

Exploring brain connectivity with two-dimensional neural maps

Radu Jianu, *Member, IEEE*, Çağatay Demiralp, and David H Laidlaw, *Senior Member, IEEE*

Abstract—We introduce two-dimensional neural maps for exploring connectivity in the brain. For this, we create standard streamtube models from diffusion-weighted brain imaging data sets along with neural paths hierarchically projected into the plane. These planar neural maps combine desirable properties of low-dimensional representations, such as visual clarity and ease of tract-of-interest selection, with the anatomical familiarity of 3D brain models and planar sectional views. We distribute this type of visualization both in a traditional standalone interactive application and as a novel, lightweight web-accessible system. The web interface integrates precomputed neural-path representations into a geographical digital-maps framework with associated labels, metrics, statistics, and linkouts. Anecdotal and quantitative comparisons of the present method with a recently proposed 2D point representation suggest that our representation is more intuitive and easier to use and learn. Similarly, users are faster and more accurate in selecting bundles using the 2D path representation than the 2D point representation. Finally, expert feedback on the web interface suggests that it can be useful for collaboration as well as quick exploration of data.

Index Terms—DTI fiber tracts, Abstraction, Filtration, Path Immersion, Interaction, Coloring

1 INTRODUCTION

LOW-DIMENSIONAL point representations have been proposed recently for better interaction with fiber tracts obtained from diffusion-weighted magnetic resonance imaging (DWI) datasets [1], [2]. Driven by known embedding methods, these representations provide an interesting window into the manifold space of neural connectivity and help in fine selection of tracts. A significant drawback of point representations is, however, that coordinate axes in the low-dimensional space lack an anatomical interpretation. It is clear from evaluations in [1], [2] that having a frame of reference, anatomical or otherwise, is important for users. Motivated by this problem, we present here two-dimensional neural maps that have the desirable properties of low-dimensional representations while preserving meaningful and familiar coordinates.

DWI enables neural pathways in the *in vivo* brain to be estimated as a collection of space curves, called a tractogram. The study of tractograms (i.e., tractography) has important applications in both clinical and basic neuroscience research on the brain. Tractograms have visual complexity proportional to the intricacy of the axonal brain connectivity and, with increasing DWI resolutions, this complexity is becoming greater and greater. It is thus often difficult for practitioners to see tract projections clearly or identify anatomical and functional structures easily in these dense curve collections. This is important because, for example, a clinical study into a neurodegenerative disease typically involves selecting more than 30 TOIs manually across different datasets. Therefore, it is necessary for

tractography visualization tools to provide means to reduce and help cope with visual complexity at interaction and representation levels. We believe that two simple but powerful ideas, *abstraction* and *filtration*, can help users overcome the difficulties of visual complexity. While abstraction involves simplification and generalization, filtration here entails clustering and hierarchization. Applying these two ideas without sacrificing the intuitiveness of representations is a theme of the present work. To this end, we first obtain a hierarchy of two-dimensional neural diagrams from a whole-brain tractogram, inspired in part by illustrations in medical textbooks [3]. These can be considered immersions of neural paths in the plane. We then link the two-dimensional neural maps with the three-dimensional streamtube representations. We also create a web interface for our representation by integrating it into a standard geographical digital-map framework [4] and enhancing this with labels, statistics, and links (see Figure 1). We assess the usefulness of our approach in comparison to the 2D point representation in two studies: one anecdotal and one quantitative. Anecdotal study results indicate that our new representation is more intuitive and easier to use and learn. Results of the quantitative study shows that users are faster and more confident with the 2D path representation than with the 2D point representation. Also, users characterized the web interface with its digital-map-like interaction as useful for collaboration and quick exploration of data.

In this context, our work has three main contributions, which are primarily in the field of tractography visualization. Our first contribution is a two-dimensional path representation of tractography datasets that, in contrast to the previously proposed 2D point representation, preserves anatomically meaningful and familiar coordinates while possessing the advantages of low-dimensional representations. By construction, this representation is conducive to

• R. Jianu, C. Demiralp and D. H. Laidlaw are with the Department of Computer Science, Brown University.
E-mail: jr,cad,dhl@cs.brown.edu

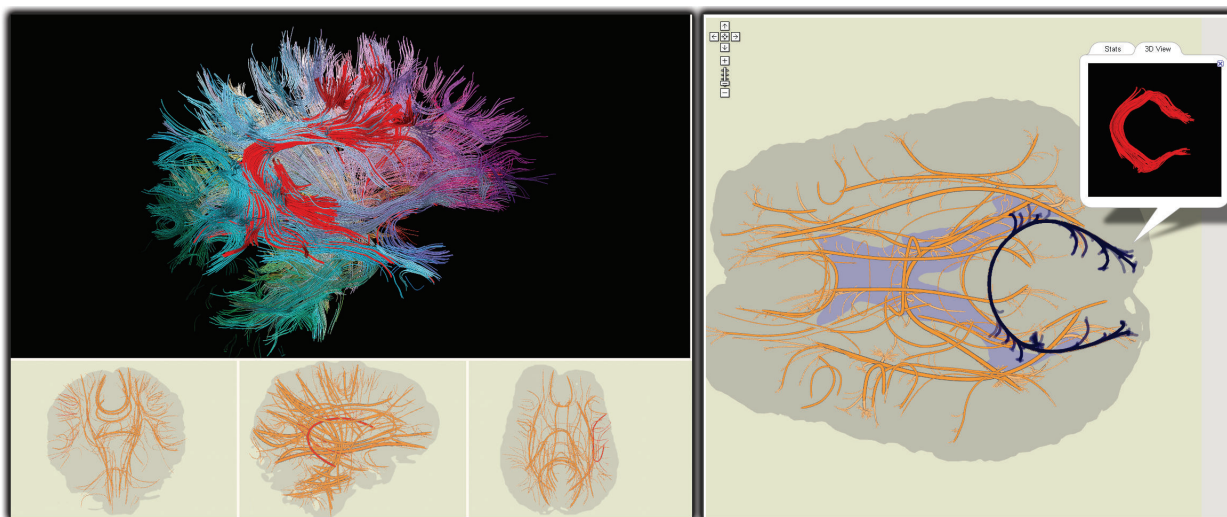


Fig. 1: Schematic planar projections of DTI tractograms as part of a standalone interactive system (left) and as a web-accessible digital map (right). The digital map interface easily incorporates any tract-associated information, including labels, links, metrics, and statistics. Shown in the pop-up window on the right is the “brain view” of the selected tract.

be navigated and interacted with like a digital geographic map. Consequently, the second contribution is a web-interface that allows the exploration and annotation of tractography datasets like online geographic maps. This opens up exciting possibilities of, not just sharing, but also enriching tractography datasets by tapping into the vast knowledge base of the Web. Our third contribution is an anecdotal and quantitative comparison of the 2D point and path representations, results of which provide a handle on the relative merits of each representation.

DWI is the source of all the data used in the work presented here. We first give an elementary introduction to DWI and then briefly discuss related work on tractography visualization.

2 RELATED WORK

Diffusion-weighted magnetic resonance imaging (DWI) measures the diffusion rate of water molecules in biological tissues *in vivo* [5]. Since tissue characteristics, geometric or otherwise, at a given point affect the diffusion rate, measured diffusion-rate information is an indicator of the tissue characteristics at the point. In particular, water in fibrous tissues such as brain white matter (i.e., a collection of myelinated axons) diffuses faster along fibers than orthogonal to them. Therefore, it is possible to estimate fiber trajectories computationally using diffusion models such as the tensor model that quantify anisotropic diffusion. Diffusion imaging based on fitting second-order tensors to DWI sequences is known as diffusion-tensor magnetic resonance imaging (DTI) [6]. Fiber trajectories are computed from DTI data by integrating bidirectionally along the principal eigenvector of the underlying tensor field. This process, called fiber tracking, yields a dense collection of integral

curves (i.e., a tractogram). All the tractograms used in our work were obtained by fiber-tracking in DTI volumes.

Tractograms are often visualized with streamlines or variations of streamlines in 3D [7], [8]. Reflecting the intricacy of the connectivity in the brain, these 3D models are generally visually dense. Consequently, typical interaction tasks over tracts, such as fine bundle selection, are often difficult to perform and have been a focus of recent research [9], [10]. Planar point representations have been proposed for improving interaction with DTI fiber tracts [1], [2]. The representation introduced here is also a projection of fiber tracts into a plane, but as planar curves rather than points.

The complex structure of tractography datasets has motivated earlier work to apply clustering techniques to tractograms (e.g., [11], [12], [13], [14], [15]). The success of a tract clustering is often determined by the degree to which the similarity measure used in the process can capture anatomical features that are of interest to a specific user. A comparison of some of the existing tractography clustering methods and similarity measures can be found in [16]. Clustering is an important tool for creating abstractions and filtrations in general and is central to our representation in particular. We use an agglomerative hierarchical clustering algorithm. The output of the clustering algorithm is a hierarchical tree called dendrogram. The height of the tree can be thought as the radius of the bounding ball of the dataset—in the units of the similarity measure used. Any horizontal cut on this tree provides a clustering of the dataset. To exemplify, the root node can represent a clustering where all data-points are contained in a single cluster. Conversely, the leaf nodes correspond to a clustering where each data point is a cluster. The similarity measure parametrizes the nested space of clusterings. In this sense, the dendrogram embodies a filtration of the dataset.

Use of simplified, multiscale, and multiview representations has a long history in interaction and visualization research (e.g., [17], [18], [19]). For example, multiple orthogonal 3D views is fairly common in computer aided design (CAD), modeling, and animation applications. Similarly, use of 2D orthogonal cross-sectional views is a de-facto standard in volumetric medical data visualization [20]. Our system interface of multiscale orthogonal 2D neural path projections linked with each other and with a view of the 3D model can be considered an implicit combination of these earlier research ideas as well as the current standard practices.

Our 2D rendering of neural paths is partly inspired by 2D illustrations of white-matter structures in medical textbooks [3]. The appeal of hand-drawn illustrations can be attributed to the power of abstraction based on perceptual simplification and generalization. Therefore, there have been several works in illustrative visualization [21], [22], [23]. Also, stylized line renderings of general streamlines and tractograms in 3D have been used in the past [24], [25]. Recently, Otten *et al.* applied illustrative rendering techniques to visualize white-matter bundles in 3D to ease the exploration of complex datasets [26]. We use simple, contoured 2D spline curves to draw our neural paths representations.

One of the advantages of 2D projections is that they can be naturally integrated into a web-based digital geographic map framework. Basic data visualization has been available on the web for many years but was usually limited to traditional techniques such as bar graphs and charts. More recently, however, visualization research started targeting this environment and advanced applications have emerged. ManyEyes [27] paved the way for everyday data visualization, with subsequent studies such as [28] and [29] proving the need for accessible web visualization. While web-development toolkits such as [30] greatly aid web visualization development, large-scale web-visualization is limited by inherent browser capabilities, as demonstrated in [31]. Alternatively, stand-alone systems have been made available in applet form or can be run as client applications directly from websites. However, users still must control the parameters involved in producing visualizations, specify their data queries and learn the system features. This often constitutes an undesirable overhead. Yet another approach, most similar to our work from an implementation standpoint, is to use Ajax (asynchronous JavaScript and XML) technology to perform the rendering on the server side and serve images asynchronously to the client browser. The essential difference between the present work and traditional offline visualization systems is that the former separates interaction and display from rendering and computation. Our brain maps differ by eliminating users' efforts in creating visualizations and assigning this task to experts and by using the Google Maps API, which is an Ajax framework for interactive display of pre-rendered images. Closest to our work in this latter aspect are X:MAP [32] and Genome Projector [33], which present genome browser




Fig. 2: Schematic tract-cluster representation. (Top) 2D projections of a tract-bundle, with an associated centroid curve (orange), are determined from a hierarchical clustering of initial 3D tracts. (Middle) The centroid curve is smoothed by a spline and the endpoints of non-centroid curves are clustered using their initial 3D coordinates (four clusters); for each cluster, three control points linking the center of the cluster to the centroid spline are computed. (Bottom) Splines are run from each curve endpoint through the control points of its corresponding cluster.

tools implemented using the Google Maps API. We extend this idea to a new domain and demonstrate its usefulness for tractography datasets.

3 METHODS

We create neural map representations in five basic steps. First, we obtain a whole-brain tractogram by fiber tracking in a diffusion-tensor volume fitted to a given DWI brain sequence. Second, we compute similarities between all pairs of tracts within the tractogram, obtaining a similarity (or affinity) matrix. Third, using the similarity matrix from the previous step, we run a hierarchical clustering algorithm on the tractogram, obtaining a clustering tree (i.e., dendrogram). Fourth, primarily using cluster centroids at different levels of the clustering tree as pivots, we create hierarchical projections of tractograms onto the major orthogonal planes. Finally, we render these 2D curves stylistically using heuristics determined by the topology and geometry of the corresponding tracts and tract clusters, creating neural maps. We give details on these steps in subsequent sections.




Fig. 3: Depth ordering of 2D paths. For each segment of a 2D spline, we locate a corresponding segment on the 3D curve from which the spline was derived by traveling the same fractional distance along both curves. The depth of the 2D segment is the same as the depth of the middle of its corresponding 3D segment.

3.1 Image Acquisition and Fiber Tract Generation

DWI brain datasets used in this paper were acquired from healthy volunteers on a 1.5T Siemens Symphony scanner with the following acquisition parameters in twelve bipolar diffusion-encoding gradient directions: thickness = 1.7 mm, FOV = 21.7 cm \times 21.7 cm, TR = 7200 ms, TE = 156 ms, b = 1000, and NEX = 3. For each DWI sequence, the corresponding DTI volume was then obtained by fitting six independent parameters of a single second-order tensor at each voxel to the twelve measurements from the DWI sequence [6]. We generate fiber-tract models of the whole brain by integrating (second-order Runge-Kutta integration) the major eigenvector field of the diffusion tensor field bidirectionally starting at seed points. We integrate with a constant step size of 0.5 mm and stop integration when a gray-matter area is reached or the linear anisotropy or signal-to-noise (SNR) ratio at the current point passes predetermined thresholds. Details of our integration scheme can be found in [8].

3.2 Similarities Between Fiber Tracts

We quantify the similarity between two tracts using the weighted chamfer distance discussed in [34]. This measure tries to capture how much any given two tracts follow a similar path, while giving more weight to the points closer to tract ends. We compute distances between each pair of fiber tract using $\lambda = 0.5$ as described in [34] and assemble the measures to create a distance matrix used in clustering and projection.

While our approach is independent of a particular similarity measure, good results in practice require a good similarity measure – one that reflects users’ understanding of the similarity between data points (i.e., tracts) and works well for the task at hand.

3.3 Clustering and Projection

For a given 3D DTI streamtube model we create schematic views of major tract bundles projected on a few selected planes. In this work we choose the sagittal, coronal, and transverse planes as our main modes of representation.

We first compute a clustering tree using an average-linkage hierarchical clustering algorithm on the tract distance matrix (e.g., [35]). We choose the average-linkage criterion because it is less sensitive than the minimum-linkage to broken tracts due to tracking errors. We obtain a clustering of tracts by manually setting a cut threshold on the dendrogram. This threshold can be also interactively changed by user to control the coarseness of the clustering. A constant cut at 60% of the clustering tree’s height gave consistent results across the six datasets we experimented with.

Next, we create simple orthogonal projections of tracts on each plane. We cull out tracts that do not contribute significantly to the projection. If the ratio of projected tract length to true tract length is under a threshold value, we remove the tract from the corresponding cluster. We set the culling threshold to 0.65 for the projections used in our experiments.

Finally, we compute a centroid for each cluster by choosing the tract with the smallest maximum distance to any other tract in the cluster. We found that for illustration purposes it is desirable to avoid broken tracts. We therefore weigh the centroid selection to favor longer tracts by dividing the maximum distance from each tract to any other tract by the tract’s length.

3.4 Rendering

We opted for an illustrative rendering of brain projections. Illustrative visualization uses abstraction to reveal structure in dense visualizations and to harness scientists’ familiarity with textbook representations [36]. Both criteria apply to white matter tractograms: fiber bundles provide a natural abstraction of 3D anatomy that avoids the clutter of large streamtube collections, while textbook illustrations [3] shape the intuition of neuroscientists. These advantages have also been recognized and explored by Otten et al [26].

The rendering assumes a given clustering with assigned centroid tracts, which can be computed as described in the previous section. Our approach is inspired by Holten’s hierarchical edge bundles [37] in attempting to group all fiber-tracts from a bundle into one, visually salient structure. However, hierarchical edge bundles operate on abstract connections that are unconstrained by concrete geometrical

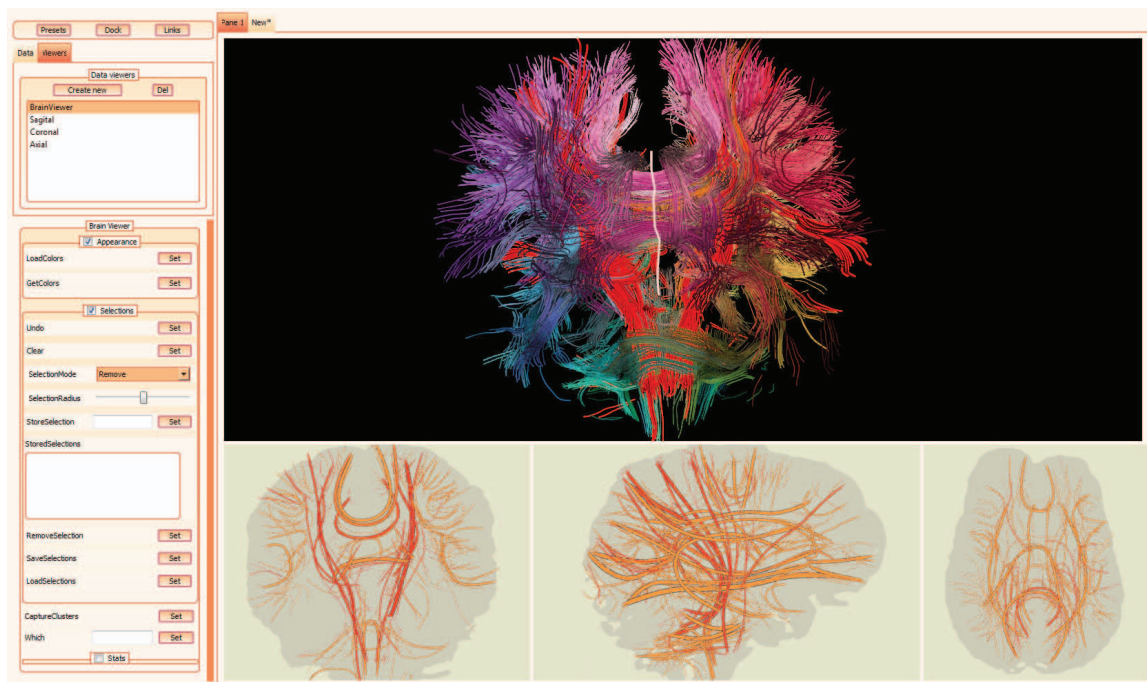


Fig. 4: An interactive analysis system using linked views and planar tract-bundle projections. Three planar representations, along the coronal, transverse and sagittal planes (bottom panels), are linked to a 3D streamtube model (upper left) and a 2D point embedding of tract similarities (upper right). Selections in the projection views can be performed by clicking or cutting across cluster curves and are mirrored in the 3D view. Points corresponding to the selected tracts are interactively embedded into the plane and used to refine selections at tract level.

shapes. They can therefore be drawn according to visual aesthetics principles alone. Conversely, fiber tract paths play important anatomical functions and should be preserved in tractogram visualizations. To this end, we perform our bundling by routing tracts along the path of the most representative tract in their bundle. Thus, the centroid tracts will define a schematic neural skeleton on top of which the non-centroid tracts are scaffolded.

Projections of centroid curves are smoothed prior to rendering to achieve a schematic representation and to reduce clutter. This is done by sampling a number of evenly distributed control points (five in our implementation) along the tract projection and using them as control points for a spline. In our implementation the spline is piecewise cubic and consists of 30 segments. The thickness of a centroid curve is proportional to the square root of the number of tracts in the bundle.

Once centroid tracts are represented as 2D splines, end-points of non-centroid curves are linked to their cluster's centroid spline following the procedure illustrated in Figure 2. First, the end-points of non-centroid curves in a bundle are clustered based on the end-points' initial 3D coordinates. Two endpoints are placed in the same cluster if the distance between them is less than 2 mm. Then, for each such endpoint cluster we compute three control points that link the geometrical center of the endpoint cluster to the centroid spline: the first point is the center itself, the second is a point on the centroid spline closest to the center point,

and the third is determined by traveling from the second point down the centroid spline, towards each curve's other endpoint, for a predefined distance (e.g., half of the distance between the first two points). Ultimately, splines are run from each tract endpoint through its cluster's three control points, thus linking each endpoint to the centroid path. The thickness of these endpoint linkage splines gradually increases from unit thickness (i.e., single-tract thickness) at the tract endpoint to a thickness proportional to the square root of the endpoint cluster size, where it merges with the centroid spline.

We depth-order spline segments so that 2D centroid splines crossings can indicate the depth ordering of their corresponding 3D shapes. The depth ordering is done differently for centroid splines and non-centroid splines, since while centroid curves are close representations of actual 3D tracts, non-centroid curves are abstract representations obtained through the process described above. Furthermore, the depth ordering is approximate (as discussed in the following paragraph) and may produce artifacts.

For centroid splines, the depth of a spline segment is computed by finding a matching segment on the 3D tract from which the spline was derived, and taking the depth of that segment's center (Figure 3). The matching segment on the 3D tract has its endpoints at the same fractional distance from the start of the 3D tract as the 2D segment's distance from the start of the 2D spline. This per-segment ordering was chosen because of the intricacy of white mat-

ter tractograms. Tracts often wrap around each other such that a correct per-tract depth ordering cannot be determined. Treating each curve segment independently maximizes the probability that the 2D rendering remains truthful, at least within the resolution of the tract segmentation. Conversely, non-centroid splines are completely abstract 2D representations. The depth of any non-centroid spline is determined by averaging the depth of the corresponding 3D tract. Wrapping fiber tracts are therefore not captured by this latter process.

Finally, bundle-color, texture or thickness can be used as additional depth cues. While we have not fully integrated and evaluated such encodings in our current prototypes, we have experimented with color cues and found those to be useful.

In the following two sections, we give details on how we use 2D neural path representations as part of an interactive application and as standalone digital maps.

3.5 Interactive Application

We expect a typical use of low-dimensional representations to be as part of interactive applications where views and interactions of conventional representations are linked with that of low-dimensional representations. Therefore, we have developed an interactive visualization system using the 2D path representation to demonstrate this mode of use (see Figure 4). We link the three views of the 2D path representations with a view of the 3D streamtube representation. We color streamtubes by embedding the tract similarity matrix in the $L^*a^*b^*$ color space [2].

The clustering cut threshold that defines the specificity of the projected bundles can be altered interactively during visualization. Tract clusters in the planar projections can be selected by drawing line segments that select intersecting bundles. A selection in any of the planar views is mirrored in the 3D model view as well as all other 2D projections.

In addition to the standard 3D viewing interactions, we have two basic selection/deselection interactions on streamtube models: sphere-selection and brushing. Sphere-selection, like box-selection, enables user to select the intersecting tracts by moving a sphere of desired radius. Brushing lets user draw 2D curve on the viewing plane and select the intersecting tracts.

Although both sphere-selection and brushing can be used to further prune the current selection, they cannot be used to grow it. For this purpose, we provide a selection-growing interaction that gradually adds tracts closest to the current selection. Proximity is determined again by the distance measure discussed above. Selections can be saved for future analysis. Moreover, statistics such as average tube length, number of tubes or average fractional anisotropy can be computed interactively on sets of selected tracts.

3.6 Digital Map Interface

Brain mapping is one of the quintessential problems in neurosciences. We believe that a geographical map metaphor is well-suited to the visualization and analysis of results obtained in this direction. Therefore, having a representation of the brain that is viewed, interacted, queried, and enriched like an online geographical map was one of the motivations behind our creation of the 2D path representation.

For this, we use the Google Maps API, an Ajax framework used to render large maps, to interactively display our tractogram maps on the web. The Google Maps API receives input image data in the form of a set of small images, called tiles, that when assembled together form the different zoom levels of the map. Each zoom level, z , consists of a rectangular grid of tiles of size $2^z \times 2^z$. The API decodes the zoom level and coordinates of the currently viewed map region to retrieve and display the visible tiles. The developer can load a custom set of tiles in the API by implementing a callback function that translates numerical tile coordinates and zoom level into unique paths to the custom tiles.

The API provides basic functionality such as zooming and panning and allows programmatic extension or customization with markers and polyline overlays, information pop-ups and event management. The API can easily be integrated into any webpage supporting Javascripts.

Our visualization system can render our 2D projections into a set of image tiles instead of the screen. For each cluster, including both tract-bundled and endpoint clusters, we export information required for interaction and browsing. Selection information consisting of evenly spaced points along splines and thickness radii for splines contained in a cluster is exported. In line with the tile paradigm, instead of exporting this information to a single large file, we divide it geometrically across corresponding tiles and write it as multiple tile-content text files. Upon user selection, the content file of a clicked tile is fetched from the server and its data analyzed for an intersection. This approach avoids loading and searching through large files. A valid cluster selection is marked on the map with polyline overlays running over tract splines contained in the selected cluster (see Figure 5). For this purpose, spline coordinates for each cluster are exported to files indexed by a unique cluster identifier.

Finally, for each tract cluster we export a variety of meta-data accessible during map browsing in information boxes, as shown in Figure 5. A short description and links to the most relevant publications or research can be manually added for major tracts. A few 3D poses of each tract bundle are prerendered and exported as animated GIF images, indexed by the cluster identifier. Statistical data, in both textual and graphical form, are computed for each cluster and written as HTML content to cluster indexed files. This information is loaded and displayed in tabbed information boxes at the user's request.




Fig. 5: DTI tractography data projected onto the sagittal, coronal and transverse planes. Major tract bundles are represented schematically by their centroid tract; individual tracts in bundles are linked from the centroid tract to their projected end points. Zooming in allows access to smaller clusters of tracts. Bundles can be selected and pre-computed statistical data along with 3D views of the tract bundle (“brain view”) can be displayed.

3.7 Implementation

We have implemented a prototype of our method in C++ using G3D and Qt libraries [38], [39], [4]. The web interface based on the Google Maps API of this prototype can be accessed via the url link [40]. Note that the current web interface implementation provides only a single view at a time. It is, however, possible to have linked multiple views, as in the standalone application, using a cookie-pooling mechanism.

4 USER EVALUATION

We compared our method to a 2D point-based representation both anecdotally and quantitatively.

4.1 Anecdotal Study: Methods and Results

In the anecdotal study we showed a prototype that implements both 2D point and neural map representations to three neuropsychologists. They were all interested in the relationship between fiber tracts and cognitive and behavioral function in the brain. Similarly, all have used computational tools for analyzing DTI data, though only one of them had used fiber-tract visualization tools in his clinical research. The participants had research interests in vascular cognitive impairment, early Alzheimer’s disease, and HIV, focusing on specific tracts such as the corpus callosum (CC), frontal lobe, basal ganglia, cingulate bundle, superior and inferior longitudinal fasciculi, anterior internal capsule, and the uncinate fasciculus.

Our anecdotal evaluation protocol was straightforward: we demonstrated the prototype while asking questions and collecting participant’s feedback. Two of the experts also tried both interfaces themselves by selecting a set of major TOIs, the CC, cingulate bundle, uncinate anterior internal capsule, and the corticospinal tract. There was agreement that our new interface is more intuitive and easier to use and learn than the 2D point representation.

Our experts also found the web interface with the digital map interaction useful. Although they believed that the standalone application with linked representations would remain necessary for quantitative analyses that require fine selection, they thought the web accessibility opened up interesting possibilities. They were particularly excited about browsing through datasets while commuting or at home, of quickly inspecting unfamiliar datasets, and of sharing such visualizations with collaborators.

4.2 Quantitative Study

In this experiment we compared the point and path representations by measuring user performance on a bundle selection task with the two representations.

4.2.1 Application

We compared our path-visualization system (Section 3.5) against the one presented in [1], which consists of a 3D streamtube model view and a linked 2D point representation. The system offers a brush tool that works similarly to ours in 3D and as a lasso tool on the 2D point

representation. Users can select tracts or points and then remove them or, conversely, remove everything else from a current selection.

4.2.2 Participant Pool

The four subjects were all familiar with neuroanatomy and tractography. They also had experience with one or more tractography visualization tools. Our first subject was a neuroscience graduate student, working on tracing white-matter tracts from frontal subregions to basal ganglia and the medial temporal lobe. Our second user was a postdoc in neuropsychology and had five years of experience with DWI in clinical research. This user, who participated in the anecdotal study as well, studied white matter variation with neurodegenerative diseases as specified above. Our third subject was a biomedical engineering graduate student and had significant tract-selection experience working as a rater for a neuroscientist. Our last subject was a computer-science graduate student doing research on computational DWI algorithms. Two of the users were male and two female.

4.2.3 Task

We measured user performance on bundle selection, a typical real world task in tractography tools. Users were asked to select three major bundles, the cingulate bundle (cb), corticospinal tract (cst), and right superior longitudinal fasciculus (slf), in two different brain datasets. We choose these bundles because they represent the easy-to-hard selection-difficulty range well and were used for evaluation in [1].

For each system, we explained to users the underlying visualization concepts and demonstrated the basic interactions, mainly involving brushing on 2D and streamtube representations. After this introduction, users were asked to select the bundles (cb, cst, and slf) on two different training datasets. Following training, the users performed the task on two different test datasets while we collected their task-completion times. After each selection they provided subjective confidence estimate in the range 1-5 (1: not confident, 5: very confident) for their selection. They could give fractional estimates. After completing the task on both systems, users were asked to complete a post-questionnaire giving qualitative feedback on their experience. Half the users performed the task first on the 2D point-representation tool and the other half on the 2D path tool.

4.2.4 Factors and Measures

The sole factor considered in our quantitative experiment was the type of low-dimensional representation: 2D point and 2D path. All subjects used both types of representation. We recorded the users' bundle-selection times and subjective confidence values as measures of performance.

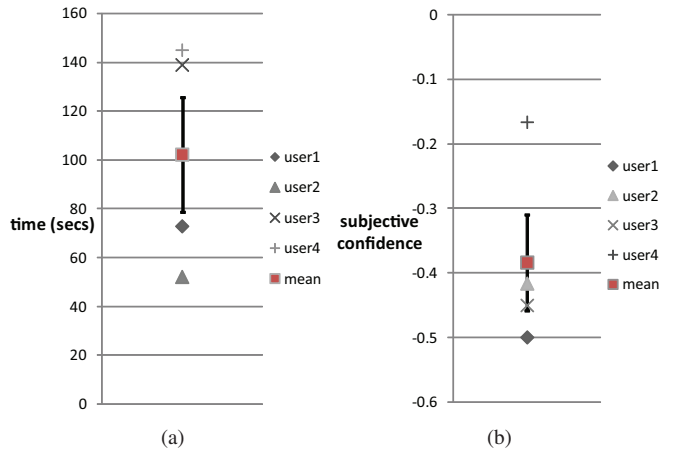


Fig. 6: Per-user differences between (a) time and (b) confidence measurements with the two tools. Differences are obtained by subtracting 2D-point-tool performance values from 2D-path-tool performance values. Red squares show the mean performance difference between the tools. Errors bar around the red squares indicate the standard error of per-user differences.

	time (secs)				confidence			
	cb	cs	slf	mean	cb	cst	slf	mean
2D point	227	361	234	274	4.1	3.3	3.1	3.5
2D path	136	165	215	172	4.1	3.8	3.7	3.9

TABLE 1: User performance on bundle selection task.

4.2.5 Results

In order to understand if the differences between user performances on the two tools were significant, we ran the paired t-test on our measurements. Results show that users were significantly faster on the 2D path tool than the 2D point tool ($p = 0.02$). Users were also significantly more confident with using the 2D path representation than the 2D point representation ($p = 0.01$). Table 1 summarizes users' overall and per-bundle mean performances on each tool. Figure 6 shows the difference between the means of performance measures per user (2D-path-performance values are subtracted from 2D-point-performance values) and the mean over users. Error bars indicate the standard error of per-user differences.

We observed some interaction patterns worth reporting. We noticed two distinct selection strategies used with the 2D path tool. Two of the users consistently brushed over large areas of the projection to ensure that the targeted bundle was selected and then relied on the 3D view to clean up the selection. The other two users aimed for fine selections in the 2D projections and then inspected the 3D view to determine whether any fibers the selection. They added the missing tracts using short, targeted brushstrokes and then removed tubes that were erroneously added during this operation. These users seemed to have a better under-

standing of the mapping between the 3D view and the 2D projections, perhaps explaining the difference in strategies.

All subjects used the 2D point representation relatively rarely. The most common operation was to remove points they were completely confident were not part of the selection (e.g., half of the brain, or peripheral U-shaped bundles). However, in the absence of a clear contextual mapping between the 2D point and streamtube views, subjects were hesitant to perform bold operations in 2D, at least in the short run.

5 DISCUSSION

It is important to note that our representation relies on the anatomical fidelity of the intermediate results at each step. For example, broken trajectories due to fiber tracking errors or noise can reduce the effectiveness of the representation. In general, the overall robustness of our representation to noise will be determined by the fiber tracking algorithm used to generate the tractograms.

Our method also expects the clustering algorithm and similarity measure to provide anatomically plausible results. However, it is difficult for a single distance measure to capture the anatomical similarity completely. Furthermore, on the same data, a good similarity measure for one purpose can be entirely irrelevant for another. While the choice of similarity measure makes clustering a subjective task, clustering algorithms themselves also have intrinsic limitations [41].

On the other hand, we believe that our representation, which is driven by cluster centroids, is stable. It is important to observe that our cluster centroids are existing tracts in datasets. They are not, for example, the medial-axes of the corresponding cluster surfaces, which tend to be less stable [42]. Therefore, we expect small perturbations to thresholds or datasets to cause only ignorable changes in the representation.

One of the limitations of our visualization pipeline is its reliance on adjustable thresholds. Each computational step introduces variability that is hard to account for automatically. More robustness in tractography, similarity measures and clustering could lead to anatomically motivated thresholds. However, we note that variability is also present in neuroscientists' assessments of what constitutes a valid bundle and that different tasks often require different clustering granularity. This suggests that manually adjustable thresholds are perhaps unavoidable. Still, one can imagine providing adaptive tools based on persistence that can help users to adjust these the thresholds more efficiently, minimizing the number of trial-and-error iterations.

Another potential limitation of the current method is that bundles surrounded by other bundles similar in orientation and shape may not be clearly visible. While we have not found this to be an issue in practice, moving projection

planes along major axes while restricting the projecting tracts to a volumetric window around the projection plane can help solve potential problems.

Also, results of user studies should be taken with a grain of salt. In general, it is difficult to run experiments that vary one factor while keeping all the other factors constant. For example, an earlier study [1] compared the 2D point representation tool with other tractography applications and reported that users were faster with the former. However, in our user evaluation, we observed that users rarely used the 2D point representation and the brush tool dominated their interaction. This brings up an open question of whether the performance difference in the reported evaluation was due mainly to the brush tool or to the 2D point representation. An experiment that replaced the brush tool with a more standard box-selection tool, say, might resolve this question. In either case, we believe that abstract representations, including the 2D point representation, are useful in the long run, as users gain more experience with the mapping between brain tractograms and low-dimensional representation primitives.

While DWI is the only imaging protocol to estimate the brain neural architecture *in vivo*, there are *in vitro* imaging techniques, such as the three-dimensional electron microscopy particularly used in the emerging field of connectomics, with which it is possible to extract neural structures on much smaller scales (e.g., individual axon bodies) [43]. We believe that the general ideas as well as specific techniques presented in this paper can extend to the visualization and analysis of visually complex axonal structures originating from these high-throughput imaging techniques.

6 CONCLUSIONS

We have presented 2D neural maps, an intuitive, low-dimensional representation for tractograms, that facilitates exploration and analysis of brain connectivity. Our representation is essentially an application of abstraction and filtration concepts to tractograms with preservation of anatomical context in mind. We achieve abstraction by both geometrically and topologically simplifying and generalizing fiber tracts with 2D schematic curves. We create filtrations of tractograms by computing hierarchical clustering trees. These help create better abstractions as well as provide a multiscale view of data, which is important in reducing visual complexity and noise. Both qualitative and quantitative user evaluation results suggest that our new representation can significantly improve user performance in fiber bundle selection.

We also introduced a novel way of making tractography data accessible by publishing neural maps online through a digital map framework. Our representation is conducive to such a geographic map interface by construction. This interface leads to new possibilities for enriching tractography datasets using the mass knowledge base available on the

web. User feedback indicates that our web interface can be particularly useful for browsing unfamiliar datasets quickly, for analysis tasks that do not require fine selection and for sharing data in collaborative settings.

ACKNOWLEDGMENT

We thank Song Zhang for generously providing us the 2D point representation tool used for comparison in our quantitative evaluation. This work was supported by NIH grant 1R01EB00415501A1.

REFERENCES

- [1] 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 TVCG (Proc. of Visualization)*, 2009.
- [2] R. Jianu, C. Demiralp, and D. Laidlaw, "Exploring 3D DTI fiber tracts with linked 2D representations," *IEEE TVCG (Proc. of Visualization)*, vol. 15, no. 6, pp. 1449–1456, 2009.
- [3] H. Gray, *Anatomy of the Human Body*. Lea & Febiger, 1918.
- [4] GoogleMapsAPI, <http://code.google.com/apis/maps/>.
- [5] J. E. Tanner, "Transient diffusion in system partitioned by permeable barriers. Application to NMR measurements with a pulsed field gradient," *Journal of Chemical Physics*, vol. 69, no. 4, pp. 1748–1754, January 1978.
- [6] P. J. Basser, J. Mattiello, and D. LeBihan, "Estimation of the effective self-diffusion tensor from the NMR spin echo," *J Magn Reson B*, vol. 103, no. 3, pp. 247–254, March 1994.
- [7] S. Mori and P. Van Zijl, "Fiber tracking: principles and strategies—a technical review," *NMR in Biomedicine*, vol. 15, no. 7-8, pp. 468–480, 2002.
- [8] S. Zhang, C. Demiralp, and D. Laidlaw, "Visualizing diffusion tensor MR images using streamtubes and streamsurfaces," *IEEE TVCG*, vol. 9, no. 4, pp. 454–462, 2003.
- [9] D. Akers, "Wizard of Oz for participatory design: Inventing an interface for 3D selection of neural pathway estimates," in *Proceedings of CHI 2006 Extended Abstracts*, 2006, pp. 454–459.
- [10] D. Akers, A. Sherbondy, R. Mackenzie, R. Dougherty, and B. Wandell, "Exploration of the brain's white matter pathways with dynamic queries," in *Proc. of Visualization*, 2004, pp. 377–384.
- [11] J. S. Shimony, A. Z. Snyder, N. Lori, and T. E. Conturo, "Automated fuzzy clustering of neuronal pathways in diffusion tensor tracking," in *Proc. Intl. Soc. Mag. Reson. Med.*, 2002.
- [12] A. Brun, H. Knutsson, H. J. Park, M. E. Shenton, and C.-F. Westin, "Clustering fiber tracts using normalized cuts," in *Seventh International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'04)*, ser. Lecture Notes in Computer Science, Rennes - Saint Malo, France, September 2004, pp. 368–375.
- [13] L. O'Donnell and C.-F. Westin, "White matter tract clustering and correspondence in populations," in *Eighth International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'05)*, ser. Lecture Notes in Computer Science 3749, Palm Springs, CA, USA, October 2005, pp. 140–147.
- [14] 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*, vol. 14, pp. 1044–1053, 2008.
- [15] 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*, vol. 12, no. 2, pp. 191–202, 2008.
- [16] B. Moberts, A. Vilanova, and J. J. van Wijk, "Evaluation of fiber clustering methods for diffusion tensor imaging," in *Procs. of Vis'05*, 2005, pp. 65–72.
- [17] R. Stoakley, M. Conway, and R. F. Pausch, "Virtual reality on a wim: Interactive worlds in miniature," in *CHI*, 1995, pp. 265–272.
- [18] S. K. Card, J. D. Mackinlay, and B. Shneiderman, Eds., *Readings in information visualization: using vision to think*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1999.
- [19] C. Ware, *Information visualization: perception for design*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2000.
- [20] W. Schroeder, K. M. Martin, and W. E. Lorensen, *The Visualization Toolkit*, 2nd ed. Prentice-Hall, Inc., 1998.
- [21] D. Ebert and P. Rheingans, "Volume illustration: non-photorealistic rendering of volume models," in *Proceedings of Visualization 2000*, Oct 2000, pp. 195–202.
- [22] A. Lu, C. J. Morris, J. Taylor, D. S. Ebert, C. Hansen, P. Rheingans, and M. Hartner, "Illustrative interactive stipple rendering," *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, pp. 127–138, 2003.
- [23] W. Li, L. Ritter, M. Agrawala, B. Curless, and D. Salesin, "Interactive cutaway illustrations of complex 3d models," *ACM Trans. Graph.*, vol. 26, July 2007. [Online]. Available: <http://doi.acm.org/10.1145/1276377.1276416>
- [24] C. Stoll, S. Gumhold, and H.-P. Seidel, "Visualization with stylized line primitives," *Visualization Conference, IEEE*, vol. 0, p. 88, 2005.
- [25] M. H. Everts, H. Bekker, J. B. Roerdink, and T. Isenberg, "Depth-dependent halos: Illustrative rendering of dense line data," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 1299–1306, 2009.
- [26] R. Otten, A. Vilanova, and H. Van De Wetering, "Illustrative White Matter Fiber Bundles," *Computer Graphics Forum*, vol. 29, no. 3, pp. 1013–1022, 2010.
- [27] F. Viegas, M. Wattenberg, F. Van Ham, J. Kriss, and M. McKeon, "ManyEyes: a site for visualization at internet scale," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, p. 1121, 2007.
- [28] F. Viégas, M. Wattenberg, M. McKeon, F. Van Ham, and J. Kriss, "Harry Potter and the meat-filled freezer: A case study of spontaneous usage of visualization tools," in *Proc. HICSS*, 2008.
- [29] C. Danis, F. Viegas, and M. Wattenberg, "Your place or mine? Visualization as a community component," in *Proceedings of CHI*, 2008.
- [30] M. Bostock and J. Heer, "Protovis: A graphical toolkit for visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, no. 6, pp. 1121–1128, 2009.
- [31] D. Johnson and T. Jankun-Kelly, "A scalability study of web-native information visualization," in *Proceedings of Graphics Interface*, 2008, pp. 163–168.
- [32] T. Yates, M. Okoniewski, and C. Miller, "X: Map: annotation and visualization of genome structure for Affymetrix exon array analysis," *Nucleic Acids Research*, vol. 36, no. Database issue, p. D780, 2008.
- [33] K. Arakawa, S. Tamaki, N. Kono, N. Kido, K. Ikegami, R. Ogawa, and M. Tomita, "Genome Projector: zoomable genome map with multiple views," *BMC bioinformatics*, vol. 10, no. 1, p. 31, 2009.
- [34] C. Demiralp and D. H. Laidlaw, "Similarity coloring of DTI fiber tracts," in *Proceedings of DMFC Workshop at MICCAI*, 2009.
- [35] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. Wiley-Interscience, 2000.

- [36] I. Viola, M. Groller, M. Hadwiger, K. Buhler, B. Preim, M. Sousa, D. Ebert, and D. Stredney, "Illustrative visualization," in *IEEE Visualization*, 2005, p. 124.
- [37] D. Holten, "Hierarchical edge bundles: Visualization of adjacency relations in hierarchical data," *IEEE Transactions on Visualization and Computer Graphics*, pp. 741–748, 2006.
- [38] G3D, <http://g3d-cpp.sourceforge.net/>.
- [39] Qt, <http://www.qtsoftware.com/>.
- [40] <http://graphics.cs.brown.edu/research/scviz/newbraininteraction/>.
- [41] J. M. Kleinberg, "An impossibility theorem for clustering," in *NIPS*, 2002, pp. 446–453.
- [42] D. Attali, J.-D. Boissonnat, and H. Edelsbrunner, "Stability and Computation of Medial Axes - a State-of-the-Art Report," in *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Springer Berlin Heidelberg, 2009, ch. 6, pp. 109–125. [Online]. Available: <http://hal.archives-ouvertes.fr/docs/00/46/86/90/PDF/04-medial-axes.pdf>
- [43] W.-K. Jeong, J. Beyer, M. Hadwiger, A. Vazquez, H. Pfister, and R. T. Whitaker, "Scalable and interactive segmentation and visualization of neural processes in EM datasets," *IEEE Transactions on Visualization and Computer Graphics*, vol. 15, pp. 1505–1514, 2009.



Radu Jianu Radu Jianu received the Dipl. Ing. degree in Computer Science in 2005 from Polytechnic University in Timisoara, Romania. In 2007 he received the MS degree at the Computer Science Department at Brown University where he currently pursues a PhD degree in Computer Science. His research interests are mainly in graphics and visualization with specific applications in genomic and proteomic biology.



Çağatay Demiralp Çağatay Demiralp is a PhD student in computer science at Brown University. His research interests are in characterizing the pattern and structure in data. Computational brain connectivity is the focus of his current research.



David H. Laidlaw David H. Laidlaw received the PhD degree in computer science from the California Institute of Technology, where he also did post-doctoral work in the Division of Biology. He is a professor in the Computer Science Department at Brown University. His research centers on applications of visualization, modeling, computer graphics, and computer science to other scientific disciplines. He is a senior member of the IEEE and the IEEE Computer Society.