Network Connectivity via Inference over Curvature-regularizing Line Graphs

Maxwell D. Collins[1,2], Vikas Singh[2,1], and Andrew L. Alexander[3]

[1] Department of Computer Sciences
[2] Department of Biostatistics and Medical Informatics
[3] Waisman Laboratory for Brain Imaging, Departments of Medical Physics and Psychiatry
University of Wisconsin-Madison, Madison, WI
mcollins@cs.wisc.edu, vsingh@biostat.wisc.edu,
alalexander2@wisc.edu

Abstract. Diffusion Tensor Imaging (DTI) provides estimates of local directional information regarding paths of white matter tracts in the human brain. An important problem in DTI is to infer tract connectivity (and networks) from given image data. We propose a method that infers high-level network structures and connectivity information from Diffusion Tensor images. Our algorithm extends principles from perceptual contours to construct a weighted line-graph based on how well the tensors agree with a set of proposal curves (regularized by length and curvature). The problem of extracting high-level anatomical connectivity is then posed as an optimization problem over this curvature-regularizing graph – which gives subgraphs which comprise a representation of the tracts' network topology. We present experimental results and an open-source implementation of the algorithm.

1 Introduction

Diffusion-tensor imaging (DT-MR or DTI) is an imaging modality that measures the diffusion of water molecules in brain tissues [1]. DTI exploits the fact that bundles of neural tissues with a certain orientation *preferentially restrict* water diffusivity (especially perpendicular to the direction of the fibers), which is otherwise isotropic in an unrestricted medium [1, 2]. The diffusion data is given as a 3×3 positive semidefinite matrix at each voxel [3, 4], and provides an estimate of the microstructural organization in the brain. DT-MR images are important to quantify how the neural fiber organization varies with cognitive change, age, and diseases [5], and therefore are very promising in the context of many neuroscience questions.

Research in DTI has extensively focused on the design of tools to facilitate the process of obtaining (from raw DTI data) connectivity maps of the entire human brain; in other words, the strength of connectivity in axonal brain networks. One approach toward deriving such information is to calculate, as a first step, the "network pathways" (or fiber tracts) between regions. This procedure is referred to as *tractography* [6–8]. It is reasonable to expect that if the underlying diffusion signal is ideal, a simple streamline propagation process (along orientation specific diffusion) will lead to the desired
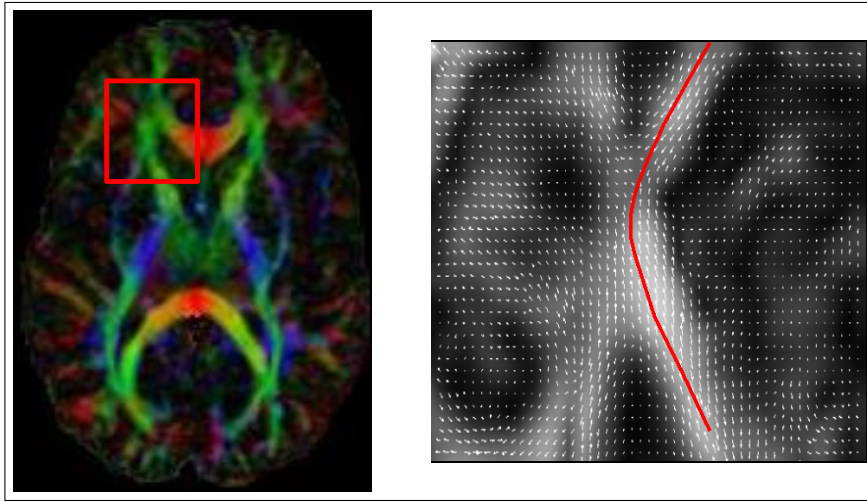
Fig. 1: A color-map of the orientation of White Matter in an image with a selected region shown in a red box (left); The vector field of the selected region and a possible pathway in red (right).

solution. That is, one sequentially follows the principal eigenvector of the diffusion tensor at each voxel to reconstruct the underlying tract [9, 10], see Fig. 1. Unfortunately, there is significant signal drop-off in areas where the diffusion is isotropic (common in regions containing "crossing" fibers). Noise in the estimation of the tensors (or in the acquisition itself) further exacerbates the problem of estimating the underlying pathway. While several local methods have been proposed for tractography, they occasionally make mistakes in the presence of noise and ambiguity (see [11] for a discussion). These errors accumulate as the tracking proceeds, and may also mislead the process into pursuing erroneous paths [12]. Further, the tracking may get lost when it passes through "uncertain" regions (i.e., where the magnitudes of the first and second eigenvectors are similar). Such limitations are common in most local tractography methods which is why recent work in this area suggests a preference toward strategies that lead to more *global* solutions [13]. While this idea is interesting and seems to be an appropriate solution to the problem, there is a significant associated cost. Many of the global methods proposed in the literature are very computationally demanding, and some take more than *one month* of processing time per image [13]. Our primary focus in this paper is to come up with efficient methods for this problem: *to infer reliable long range connections globally from (potentially erroneous or ambiguous) local orientation information.*

1.1   Related Work

Local methods track fibers through a series of small steps, where each step provides an estimate of the local fiber direction within a pre-specified neighborhood. The simplest tractography methods, known as *streamline tracking* [14–16], employ path integration based on the diffusion direction. Noise in the principal eigenvector (due to estimation

or acquisition), however, may lead to significant inaccuracies in this streamline process. To address this issue, Tensor Deflection (TEND) [11] takes into account the shape of the diffusion tensor as well in each update step (i.e., not relying on only the principal eigenvector). A number of subsequent papers have demonstrated that introducing some degree of randomness/stochasticity in the local estimation process helps in reducing errors, and yields better results in general [17].

Local methods described above have been extensively used in Neuroscience studies using DTI data. However, in an effort to identify more subtle group-level differences in statistical evaluations (which clearly require more accurate tractography solutions), there is a great deal of interest in leveraging more global methods for this problem. To this end, some authors have made use of Bayesian methods [18], while others have proposed the incorporation of strong priors on the tracking by assuming probability density models of the fiber directions [19, 20]. However, these strategies have certain limitations, especially in cases where the pre-specified model is far from the actual orientation distribution in the given image volume. A recently proposed technique for this problem [13], Gibbs Tracking, formulates the problem as an energy minimization, which is solved via a variant of Markov Chain Monte Carlo methods. The approach is quite interesting and seeks to consider the entire image at once in order to handle fiber crossings, while incorporating local agreement with the data and higher order connectedness properties. Unfortunately, such algorithms are known to require significant computational resources and turn out to be rather inefficient for practical applications.

The efficiency/quality trade-off described above is quite significant in many cases. This has led a number of authors [21–24] to investigate the utility of concepts from graph theory – that is, by constructing a weighted neighborhood graph over the voxels. For example, in [24], the orientation distribution function at a given voxel is evaluated for the vector in the direction of the outgoing edge, to obtain a weight which corresponds to the likelihood that a tract connects the two voxels. Then, a simple shortest-path algorithm on this graph finds the most probable tract between any pair of locations. Observe that the success of such a method varies with the richness of the underlying graph representation (e.g., whether it is curvature regularized or not, and how/whether the higher order dependencies are modeled). It also raises the question whether more powerful graph algorithms can lead to tangible improvements. We seek to address both these issues in this paper.

A recent work that is relevant to ours is the method in [25]. Here, the authors model tracts as helices between *triplets* of tensors. They introduce *co-helicity* to model the setting where the orientations at each point in the triplet can be joined by a helix. Then, a local search method is used to find orientations for each tensor that leads to the "most probable" helices that match the data. The choice of a helix representation here seems to be a distinct weakness. While our algorithm is similar to [25] in spirit, we present a different geometrically based scheme for generation of primitives (which offers some distinct advantages over [25]). Our most probable set of primitives is then optimized *globally* rather than via local search.

The **main contributions** of this paper are: **(i)** We propose construction of line graph primitives as basic building blocks of a tractography solution. Our strategy adapts ideas from Perceptual Contours and Tensor Voting for the generation of a set of "proposal

splines" to geometrically describe the local context in an image volume. **(ii)** We propose a simple optimization model (equipped with connectivity and branching constraints) to select a subset of edges from a curvature regularized graph to infer the final solution, from the given data. **(iii)** We present results on a set of Diffusion Tensor images noting that the proposed solution is of independent interest in the context of general Vision problems dealing with connectivity inference (for example, see [26, 27]).

2   Preliminaries

Certain ideas from a body of Computer Vision literature (called *perceptual contours*) can inform approaches to DTI connectivity and tractography; we briefly review the relevant details, and then move to the presentation of our formalization.

2.1   Perceptual Contours/Perceptual Grouping

The use of energy functions that prefer connections or curves minimal length and curvature has a long history in the perceptual contour literature [28]. Recent work has exploited such functions in a graph setting (with different types of regularization) for a variety of problems and applications [29–31]. For example, [31] presents a linear program over a graph of region and boundary segments to find a segmentation of the image – this minimizes an objective which includes curvature regularity on the boundary. *Tensor Voting* [32] is a type of perceptual grouping which takes as input a tensor field[1]. Each tensor field provides a distribution over *normals* rather than *tangents*. The tensor voting algorithm consists of iterations where each member of a tensor field casts a *vote* to its neighbors. The vote itself consists of a tensor based on the orientations of "proposal curves". In the case of the *stick vote*, which is the vote cast by a purely anisotropic stick tensor, we simply consider a proposal curve between the voting tensor and its neighbor with *known* orientation at all points. A stick vote is thus cast to the neighboring tensor in the same direction as the proposal curve's orientation at the neighbor. The magnitude of the stick vote is scaled by a curvature regularity function termed the *saliency decay function*:

$$\exp\left(\frac{s^2 + c\kappa^2}{\sigma}\right),\tag{1}$$

where s and κ are the length and curvature of the proposal arc, and c and σ are parameters expressing a neighborhood size. In the presence of uncertainty, the vote consists of an integration over multiple proposal curves. The magnitude of the resulting vote is then a function of how well the two tensors can be linked by simple and probable proposal curves. We make use of this idea in our modeling next.

---

[1] The term "tensor" here is used in a more general sense, but is similar in terms of interpretation to the usage in the context of DTI.
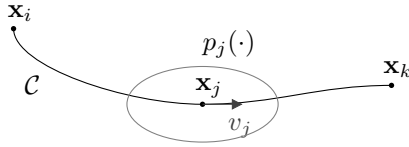
Fig. 2: Information in a triplet as described in Section 3.1. Probability distributions p_i, p_k and unit tangents v_i, v_k not shown.

3 Weighted Line Graph

Our strategy is to encode curvature regularity in the context of weights on a certain line graph over the tensor field. We will first outline a graph construction, and then discuss how a proposal set of curves can be calculated. Later, in Section 4, we will describe the optimization model to obtain our final solution.

3.1 Graph Construction

Consider a set of voxels $\mathcal{V} = \{\mathbf{x}_1, ..., \mathbf{x}_n\}$. Over these voxels an orientation distribution field is defined, so at $\mathbf{x}_i$ we have a distribution function $p_i(\hat{\mathbf{v}})$ corresponding to the local probability that a fiber *through* $\mathbf{x}_i$ has *tangent* along $\hat{\mathbf{v}}$ in that voxel. Adapting DTI data to such a distribution can be done by using $p_i(\hat{\mathbf{v}}) \propto \hat{\mathbf{v}}^T D_i \hat{\mathbf{v}}$, where D_i is the symmetric PSD matrix representing the tensor at voxel i [33]. Note that the formulation easily generalizes to other diffusion imaging modalities [34, 35].

We now define a neighborhood graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ over voxels in the input image. For simplicity, we can assume that $\mathcal{G}$ is the complete graph $\mathcal{K}_n$, though one may also use a threshold to limit the number of edges introduced (i.e. $\mathcal{E} = \{(ij) \mid \|\mathbf{x}_i - \mathbf{x}_j\| < \delta\}$ for some δ). Let $(\mathcal{E}, \mathcal{L})$ denote the *line graph* of $\mathcal{G}$, the edge set $\mathcal{L}$ is the set of *pairs* of edges of $\mathcal{G}$ which share a common endpoint. It is convenient to view an edge $((ij), (jk)) \in \mathcal{L}$ as a triplet $(ijk) \in \mathcal{L}$ of vertices such that $(ij) \in \mathcal{E}$ and $(jk) \in \mathcal{E}$, which elides the repeat of j. With the above notation in place, we can represent groups of white matter tracts at a global level as subgraphs of $\mathcal{L}$. If $\mathcal{H} \subset \mathcal{L}$ is a tractography and $(ijk) \in \mathcal{H}$, this expresses the belief that there is a tract segment $\mathcal{C}$ which connects the voxels $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$ (in order) and passes through *no other voxels* in between, see Fig. 2 for an illustration. Using the line graph enables us to explicitly model the topology of the tracts, as discussed in [36] and used in Section 4.1.

3.2 Minimum Energy Proposal Curves

To equip the graph, $\mathcal{G}$, with edge weights we use an energy function over curves $\mathcal{C}$ which consists of a weighted sum of total curvature and squared speed [37, 28].

$$\mathrm{E}(\mathcal{C}) = \int_{-1}^{1} K \cdot \kappa_{\mathcal{C}}(t)^2 + \|\mathcal{C}'(t)\|^2 dt, \tag{2}$$
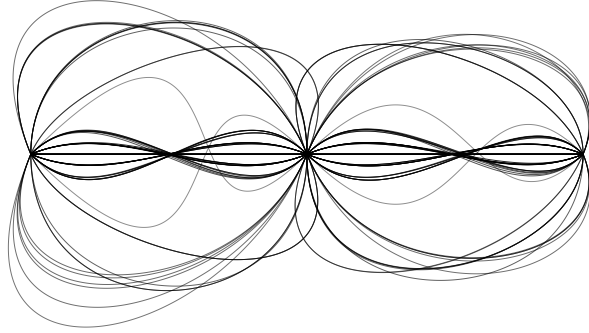
Fig. 3: Plot of the proposal spline function (3) when in two dimensions. We take a set of splines constructed and weighted as in section 3.2 passing through $(-1, 0), (0, 0), (1, 0)$, with orientation constraints ranging over a discretization of $\left[0, \frac{\pi}{2}\right]$. Lower-energy curves are shaded darker.

where K is a user-defined weighting constant which determines a neighborhood size. This energy function serves as a regularizer to prefer shorter and smooth curves.

Given known orientations at each point in the triplet $\hat{v}_i, \hat{v}_j, \hat{v}_k$, we can propose a most likely curve. This is equivalent to the proposal arc in tensor voting's stick vote [32] and is the basic building block in our weight construction. Briefly, given a family of curves parameterized over $[-1, 1]$, we can generate a proposal curve by choosing $\mathcal{C}$ according to:

$$
\begin{aligned}
\underset{\mathcal{C}}{\text{argmin}} \quad & E(\mathcal{C}) \\
\text{subject to} \quad & \mathcal{C}(-1) = \mathbf{x}_i, \mathcal{C}(0) = \mathbf{x}_j, \mathcal{C}(1) = \mathbf{x}_k \\
& \frac{\mathcal{C}'(-1)}{\|\mathcal{C}'(-1)\|} = \hat{v}_i, \quad \frac{\mathcal{C}'(0)}{\|\mathcal{C}'(0)\|} = \hat{v}_j, \quad \frac{\mathcal{C}'(1)}{\|\mathcal{C}'(1)\|} = \hat{v}_k,
\end{aligned}
\tag{3}
$$

where $E(\cdot)$ denotes the energy. The proposal curves used here are cubic hermite splines with two segments. The knots have positions $\mathbf{x}_i, \mathbf{x}_j, \mathbf{x}_k$ and derivatives $m_i \hat{v}_i, m_j \hat{v}_j, m_k \hat{v}_k$ respectively. Note that we must distinguish between an orientation and a tangent. The tensors provide a distribution over orientations which we express as a unit vector $\hat{v}.$, parallel (or antiparallel) to the tangent of any proposal curve. The tensors provide no information on the magnitudes $m.$ of the tangents, so they are chosen to minimize the curve energy in (2) through gradient descent. We calculate the gradient for $m.$ by approximating the integral (the longer version of the paper includes all relevant details):

$$
\frac{\partial}{\partial m.}(E(\mathcal{C})) = \int_{-1}^{1} \frac{\partial}{\partial m.} \left(K \kappa_{\mathcal{C}}(t)^2 + \|\mathcal{C}'(t)\|^2 \right) dt
\tag{4}
$$

Thus, given the positions and orientations at each point we can find a proposal spline and its corresponding energy. An illustration is provided in Fig. 3.
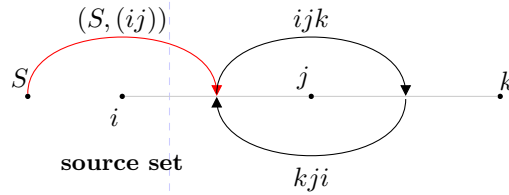
Fig. 4: Elements of the flow graph used in (7). Displayed is a source set of $\{i\}$ and a neighborhood graph G with the edes (ij), (jk), inducing the directed line graph edges (ijk), (kji). An edge $(S, (ij))$ is added to introduce flow from the source to (ij).

3.3   Expected Curve Energy

The local information from the tensors in a triplet and (4) can now be used to calculate weight for the corresponding line edge. If we call a solution to (3) $\text{Curve}(\mathbf{x}_{\{i,j,k\}}, \hat{v}_{\{i,j,k\}})$, then by taking the expectation

$$w_{ijk} = \underset{\hat{v}_{\{i,j,k\}} \sim p_{\{i,j,k\}}}{\mathbb{E}} \left[E(\text{Curve}(\mathbf{x}_{\{i,j,k\}}, \hat{v}_{\{i,j,k\}})) \right] \tag{5}$$

This yields weights for our line edges, and are analogous to edge weights calculated in other graph-based methods in tractography [21, 23, 24] and segmentation [30, 31]. Observe that it is possible to quickly calculate the weights for a large connectivity graph over a regular grid (via a preprocessing step) by taking advantage of the fact that the energy function is invariant to rigid body transformations of the corresponding curve.

4   Inferring Connectivity

We can find the most probable (least-weight) tracts connecting an arbitrary pair of regions by solving an *augmented* min-cost flow problem over the digraph $(\mathcal{V}_f, \mathcal{E}_f)$. $\mathcal{V}_f = \mathcal{E} \cup \{S, T\}$, where S and T are nodes for the source and sink respectively. $\mathcal{E}_f$ is the union of the symmetric digraph equivalent to the line graph $\mathcal{L}_{\pm}$, and a set of edges $(S, (ij))$ or $((ij), T)$ for all $(ij) \in \mathcal{E}$ incident to voxels in the source or sink set respectively, see Figure 4. We set d_v for $v \in \mathcal{V}_f$ to be the flow divergences:

$$d_e = \begin{cases} N & \text{if } e = S \\ -N & \text{if } e = T \\ 0 & \text{if } e = (ij) \in \mathcal{E} \end{cases} \tag{6}$$

where we wish to recover N most-likely tracts.

We use α_{ijk} as indicator variables to give the presence of flow across voxel j moving from i to k. A penalty for branching or crossing is imposed by adding variables β_j for each voxel $j = \{1, \cdots, n\}$. Each β_j is the total amount of flow passing across that voxel. The model imposes a (user-specified) penalty of λ for each unit of flow over 1. The effect is a hinge loss which produces topologically simple tractographies with large groups of parallel tracts, except where the data strongly suggests that the
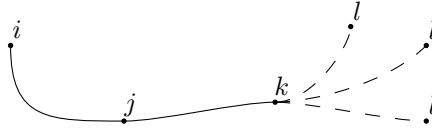
Fig. 5: Illustration of continuation constraints described in Section 4.1.

reduction in the weight term introduced by a crossing (or branching) is greater than λ. This hyperparameter encodes a prior based on the topology of the tracts, and can be tuned by a qualitative analysis on how much branching is expected in the true tracts.

$$
\begin{aligned}
\min_{\alpha,\beta} \quad & \sum_{(ijk)\in\mathcal{L}_\pm} w_{ijk}\alpha_{ijk} + \lambda \sum_{j\in\mathcal{V}} \beta_j & (7)\\
\text{subject to} \quad & \sum_{(kij)\in\mathcal{L}_\pm} \alpha_{kij} - \sum_{(ijk)\in\mathcal{L}_\pm} \alpha_{ijk} = d_{(ij)} \quad \forall(ij)\in\mathcal{E}_f \\
& \beta_j \geq \sum_{(ijk)\in\mathcal{L}_\pm} \alpha_{ijk} \quad \forall j \in \mathcal{V} \\
& \beta_j \geq 1 \quad \forall j \\
& \alpha_{ijk} \in \{0,1\} \quad \forall(ijk)\in\mathcal{L}_\pm
\end{aligned}
$$

We can solve the above IP using a number of solvers; from the solution the resulting tractography is interpreted as $\mathcal{H} = \{(ijk) \in \mathcal{L} \mid \alpha_{ijk}^* + \alpha_{kji}^* \geq 1\}$ where α^* is the solution to (7). Note that the minimum-cost flow problem is a special case of (7) for $\lambda = 0$, and can be solved exactly with the LP relaxation $\alpha_{ijk} \in [0,1]$.

4.1   Continuation Constraints

We can derive a new model by modifying (7), relaxing the constraints on the endpoints of the tracts. In such a relaxation, the flow constraints will be replaced with *continuation* constraints similar to [31], see Fig. 5. The continuation constraints express the dichotomy that for a given edge pair and direction, there is another edge pair the continues that tract, otherwise selecting this edge pair introduces at least one "endpoint" to the tractography. We use variables γ_{ijk} as indicator variables equal to 1 in such a situation (i.e., for an endpoint), and 0 otherwise. Each γ incurs a penalty of μ, optionally relaxed at a set of voxels M considered likely endpoints (i.e. the GM-WM boundary). This hyperparameter allows for "soft" endpoint constraints, so that M can be specified approximately. Note that as $\mu \to +\infty$, the optimal $\gamma^* = \mathbf{0}$ for all $(ijk) \in \mathcal{L}$, and all tracts found will either be cycles or end within M. The reverse of this is that if connecting a tract will incur a penalty in the other remaining terms of less than μ, the corresponding connection *will* be made in the optimal subgraph. Our model is given as

$$\min_{\alpha,\beta,\gamma} \quad -\sum_{(ijk)\in\mathcal{L}} w_{ijk}\alpha_{ijk} + \lambda\sum_{j\in\mathcal{V}}\beta_j^s + \mu\sum_{(ijk)\in\mathcal{L}\setminus\mathcal{L}(M)}\gamma_{ijk}$$

$$\text{subject to} \quad \alpha_{ijk} \in \{0,1\}$$

$$\beta_j \geq \sum_{(ijk)\in\mathcal{L}}\alpha_{ijk}$$

$$\gamma_{ijk} \geq 0$$

$$\gamma_{ijk} \geq \alpha_{ijk} - \sum_l \alpha_{lij}$$

$$\gamma_{ijk} \geq \alpha_{ijk} - \sum_l \alpha_{jkl},$$

(8)

where $\mathcal{L}(M) = \{(ijk) \in \mathcal{L} \mid i \text{ or } k \in M\}$, and $s \in \{1,2\}$ is a hyperparameter governing the kind of sparsity penalty. For instance, $s = 2$ penalizes high levels of branching and crossing. Further, the binary constraint on $\alpha_{ijk} \in \{0,1\}$ can be relaxed to $\alpha_{ijk} \in [0,1]$ (in which case we obtain a LP for $s = 1$ and QP for $s = 2$).

5   Experiments

Our experiments were designed to evaluate where the proposed model can reliably (and efficiently) recover tract connections among different brain regions by incorporating local geometric context within a global optimization model. To this end, we first compared results from our algorithm relative to other streamline-based tracking methods on a number of synthetic tensor datasets. These experiments are useful to answer if the method can correctly resolve crossing fibers, as well as its applicability in general curve inference problems. We also evaluated our results on a set of DT-MR images. Since obtaining ground-truth data on such images is clearly impractical, our evaluations were mainly qualitative – by focusing on pairs of some important regions (e.g., Corpus Callosum) we can reliably assess consistency between our solution and known organizations of tracts in those regions. We present our experimental results next.

5.1   Simulated Data

Synthetic tensor fields were constructed from manually specified linear paths. At each voxel along the line, we add the stick tensor in the direction of the ground truth line and a random tensor of magnitude up to a given SNR (10:1 in our experiments). At voxels containing a crossing we take the average of all of these tensors for the crossing tracts. Finally, areas not occupied by a tract are filled with random noise.

In Fig. 6, we illustrate the necessity of a global approach to tractography by showing an example where two fiber cross at an angle. A source set was placed at the lower-left tract endpoints and a sink set was placed at the upper-right endpoints, the union of these two sets was used to seed TEND [11]. Notice that due to partial voluming, the principal diffusion direction in the region of the crossing is between the two tracts. Local methods will typically infer tracts which follow this average direction and infer higher-curvature tracts or leave the region occupied by the tracts entirely. This is compared relative to

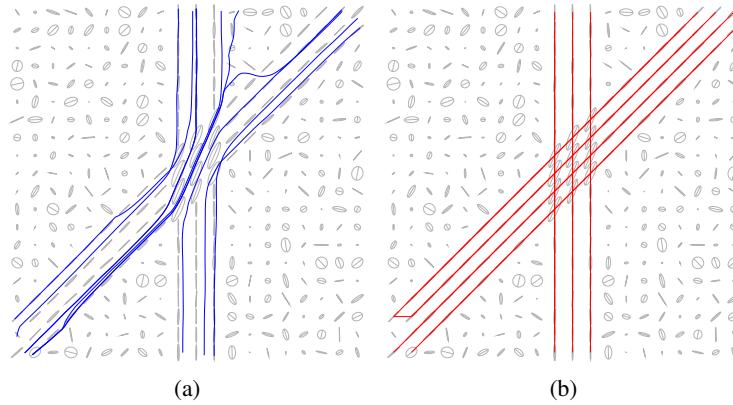(a)   (b)

Fig. 6: Comparison of local tractography (a) with ROI pair method from section 4 (b) on a basic problem with a crossing.
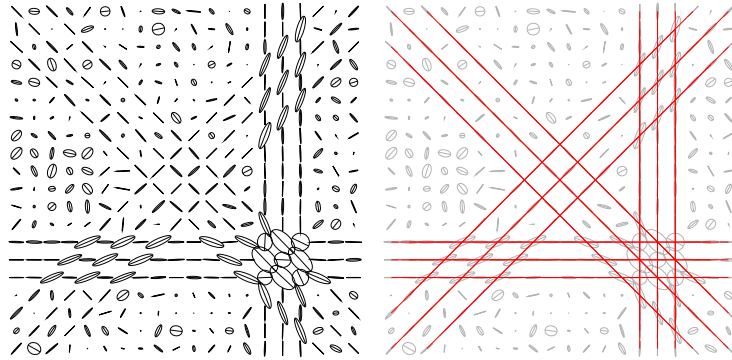

Fig. 7: Demonstration of (8).

our global method which extracts the minimum-curvature tracts that match the ground truth rather well.

We also demonstrate the usefulness of the model in (8) on such a synthetic dataset in Fig. 7. With M set to the tensors at the edges of the image, we infer the tracts with no user intervention beyond specifying the appropriate hyperparameters for these artificial tensor fields.

5.2   Anatomical Structures

Acquisition Setup. DT-MR images used in our evaluations were acquired with a GE SIGNA 3-Tesla scanner. DW images were taken with 12 non-collinear diffusion directions and a diffusion weighting factor of $b = 1000\text{s/mm}^2$. Eddy current related distortion and head motion of each data set were corrected, using standard methods. Distortions from field inhomogeneities were corrected using field maps. From the raw

data, tensor elements were estimated using methods available in Camino, and the data was registered to a common template [38]. The resulting DTI images were then resampled to $128 \times 128 \times 64$ voxels, each of size 1.5 mm $\times 1.75$ mm $\times 2.25$ mm. White matter was segmented using the FAST tool available as part of FSL [39]. The tract networks inferred by our algorithm were restricted to lie within this mask.

We constructed a 6-neighborhood graph over the white matter region, and the graph weights were calculated using (5). Source and sink sets were specified for two prominent tract groups: the callosal and projective fibers [40]. We set the minimum flows to recover a total of 200 pairs of tract endpoints, with 120 in the callosal fibers. Tracts were extracted from the optimal line graph by finding the endpoint-to-endpoint paths between voxels within the subgraph, using ideas discussed in Section 4.1. We find the edge pair leading up to an endpoint as defined there, then repeatedly consider those edge pairs which share the next edge. This leads to paths among the voxels, and B-splines are then fit to the ordered sequence of their centers to obtain the final tracts. Representative results of our method are presented in Figure 8. In general, our results are consistent with known/expected connection pathways in these regions, as well as results obtained via other methods. Some local artifacts are seen due to the low angular resolution of a 6-neighborhood graph over the voxels. This can be addressed by considering a more computationally-intensive higher-degree graph or by increasing the smoothing when fitting B-spline streamlines during postprocessing.

We note, however, that the proposed method is global and is unaffected by voxels with crossing fibers (which occasionally leads to inaccuracies and errors in other methods). In addition, while streamlines serve as a visualization of inferred tracks, our core algorithm outputs a line graph from which a wide range of measures can be calculated. This includes overall connectivity between ROIs for use within group studies and identifying points of crossing and branches.

6   Conclusions

In this paper, we have presented an algorithm for inferring tract connectivity information from DT-MR images. Our algorithm constructs a line graph, whose edge weights are calculated based on a set of proposal splines. This helps equip our algorithm with local geometric context. Once such primitives are generated, a global optimization procedure gives the final tractography solution. Our global model is inspired by network flow algorithms but includes additional constraints that penalize extensive branching and encourage strong region-to-region connectivity. Such requirements are imposed to ensure that the resultant solution is consistent with the topology of white matter tract networks in the brain. We have presented experimental results on synthetic data as well as on brain images, where the proposed method performs well. While extensive further evaluations on a variety of datasets are still required to assess the advantages and limitations of this method, and we believe that the proposed model provides a framework for incorporation of local *and* global context for tractography. Our C++ implementation will be made publicly available concurrently with publication.
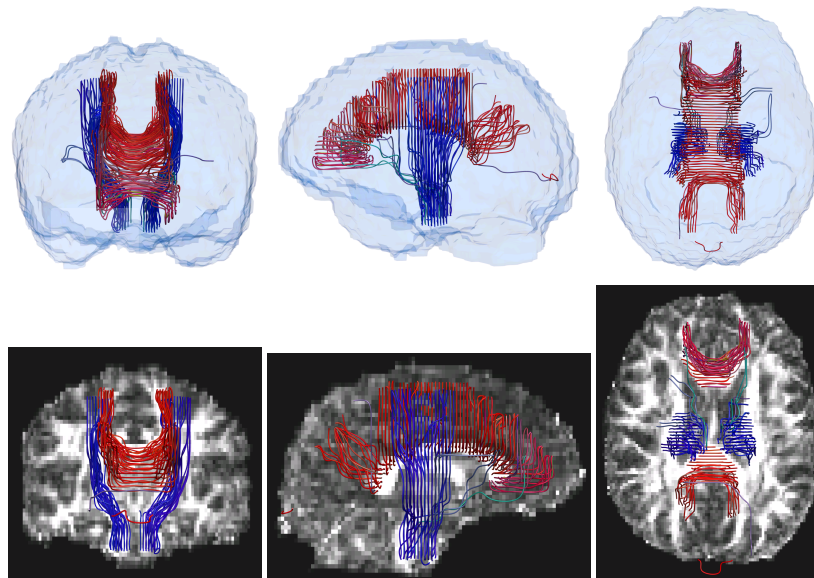
Fig. 8: A visualization of our tractography solution in six different views using Trackvis.
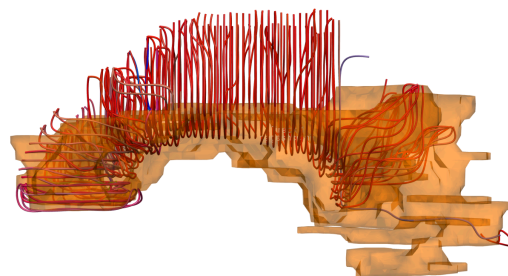

Fig. 9: Callosal fibers overlaid on automatic segmentation of the corpus callosum.
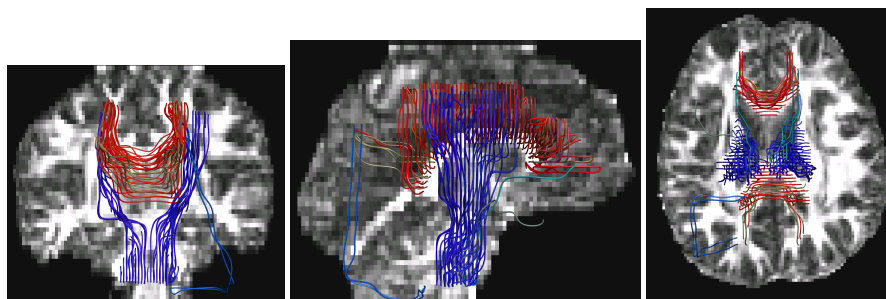

Fig. 10: Fibers from another subject using the same source and sink sets.

Acknowledgements

References

1. Bihan, D.L.: Molecular diffusion nuclear magnetic resonance imaging. Magnetic Resonance Quarterly **7** (1991) 1–30
2. Bammer, R.: Basic principles of diffusion-weighted imaging. European J. of Radiology **45** (2003) 169 – 184
3. Bihan, D.L., Mangin, J., Poupon, C., et al.: Diffusion tensor imaging: Concepts and applications. J. of Magnetic Resonance Imaging **13** (2001) 534–546
4. Basser, P.J., Jones, D.K.: Diffusion-tensor MRI: theory, experimental design and data analysis - a technical review. NMR Biomed **15** (2002) 456–467
5. Naggara, O., Oppenheim, C., Rieu, D., et al.: Diffusion tensor imaging in early Alzheimer's disease. Psychiatry Research: Neuroimaging **146** (2006) 243 – 249
6. Cook, P., Bai, Y., Nedjati-Gilani, S., Seunarine, K., Hall, M., Parker, G., Alexander, D.: Camino: Open-source diffusion-MRI reconstruction and processing. In: ISMRM, Seattle, WA, USA (2006) 2759
7. Basser, P.J., Mattiello, J., Bihan, D.L.: MR diffusion tensor spectroscopy and imaging. Biophys. J. **66** (1994) 259–267
8. Stieltjes, B., Kaufmann, W.E., Zijl, P.V., et al.: Diffusion tensor imaging and axonal tracking in the human brainstem. NeuroImage **14** (2001) 723 – 735
9. Bammer, R., Acar, B., Moseley, M.E.: In vivo MR tractography using diffusion imaging. European J. of Radiology **45** (2003) 223 – 234
10. Roebroeck, A., Galuske, R., Formisano, E., et al.: High-resolution diffusion tensor imaging and tractography of the human optic chiasm at 9.4t. NeuroImage **39** (2008) 157 – 168
11. Lazar, M., Weinstein, D.M., Tsuruda, J.S., et al.: White matter tractography using diffusion tensor deflection. Human Brain Mapping **18** (2003) 306 – 321
12. Basser, P.J., Pajevic, S., Pierpaoli, C., et al.: In vivo fiber tractography using dt-mri data. Magnetic Resonance in Medicine **44** (2000) 625–632
13. Kreher, B.W., Mader, I., Kiselev, V.G.: Gibbs Tracking: A novel approach for the reconstruction of neuronal pathways. Magnetic Resonance in Medicine **60** (2008) 953 – 963
14. Basser, P.J., Pajevic, S., Pierpaoli, C., et al.: In vivo fiber tractography using DT-MRI data. Magnetic Resonance in Medicine **44** (2000) 625 – 632
15. Conturo, T.E., Lori, N.F., Cull, T.S., et al.: Tracking neuronal fiber pathways in the living human brain. Proc. of the National Academy of Sciences **96** (1999)
16. Mori, S., Crain, B.J., Chacko, V.P., et al.: Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Annals of Neurology **45** (1999) 265 – 269
17. Bjornemo, M., Brun, A., Kikinis, R., et al.: Regularized stochastic white matter tractography using diffusion tensor MRI. MICCAI (2002) 435–442
18. Jbabdi, S., Woolrich, M.W., Andersson, J.L.R., et al.: A Bayesian framework for global tractography. Neuroimage **37** (2007) 116–129
19. Wedeen, V.J., Hagmann, P., Tseng, W.I., et al.: Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magnetic Resonance in Medicine **54** (2005) 1377–1386

20. Schmahmann, J.D., Pandya, D.N., Wang, R., et al.: Association fibre pathways of the brain: parallel observations from diffusion spectrum imaging and autoradiography. Brain **130** (2007) 630–653
21. Iturria-Medina, Y., Sotero, R.C., Canales-Rodrguez, E.J., et al.: Studying the human brain anatomical network via diffusion-weighted MRI and graph theory. NeuroImage **40** (2008) 1064 – 1076
22. Lifshits, S., Tamir, A., Assaf, Y.: Combinatorial fiber-tracking of the human brain. NeuroImage **48** (2009) 532–540
23. Sotiropoulos, S.N., Bai, L., Morgan, P.S., et al.: Brain tractography using Q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach. NeuroImage **49** (2010) 2444 – 2456
24. Zalesky, A.: DT-MRI fiber tracking: A shortest paths approach. IEEE Trans. Medical Imaging **27** (2008) 1458–1471
25. Savadjiev, P., Campbell, J., Pike, G., et al.: 3D curve inference for diffusion mri regularization and fibre tractography. Medical Image Analysis **10** (2006) 799 – 813
26. Peng, T., Jermyn, I.H., Prinet, V., et al.: Incorporating generic and specific prior knowledge in a multi-scale phase field model for road extraction from VHR images. IEEE Trans. Geoscience and Remote Sensing **1** (2008) 139–146
27. Rochery, M., Jermyn, I.H., Zerubia, J.: Higher-order active contour energies for gap closure. J. of Mathematical Imaging and Vision **29** (2007) 1–20
28. Kass, M., Witkin, A., Terzopoulos, D.: Snakes: Active contour models. International J. of Computer Vision **1** (1988) 321–331
29. Boykov, Y., Kolmogorov, V.: Computing geodesics and minimal surfaces via graph cuts. In: ICCV. (2003) 26–33 vol.1
30. Schoenemann, T., Cremers, D.: Introducing curvature into globally optimal image segmentation: Minimum ratio cycles on product graphs. In: ICCV. (2007)
31. Schoenemann, T., Kahl, F., Cremers, D.: Curvature regularity for region-based image segmentation and inpainting: A linear programming relaxation. In: ICCV. (2009)
32. Mordohai, P., Medioni, G.: Tensor Voting: A Perceptual Organization Approach to Computer Vision and Machine Learning. Morgan & Claypool Publishers (2006)
33. Parker, G., Haroon, H., Wheeler-Kingshott, C.: A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. J. of Magnetic Resonance Imaging **18** (2003) 242 – 254
34. Barmpoutis, A., Hwang, M., Howland, D., et al.: Regularized positive-definite fourth order tensor field estimation from DW-MRI. NeuroImage **45** (2009) S153 – S162 Mathematics in Brain Imaging.
35. Tuch, D.S.: Q-ball imaging. Magnetic Resonance in Medicine **52** (2004) 1358 – 1372
36. Savadjiev, P., Campbell, J.S., Descoteaux, M., Deriche, R., Pike, G.B., Siddiqi, K.: Labeling of ambiguous subvoxel fibre bundle configurations in high angular resolution diffusion MRI. NeuroImage **41** (2008) 58–68
37. Guy, G., Medioni, G.: Inferring global perceptual contours from local features. International J. of Computer Vision **20** (1996) 113–133
38. Zhang, H., Yushkevich, P., Alexander, D., et al.: Deformable registration of diffusion tensor MR images with explicit orientation optimization. Medical Image Analysis **10** (2006) 764–785
39. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain MR images through a hidden markov random field model and the expectation-maximization algorithm. IEEE Trans. Medical Imaging **20** (2001) 45 –57
40. Wakana, S., Jiang, H., Nagae-Poetscher, L.M., et al.: Fiber tract-based atlas of human white matter anatomy. Radiology **230** (2004) 77–87