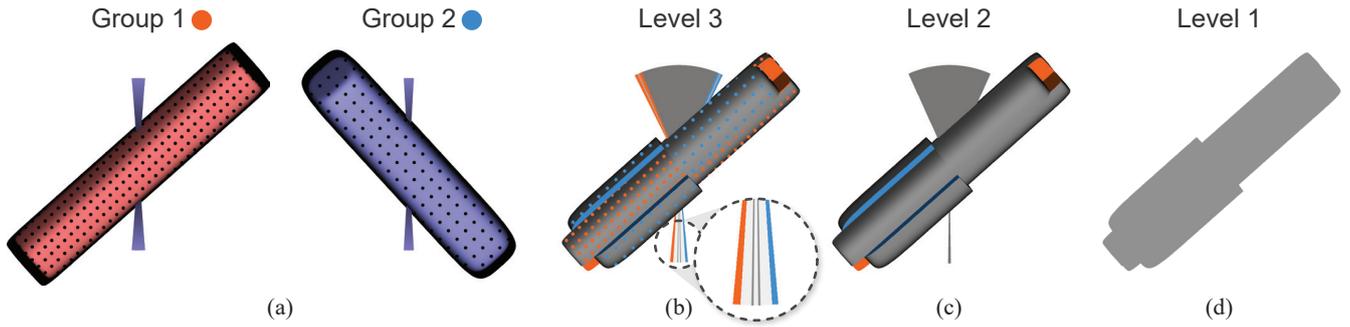


Figure 2: Level-of-detail glyph representations for tensor ensemble summary comparison. (a) displays the corresponding tensor summary glyphs for the synthetic ensembles shown in Figure 1a. (b) shows the new LOD glyph at zoom level 3 (i.e., the fine level), designed to compare the tensor ensemble summary information. (c) shows the new LOD glyph at zoom level 2 (i.e., the medium level), used to support the visual comparison of the corresponding mean tensors alone. (d) shows our LOD glyph at zoom level 1 (i.e., the coarse level) with plain color coding to indicate the tensor shape difference. (b,c) orange and blue color are for the identification of different tensor ensembles.

i.e., the upper and lower parts. Then, the angles of the upper and lower part are used to encode orientation-related and scale-related differences, respectively. Figure 3 illustrates the idea of sector representation by analogy with the 1D normal distribution. At zoom level 3, two bands are drawn with the identification color. The center of the band is aligned with the mean value (i.e., the vertical dashed line) while the length of the band is determined by the standard deviation. As shown in Figure 3a, if these two bands overlap, the overlapping region will be colored white (i.e., the addition of orange and blue). Otherwise, if there is a gap between the two bands, the in-between region is colored to with a dark gray (Figure 3b). We create the sectors for the glyph visualization by replacing the length with the angle of the arc. Figures 3a and 3b at zoom level 3 are analog to the lower and upper sectors in Figure 2b, respectively, with the standard deviation corresponding to scale and orientation variations [ZCH*17]. The major benefit with this visual design is that it facilitates the direct evaluation of the overlap between the two distributions in scale and orientation.

For the comparison of mean tensor shapes, we place them in a checkerboard style as presented for the Tender glyph [ZSL*16]. For the visual comparison of tensor shape variations σ_{shape} , we employ the halftone textures due to their ability to provide a more precise estimation [ZCH*17]. In this way, we convert the comparison of tensor shape variations into that of the density of the colored dots.

Figure 2b shows the final glyph representation at this zoom level, which combines the two ensemble summary glyphs in Figure 2a. Please note that Figures 1a and 2a are generated with different viewing settings. The shape differences between the two mean tensors are manifested as the discontinuities between the glyphs. The orange ensemble exhibits a slightly larger shape variation (i.e., the dots are packed slightly more dense). The upper gray sector shown in Figure 2b indicates a large orientation difference between the mean tensors while the two small colored sections on both sides indicate little orientation variation within both ensembles. From the lower sector shown in Figure 2b, we can deduce that two ensembles have similar scale variations (i.e., two small orange and blue ends

on both sides) and largely overlap with each other (i.e., the white region).

Aggregation Level 2 (Zoom Level 2) : Mean Tensors. As users zoom out, more voxels are shown, and hence fewer pixels are assigned per voxel. Accordingly, we reduce the amount of information to be displayed. Prior to comparing the associated tensor variation information, neuroscientists are normally interested in comparing the mean tensors. If the mean tensors are different, neuroscientists prefer to dive further to assess the overlap between the distributions of the tensor ensembles. Therefore, at this zoom level, we decide to compare the mean tensors alone. More specifically, we

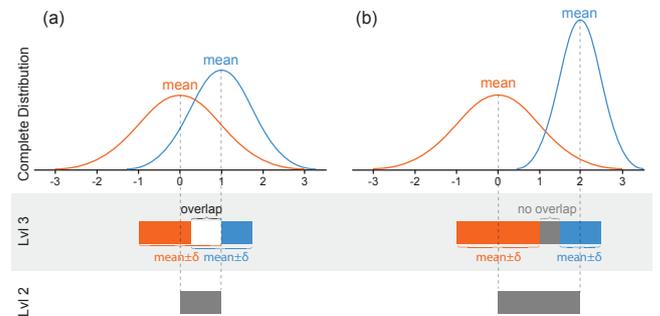


Figure 3: Illustration of the sector construction by analogy with one dimensional normal distributions for (a) a large and (b) small distribution overlap. Blue and orange colors are used to identify the corresponding normal distributions. At level 3, two colored bands are drawn, the length of which are proportional to the standard deviation of the distribution. The center of the bands are based on the mean value. (a) if these two bands are overlaid, white color is generated. (b) Otherwise, the color in-between is kept to gray. At level 2, a single gray band is drawn, the length of which indicates the difference in the mean values of the distributions. (a) and (b) are analog to the lower and upper sectors, respectively, in Figures 2b and 2c.

leave out the visual channels that are used to encode tensor variation information, and show the glyphs to only support the comparison of the mean tensors. Figure 2c shows that the glyph at this level is a simplified representation of the glyph at zoom level 3. Specifically, the textures are removed. The sectors are simplified to encode only the pair-wise scale/orientation differences to the angles, which are analog to the visual representation at level 2 in Figure 3. Furthermore, the sectors are colored using a neutral gray. At zoom levels 2 and 3 the glyph color is employed to explicitly encode the difference of the ensembles, as described in the next paragraph.

Aggregation Level 3 (Zoom Level 1): Difference Measure between Mean Tensors. At the coarsest zoom level, each voxel occupies a very limited number of pixels, and therefore requires a higher level of aggregation. The goal of this level is to help analysts to gain a first overview and identify areas of interest, into which they can dig deeper. Consequently, we decide to reduce these two mean tensors to single scalar values that can then be encoded into the color channel, which is still visible even if only few pixels are available per voxel. At this level, it is most important to choose an appropriate difference measure between the mean tensors. Such measures are strongly application-dependent. For example, Abbasloo et al. [AWHS16] use slice views to show the overall variance and generalized FA of the fourth-order covariance tensor to provide a first impression of the DTI ensemble. We allow analysts to choose among a set of pre-defined difference measures between the mean tensors $d(\bar{\mathbf{D}}_A, \bar{\mathbf{D}}_B)$, depending on the task at hand. For example, the analyst can choose the difference in a specific tensor intrinsic property such as the tensor shape difference d_{shape} . Furthermore, statistical measures like the t-statistic, with which neuroscientists are familiar in their routine research, can be used. We decide to use the shape of the glyphs instead of simple view-aligned quads to show the color at this zoom level. This guarantees a smooth transition between different glyph representations as much as possible. Since usually only a few pixels are used for the glyph at this zoom level, the shape of individual glyphs is barely noticeable, making it unnecessary to apply an illumination model. Figure 2d shows the glyph representation at the coarse level, which is colored according to the shape differences between the mean tensors.

3.3. Summary

We have introduced the three zoom levels in reverse order for the sake of the easy explanation of the aggregation of tensor ensembles. The actual visual exploration workflow starts from zoom level 1. The coarse zoom level utilizes the *explicit encoding* strategy to depict application-dependent difference measures between two mean tensors. This level provides a first impression for a large field of view, and hence can guide analysts to voxels of interest (e.g., Figure 4a). Analysts can then zoom into level 2, which is based on *juxtaposition* and *explicit encoding*. At this level, analysts can visually compare the mean tensors in all the three aspects of tensor intrinsic properties. If analysts would like to further compare the tensor ensemble variations, they can continue to zoom level 3, which is also based on *juxtaposition* and *explicit encoding*. At this level, analysts can scrutinize a small amount of voxels, harnessing the full tensor ensemble summary information. Finally, this zoom

level serves as the bridge between the comparison of the tensor ensemble summaries in spatial domain and the comparison in tensor intrinsic properties for voxels of interest. This level-of-detail glyph visualization enables the gradual exposure of the differences between two tensor ensembles in a voxel-wise manner.

4. Application

We demonstrate the usefulness of our visual analysis prototype by an exemplary comparison of two groups of diffusion tensor fields. The two groups consist of a group of 75 HIV positive (HIV+) patients and a control group with 46 subjects. Each DTI dataset has an original resolution of $224 \times 224 \times 144$ with an isotropic voxel size of 2mm. We uniformly crop each dataset to $121 \times 161 \times 111$ voxels in order to remove those outside of the skull. The diffusion weighting is $b=1000$ s/mm² along 64 uniformly distributed directions with four non-diffusion averages. The numerical range of the tensor values is scaled by 1000. The case studies were conducted and interpreted together, by the first and the third author, who is a neuroscientist specialized in DTI group analysis. Throughout the examples, we indicate the HIV+ group by orange color and the control group by blue color.

All datasets are registered with a non-rigid approach provided by DTI-TK [ZYAG06], which takes the whole tensor information into consideration. The registration quality is important for voxel-wise analysis, however, here, we focus on the comparative visualization of DTI group studies only. Therefore, the influence of the registration quality on the analysis results is beyond the scope of this work. Although we do not explicitly show the detail level comparison for several selected voxels, it is possible to do this with our prototype. Therefore, this type of local averaging can to some extent reduce the influence of potential mis-registration. Please note that our detail level comparison makes no difference whether a single voxel or multiple voxels are selected.

4.1. Case Study 1: Fornix

The neuroscientist starts the analysis with the visualization of the coarsest zoom level (Figure 4a) to identify voxels of interest. This zoom level is constructed based on the tensor shape difference d_{shape} , which represents the difference in the mean diffusion anisotropy profiles between two groups. We can see that the highlighted region is largely white, indicating that the voxels exhibit only small differences in shape. This region roughly matches the ventricle of the brain (see the dashed line in Figure 4a). Therefore, these small differences in shape are not surprising, since the diffusion tensors in the ventricles are nearly isotropic. We can see a small region in the center, highlighted by the rectangle in Figure 4a, which is notably darker, indicating a larger shape difference. The neuroscientist selects this region, which turns out to be the body of the fornix, for further analysis in zoom level 2 (Figure 4b). The fornix is a C-shaped bundle of fibers that roughly travel anteriorly. Note that in zoom level 2 the glyphs are rotated without changing the viewing settings for the sake of a clearer illustration. Here, we can see that the more elongated glyphs exhibit larger pair-wise shape differences (i.e., darker colors). For example, in the central voxels, the discontinuities between the checkerboard sides of the

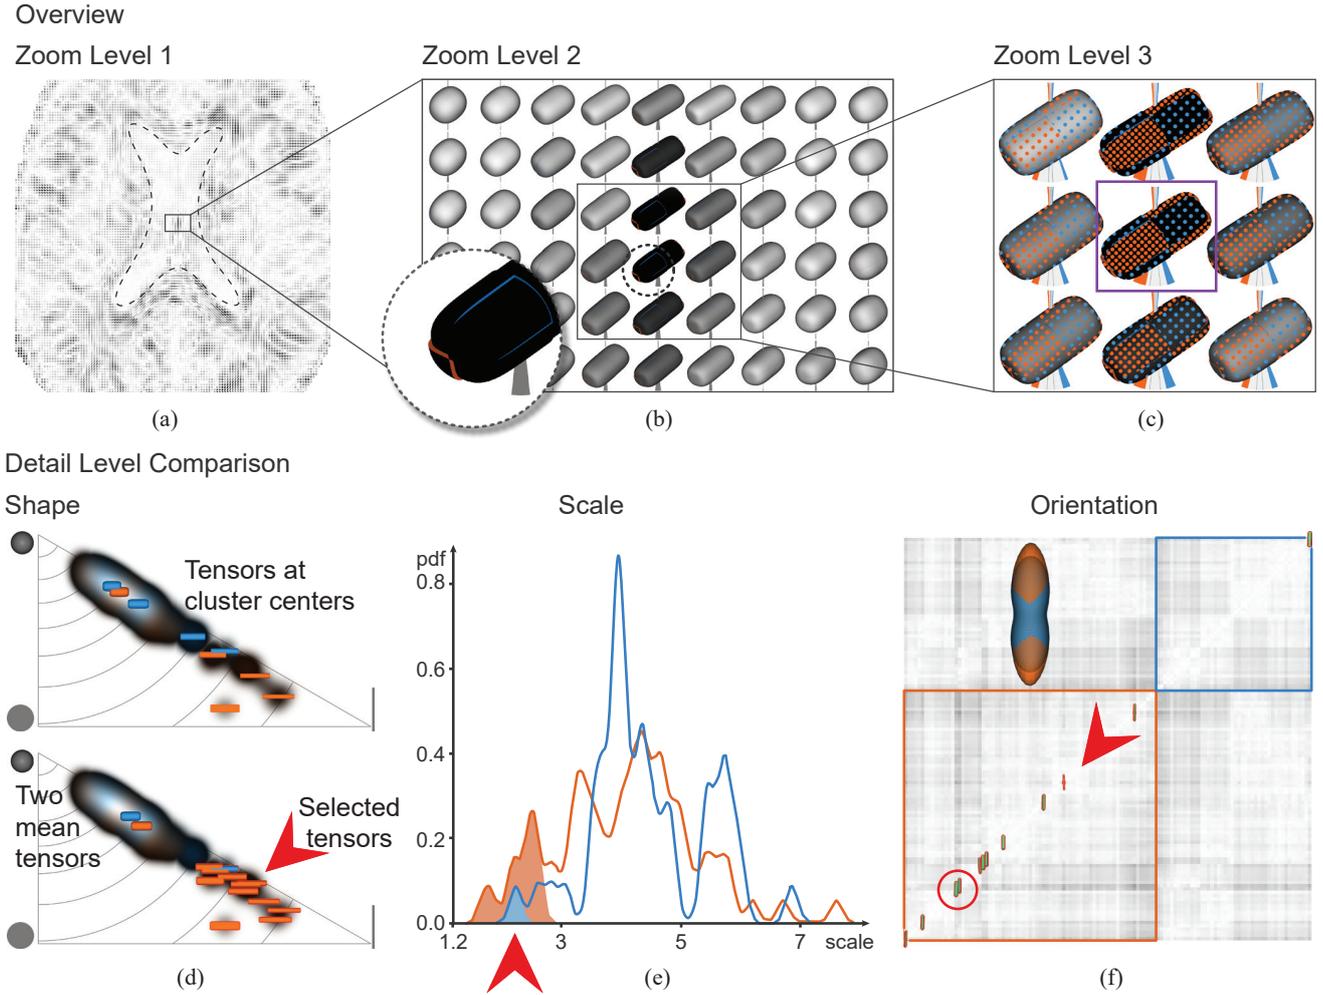


Figure 4: Exemplative exploration of the fornix (roughly indicated by the small rectangle in a). Overview visualization in the spatial domain at different zoom levels, from coarse to fine are shown in (a–c). At zoom level 2 (b), the dark glyphs indicate that the average tensors of the patient group have larger normalized first eigenvalues (e.g., see the inset). At zoom level 3 (c), it is observed that the shape variation (i.e., the colored dots) is higher for the HIV+ group than that of the control group. The voxel enclosed in the purple rectangle (c) is chosen to be further analyzed in the feature space (d–f). Here a selected group of subjects is indicated by red arrows and shown as additional glyphs with red halos (d, bottom), area under the pdf of tensor scale (e), and glyphs on the diagonal of the orientation dissimilarity matrix (f). The dODF glyphs are used to assist the comparison in tensor orientations.

glyphs can be readily perceived (due to size limitations of the figure, see the inset in Figure 4b for a clearer view). Glyphs with orange sidefaces have larger normalized major eigenvalues $\tilde{\lambda}_1$ but smaller normalized medium $\tilde{\lambda}_2$ and minor eigenvalues $\tilde{\lambda}_3$. This implies that, on average, the HIV+ group has larger linear anisotropies in the fornix. This finding surprises the neuroscientist, as it conflicts with the common hypothesis that diseases of the nervous system reduce the extent of diffusion anisotropy by damaging the underlying fiber tracts. Furthermore, we can observe that the orientation differences are consistently small for this region, while scale differences (i.e., the angles of the lower sectors) are proportional to the shape differences (i.e., the darkness of the color). The neuroscientist further zooms into the finest level to inspect variations in a

three by three voxel area in the center, as marked in Figure 4b. In Figure 4c we can see that variation in orientation is small and similar for both groups, as indicated by the angles of the upper sectors. The glyphs corresponding to the HIV+ group exhibit larger scale variations (i.e., larger angles in the orange part of the lower sectors) and larger shape variations (i.e., more densely packed orange dots). The fornix on the axial slice consists of only a few voxels and, therefore, partial voluming effects and/or mis-registration can lead to a contamination by voxels from the surrounding ventricles which normally have larger diffusion scales. The larger variation in the diffusion scales could be an indicator that the HIV+ group is more vulnerable to such contamination. More surprising is the large difference in shape variation between two groups (see the purple

rectangle in Figure 4c). Therefore, the neuroscientist selects this voxel to compare the original tensor ensembles at the detail level.

First, the neuroscientist inspects the shape distribution. Figure 4d top shows the density plot of the shape distribution, overlaid by the glyphs corresponding to the automatically generated density cluster centers. We can see that the tensor shapes are mainly linear for both groups. The overall span is quite large, with FA values ranging from 0.2 to 0.9, as indicated by the glyph locations relative to the background FA contours. Therefore, we select all the tensor glyphs corresponding to large anisotropies (i.e. the glyphs marked by the red arrow in Figure 4d, bottom). It can be seen that most of the selected glyphs originate from the HIV+ group (orange color) while only one is from the control group (blue). The neuroscientist speculates that this could be attributed to the scanner shift, as some subjects were scanned with two different scanners. However, we could exclude this hypothesis by inspecting the origin of each scan, which showed that 1) both groups contain subjects scanned with either scanner and 2) the selection includes tensors originating from either scanner (see the supplementary video). Inspecting the detail visualization for the scale (Figure 4e), we can see that all of the selected tensors also exhibit small values for tensor scale (see the colored region in Figure 4e). The orientation matrix view (Figure 4f) does not exhibit such a strong clustering. Overall, we see that, based on the average dODF glyphs with added transparency, the HIV+ group has larger diffusion probability in the vertical direction. We also find two tensors in the HIV+ group (red circle in Figure 4f) exhibiting larger orientation difference compared to the rest of the group and some subjects from the control group.

This example demonstrates that our prototype is able to help the neuroscientist to obtain new findings, i.e., the HIV+ group presents both larger linear anisotropies and tensor shape variations for the fornix structure. While the interpretation of the described findings is out of the scope of this article, they inspired the neuroscientist to further verify these findings with a quantitative analysis.

4.2. Case Study 2: Corpus Callosum

In the second case study we inspect the corpus callosum (CC), which is a flat bundle of fibers connecting the left and right cerebral hemispheres. It is indicated by the dashed line in a sagittal slice shown in Figure 5a. The neuroscientist starts with the overview visualization on zoom level 1, employing the commonly used scalar-valued statistical measure, i.e., t-statistic, color-coded using a green-white-violet colormap. We use a diverging colormap, since the values for the t-statistic are centered around zero and specifically this green to violet one, as it does not interfere with the blue and orange colors used for the different groups. The neuroscientist hypothesizes that the FA value is higher in the control group than in the HIV+ group. A positive t-statistic (violet) would indicate that the hypothesis is true while a negative t-statistic (green) indicates that it is false. The absolute value of the t-statistic indicates the extent to which the FA values are significantly different between two groups. *Significant difference* intuitively means large difference in the FA values of the mean tensors and small distribution overlaps. The voxel-wise t-statistic is constructed following the standard routine in the Tract-Based Spatial Statistics (TBSS) toolkit based on FA [SJJB*06]. In zoom level 1 (Figure 5a), it can

be seen that the anterior part of CC (i.e., the genu) appears to be more violet while the posterior part (i.e., the splenium) exhibits some green. The neuroscientist is specifically interested in these voxels with negative t-statistic values and zooms into level 2 for this region (Figure 5b). We observe that the HIV+ group, indicated by orange sidefaces (see the inset in Figure 5b), has larger normalized major eigenvalues $\tilde{\lambda}_1$ and smaller normalized medium and minor eigenvalues, similar to what we have found in the fornix. Furthermore, the scale and orientation differences are quite small, as indicated by the small lower and upper sectors, respectively. For further comparison of the group variations, we then zoom into level 3 (see Figure 5c). It can be seen that the variations in scale, shape, and orientation are quite similar between the two groups. We now inspect the complete distributions for a single voxel, indicated by the square in Figure 5c. In the shape space visualization (Figure 5d), we can immediately see that the two mean tensor glyphs are placed at some distance. In combination with the emphasized blue and orange tones on the ends of the density plot this confirms that the FA is indeed larger for the HIV+ group, as indicated by the green color in Figures 5a–c. As discussed before, diseases like HIV damage the fiber structures, which then makes the water diffuse more freely. Consequently, we expect the HIV+ patients to exhibit a larger amount of diffusion [WSC*06]. However, inspecting Figure 5e, we find that the distribution of the tensor scale for the HIV+ group (orange line) is slightly more skewed towards a lower scale value compared to the control group (blue). This means that a subject from the HIV+ group has a higher probability to present a small amount of diffusion, compared to a subject from the control group. This finding contradicts the expectation that HIV+ patients would exhibit a larger amount of diffusion. While we cannot present a solid interpretation for this finding in the scope of this paper, it shows that our prototype is able to reveal more information compared to the standard FA-based statistical analysis of DTI inter-group comparison.

5. Conclusions and Future Work

In this work, we present a level-of-detail comparative visualization technique to support the comparison between two groups of diffusion tensor fields. The comparison is divided into two parts, i.e., the comparison of tensor ensemble summaries in the spatial domain and that of the original tensor ensembles in terms of tensor intrinsic properties. For the comparison of tensor ensemble summaries, our glyph representation is based on the combination of two existing glyph designs [ZSL*16, ZCH*17]. More importantly, we integrate the details-on-demand concept to present more information as users zoom in. The coarse zoom level is designed to show large scale patterns of user specified scalar-valued difference measures, mimicking slice views. The medium zoom level facilitates the comparison of the mean tensors of the two groups in terms of their tensor intrinsic properties. The glyph representation at this level is a variant of the Tender glyph [ZSL*16]. The finest zoom level additionally supports comparison of variation information of each group. The glyph representation at this level is an extension of the tensor ensemble summary glyph [ZCH*17]. Here, the glyph-based comparative visualization helps analysts to locate voxels of interest for further comparison in tensor intrinsic properties at the detail level. Therefore, we extend the detail visualization [ZCH*17] to support

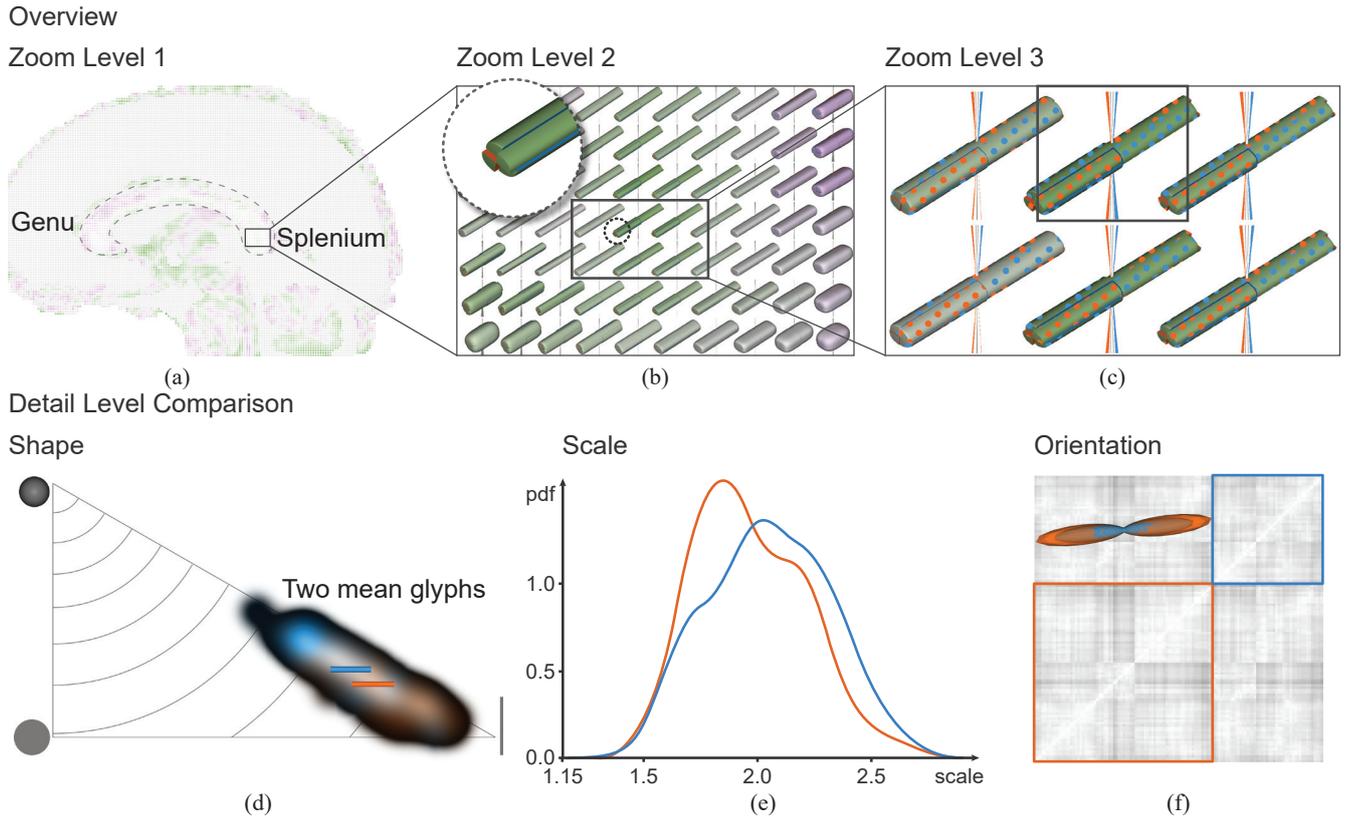


Figure 5: Comparative visualization for the corpus callosum (CC), colored according to the t -statistic values with a green-white-violet colormap. (a) shows the coarsest zoom level with the CC marked with the dashed line. (b) shows the mean tensor glyphs at zoom level 2, corresponding to the splenium that is marked by the rectangle in (a). The green color means negative t -statistic values for the splenium. The inset shows the HIV+ group has larger normalized major eigenvalues, as shown by the discontinuity, exposing the orange sideface. (c) shows the visual comparison of 6 selected voxels at zoom level 3 indicating rather small difference between the two groups. (d, e, f) show the detail level comparisons for the shape, scale, and orientation, respectively.

the visual comparison of two original tensor ensembles by exploiting the superposition and explicit encoding strategy. We performed two case studies together with a neuroscientist to illustrate the usefulness of our prototype by comparing an HIV+ patient group and a control group. Our prototype successfully reveals new findings that are otherwise unknown with standard scalar-based comparison. The case studies triggered the neuroscientist's interest to perform a formal study to further inspect these new findings.

There is still room for improvement as future work. Regarding the glyph design, the sector representation is not intuitive since it relies on the users' memory of the correspondence to tensor scale or orientation. Furthermore, although the image-space textures, used to indicate shape variation, provide a good estimation of the variation, their image-space nature can cause annoying effects during interaction. Research regarding these issues could further improve the glyph representation. In the spatial visualizations, the glyphs are placed on a regular grid, which is likely to hide the underlying continuous structures of the field. The combination of a smart sampling scheme [KW06] with the level of detail glyph representations

could further improve the overview visualizations. Finally, the next logical step would be the comparison of more than two groups.

Acknowledgements

We gratefully acknowledge the Co-morbidity in relation to Aids (COBRA) project for providing the DTI datasets that are used in Section 4.

References

- [ALS*10] AGANJ I., LENGLET C., SAPIRO G., YACOUB E., UGURBIL K., HAREL N.: Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magnetic Resonance in Medicine* 64, 2 (2010), 554–566. 3, 4
- [AWHS16] ABBASLOO A., WIENS V., HERMANN M., SCHULTZ T.: Visualizing tensor normal distributions at multiple levels of detail. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 975–984. 1, 2, 3, 6
- [BKLW06] BERGMANN Ø., KINDLMANN G., LUNDERVOLD A., WESTIN C.-F.: Diffusion k-tensor estimation from Q-ball imaging using

- discretized principal axes. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2006), pp. 268–275. 2
- [BP03] BASSER P. J., PAJEVIC S.: A normal distribution for tensor-valued random variables: applications to diffusion tensor MRI. *IEEE Transactions on Medical Imaging* 22, 7 (2003), 785–794. 1, 2
- [BP07] BASSER P., PAJEVIC S.: Spectral decomposition of a 4th-order covariance tensor: applications to diffusion tensor MRI. *Signal Processing* 87, 2 (2007), 220–236. 1, 2
- [DBB*13] DZYUBACHYK O., BLAAS J., BOTHA C. P., STARING M., REIJNIESE M., BLOEM J. L., VAN DER GEEST R. J., LELIEVELDT B. P. F.: Comparative exploration of whole-body mr through locally rigid transforms. *International Journal of Computer Assisted Radiology and Surgery* 8, 4 (2013), 635–647. 2
- [DMA*09] DOUAUD G., MACKAY C., ANDERSSON J., JAMES S., QUESTED D., RAY M. K., CONNELL J., ROBERTS N., CROW T. J., MATTHEWS P. M., SMITH S., JAMES A.: Schizophrenia delays and alters maturation of the brain in adolescence. *Brain* 132, 9 (2009), 2437–2448. 1
- [EK06] ENNIS D. B., KINDLMANN G.: Orthogonal tensor invariants and the analysis of diffusion tensor magnetic resonance images. *Magnetic Resonance in Medicine* 55, 1 (2006), 136–146. 3
- [FH75] FUKUNAGA K., HOSTETLER L.: The estimation of the gradient of a density function, with applications in pattern recognition. *IEEE Transactions on Information Theory* 21, 1 (1975), 32–40. 3
- [GAW*11] GLEICHER M., ALBERS D., WALKER R., JUSUFI I., HANSEN C. D., ROBERTS J. C.: Visual comparison for information visualization. *Information Visualization* 10, 4 (2011), 289–309. 2
- [HPvU*16] HÖLLT T., PEZZOTTI N., VAN UNEN V., KONING F., EISEMANN E., LELIEVELDT B. P., VILANOVA A.: Cytosplore: Interactive immune cell phenotyping for large single-cell datasets. *Computer Graphics Forum* 35, 3 (2016), 171–180. 3
- [HSNHH10] HOTZ I., SREEVALSAN-NAIR J., HAGEN H., HAMANN B.: Tensor field reconstruction based on eigenvector and eigenvalue interpolation. In *Scientific Visualization: Advanced Concepts*, vol. 1 of *Dagstuhl Follow-Ups*. 2010, pp. 110–123. 2
- [JGA*02] JONES D. K., GRIFFIN L. D., ALEXANDER D. C., CATANI M., HORSFIELD M. A., HOWARD R., WILLIAMS S. C.: Spatial normalization and averaging of diffusion tensor MRI data sets. *NeuroImage* 17, 2 (2002), 592–617. 2
- [JPGJ12] JIAO F., PHILLIPS J. M., GUR Y., JOHNSON C. R.: Uncertainty visualization in HARDI based on ensembles of odfs. In *IEEE Pacific Visualization Symposium* (2012), pp. 193–200. 2
- [KEWW07] KINDLMANN G., ENNIS D. B., WHITAKER R. T., WESTIN C.-F.: Diffusion tensor analysis with invariant gradients and rotation tangents. *IEEE Transactions on Medical Imaging* 26, 11 (2007), 1483–1499. 2
- [Kin04] KINDLMANN G.: Superquadric tensor glyphs. In *Joint Eurographics / IEEE TCVG Symposium on Visualization* (2004), pp. 147–154. 2, 4
- [KSJEN*07] KINDLMANN G., SAN JOSÉ ESTÉPAR R., NIETHAMMER M., HAKER S., WESTIN C.-F.: Geodesic-loxodromes for diffusion tensor interpolation and difference measurement. In *MICCAI*. 2007, pp. 1–9. 2
- [KW06] KINDLMANN G., WESTIN C.-F.: Diffusion tensor visualization with glyph packing. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (2006), 1329–1336. 9
- [LH11] LAMPE O. D., HAUSER H.: Interactive visualization of streaming data with kernel density estimation. In *2011 IEEE Pacific Visualization Symposium* (2011), pp. 171–178. 3
- [MVB*17] MEUSCHKE M., VOSS S., BEUING O., PREIM B., LAWONN K.: Glyph-based comparative stress tensor visualization in cerebral aneurysms. *Computer Graphics Forum* 36, 3 (2017), 99–108. 2
- [PP99] PAJEVIC S., PIERPAOLI C.: 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 (1999), 526–540. 3
- [SDT05] SCHWARTZMAN A., DOUGHERTY R. F., TAYLOR J. E.: Cross-subject comparison of principal diffusion direction maps. *Magnetic Resonance in Medicine* 53, 6 (2005), 1423–1431. 1, 2
- [Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. In *IEEE Symposium on Visual Languages* (1996), pp. 336–343. 1
- [SJJ*06] SMITH S. M., JENKINSON M., JOHANSEN-BERG H., RUECKERT D., NICHOLS T. E., MACKAY C. E., WATKINS K. E., CICCARELLI O., CADER M. Z., MATTHEWS P. M., BEHRENS T. E.: Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. *NeuroImage* 31, 4 (2006), 1487–1505. 1, 8
- [SSSW13] SCHULTZ T., SCHLAFFKE L., SCHÖLKOPF B., SCHMIDT-WILCKE T.: HiFiVE: A Hilbert space embedding of fiber variability estimates for uncertainty modeling and visualization. *Computer Graphics Forum* 32, 3 (2013), 121–130. 2
- [VHNE*09] VAN HECKE W., NAGELS G., EMONDS G., LEEMANS A., SIJBERS J., VAN GOETHEM J., PARIZEL P. M.: A diffusion tensor imaging group study of the spinal cord in multiple sclerosis patients with and without T2 spinal cord lesions. *Journal of Magnetic Resonance Imaging* 30, 1 (2009), 25–34. 1, 2
- [WMK*99] WESTIN C.-F., MAIER S. E., KHIDHIR B., EVERETT P., JOLESZ F. A., KIKINIS R.: Image processing for diffusion tensor magnetic resonance imaging. In *MICCAI*. 1999, pp. 441–452. 3
- [WSC*06] WU Y., STOREY P., COHEN B., EPSTEIN L., EDELMAN R., RAGIN A.: Diffusion alterations in corpus callosum of patients with hiv. *American Journal of Neuroradiology* 27, 3 (2006), 656–660. 8
- [WWHT07] WHITCHER B., WISCO J. J., HADJIKHANI N., TUCH D. S.: Statistical group comparison of diffusion tensors via multivariate hypothesis testing. *Magnetic Resonance in Medicine* 57, 6 (2007), 1065–1074. 1
- [ZCH*17] ZHANG C., CAAN M., HÖLLT T., EISEMANN E., VILANOVA A.: Overview + detail visualization for ensembles of diffusion tensors. *Computer Graphics Forum* 36, 3 (2017), 121–132. 1, 2, 3, 4, 5, 8
- [ZSL*16] ZHANG C., SCHULTZ T., LAWONN K., EISEMANN E., VILANOVA A.: Glyph-based comparative visualization for diffusion tensor fields. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 797–806. 2, 3, 4, 5, 8
- [ZYAG06] ZHANG H., YUSHKEVICH P. A., ALEXANDER D. C., GEE J. C.: Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Medical Image Analysis* 10, 5 (2006), 764–785. 6