A Minimum Cost Approach to Connectivity from Orientation Distribution Functions via Efficient Multi-directional Graph Propagation

Ipek Oguz, Alexis Boucharin, Wenyu Lu, Clement Vachet, Francois Budin, Yundi Shi, and Martin Styner

Depts. of Psychiatry and Computer Science, University of North Carolina

Abstract. Regional connectivity measurements derived from diffusion imaging datasets are of considerable interest in the neuroimaging community for better understanding white matter connectivity. Current connectivity measurements are usually either based on fiber tractography applied in Monte-Carlo fashion, or variations of the Hamilton-Jacobi approach, or are graph-based. We propose a novel, graph-based algorithm that provides a fully deterministic, efficient and stable connectivity measure. This method handles crossing fibers and deals well with multiple seed regions. The computation is based on a multi-directional graph propagation algorithm applied to sampled orientation distribution functions computed directly from the original diffusion imaging data. While a maximum probability approach is possible, here we focus on a minimum cost formulation. We present results on synthetic and real datasets to illustrate the potential of our method towards subject-specific connectivity measurements performed in an efficient, stable and reproducible manner. Such individual connectivity measurements would be well suited for neuroimaging studies.

Keywords: DTI, MRI, ODF, connectivity, graph-based, F*

1 Introduction

In recent years, the use of regional connectivity measurements computed from diffusion MRI data have become of considerable interest in the neuroimaging community in an effort to better understand cortical and subcortical white matter connectivity. While measuring cortical and subcortical white matter connectivity is an important endpoint in and of itself, efficient and stable methods for quantifying connectivity is particularly important for network analysis studies. A connectivity matrix designating the connection strength between all brain regions is typically the starting point of such studies.

The most elementary measurement of connectivity strength is the mean Fractional Anisotropy (FA) along the path between two regions. While this is a straightforward metric to compute, it is far from ideal as it cannot effectively handle crossing fibers and also because FA has a non-uniform and non-linear distribution along the fiber, which simple averaging completely disregards.

Stochastic tractography methods overcome some of these issues by repeatedly applying streamline tractography in a Monte-Carlo fashion[1, 5]. The connectivity strength at a given voxel is then defined as the number of paths reaching that voxel divided by the total number of generated paths. While this powerful method can overcome the problem of crossing fibers with an appropriate local diffusion model, it is not deterministic and can be very inefficient due to the large number of tracts that must be generated to converge to a stable result.

Various methods based on the Hamilton-Jacobi approach have been proposed to overcome some of the difficulties arising in tractography [7, 12, 10]. The main idea for these methods is to compute the shortest path where the cost associated with each path is an integral dependent on position, path orientation and local diffusion anisotropy/strength. These formulations result in first- or higher-order partial differential equations which model evolving fronts whose speeds are determined by information from the diffusion tensor [7]. These methods are inherently more robust to noise than tractography methods and also have the advantage of being fully deterministic and computationally efficient. However, these methods cannot handle crossing fibers effectively, whereas this can be incorporated into explicit tractography models. Moreover, these methods can not take into account the consistency of path orientation along the minimal path.

A final class of methods for computing connectivity consists of graph-based approaches [8, 15, 13]. These approaches treat the image as a graph by placing a vertex at each voxel and edges between neighboring voxels. The edges have weights, which can be considered as costs or probabilities of their two end voxels being connected to each other. These methods can take advantage of the vast graph processing literature for efficient computations. However, these methods are typically prone to taking "shortcuts" in the graph (or creating "phantom bridges", as Zalesky et al. calls them), i.e. preferring a shorter path that doesn't fit the data over a longer path that fits the data.

We propose a novel multi-directional graph-propagation based algorithm that computes connectivity between brain regions in a fully deterministic and efficient way, much like the Hamilton-Jacobi approach, while allowing crossing fibers. Furthermore, our method respects the local connectivity patterns in the data, which is not the case for the global-level optimization of Hamilton-Jacobi methods. This further allows our method to avoid the phantom bridge problem.

2 Methods

Motivated by the well-known $F*$ (pronounced "f-star") graph traversal algorithm [6] and previous work using $F*$ to compute non-linear distances within the brain [14], we extended $F*$ to incorporate multiple incoming and outgoing directions for computing the overall probability or cost of a voxel. While we designed two formulations of this algorithm, one for propagating costs and one for propagating probabilities, here we will focus purely on the former. For details on the probability formulation, please refer to [2]. In addition to being more computationally expensive (probabilities are multiplicative whereas cost is ad-

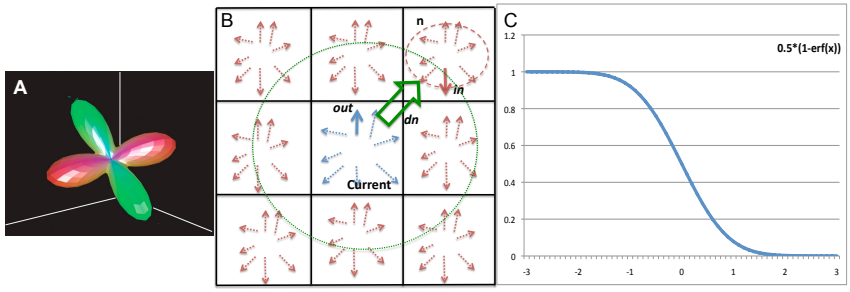

Fig. 1. An ODF representation at a single voxel (A). Multi-directional cost function diagram (B) shows the variables that affect the cost of traversing an edge in 2D. For the voxel *Current*, for each direction ***out***, we consider all neighboring voxels (green) and all the possible directions (red) at these voxels in order to identify the neighbor n, which lies in the direction $\boldsymbol{d_n}$ from the *Current* voxel, and the direction ***in*** at the neighbor n, such that the cost is minimal. (C) Angle penalty function.

ditive), one of the major issues with the probabilistic approach is its proneness to cyclic connections, whereby a small loop of strongly connected voxels can be revisited many times rather than leaving the loop to follow the fiber path. The cost approach, as discussed below, allows for a length-independent formulation, which makes such loops unadvantageous and therefore avoids this problem.

2.1 Scalar F* graph traversal

The basic F* algorithm [6] uses an adapted TV-scan processing of the image: starting from one corner of the three dimensional input image, the algorithm will visit each voxel in consecutive order. The processing will first visit each voxel of the current line until its end, then it will visit each voxel of the same line again in reverse order. This is performed for each line until the end of the current slice. At the end of the current slice, it will traverse that slice in reverse order. This interweaving of forward and backward iterations greatly reduces the overall computation time: during backward iterations, only the neighbors visited in the preceding forward iteration need to be considered; similarly, during forward iterations, only the neighbors visited in the previous backward iteration need to be considered. In the experiments presented in the Results section, the number of iterations were typically in the range of 5 to 10 before convergence.

The F* algorithm keeps track of the current optimal cost for each voxel as well as the length and the origin node of the path associated with this optimal cost. Note that this basic F* algorithm is only suitable for graph traversal problems where the cost of traversing an edge is a single scalar.

2.2 Local diffusion model

We use orientation distribution functions (ODF) for our local diffusion model. The ODF image can be computed directly from a diffusion-weighted image

(DWI) using the technique presented by Descoteaux et al. [4]. Briefly, this method provides an analytical solution by modeling the diffusion imaging signal with a spherical harmonic basis that incorporates a regularization term based on the Laplace-Beltrami operator. The ODF image thus provides a continuous diffusion distribution function at each voxel of a 3D image (Fig. 1A). This ODF image is sampled at each voxel into a spherical sampled distribution using a standard electrostatic repulsion scheme[9], which allows for a high number of directions to achieve the desired level of accuracy.

2.3 Multi-directional F* graph traversal

In order to accommodate the multi-directional local diffusion model, we extend the basic scalar F* algorithm to keep track of multiple directions at each voxel (represented by a graph vertex). Each voxel visit consists of updating the cost for each sampled ODF direction of the current voxel based on the costs of its neighbors (graph vertices connected by an edge to the current vertex). To this end, for each sampled ODF direction *out*, for each neighbor n, the cost of arriving to n in direction *in* and traversing an additional edge along *out* is computed. The minimum cost thus obtained is:

$$C_{current}^{\boldsymbol{out}} = \min_{\substack{n=neighbor \\ \boldsymbol{in}=direction}} \left[C_n^{\boldsymbol{in}} + cost(\boldsymbol{in}, \boldsymbol{out}, \boldsymbol{d_n}, ODF_{current}, ODF_n) \right] \quad (1)$$

where $\boldsymbol{d_n}$ represents the vector from the voxel-center of *current* to the voxel-center of n. The $ODF_{voxel}^{\boldsymbol{dir}}$ terms refer to the value of the input orientation distribution function sampled at direction $\boldsymbol{dir}$ for location *voxel*. Figure 1B illustrates each of the voxels and vectors that factor into this cost function. C is initialized to 0 for the seed voxels and to ∞ for all other voxels.

In contrast with the scalar F* algorithm, the multi-directional F* algorithm keeps track of the current optimal cost for each voxel *along each sampled ODF direction*, as well as the length and the starting vertex (i.e. origin) of the paths associated with each of these optimal costs at each voxel along each sampled ODF direction. This results in significantly larger memory costs than the scalar F* algorithm; however, the computation time is nonetheless considerably smaller than Monte Carlo approaches.

2.4 Length-independent connectivity cost

A major shortcoming of current connectivity assessment methods is the length-dependency of the resulting connectivity strengths; longer tracts are biased against in stochastic tractography, Hamilton-Jacobian and most graph-based methods. This essentially means that especially the Hamilton-Jacobi and graph-based algorithms may favor unrealistic "shortcuts" rather than longer paths that may be more true to the anatomy and local tensor alignment. The same weakness also applies to the probability formulation of our algorithm presented in

[2]. However, the cost formulation we present here lends itself easily to length-independent connectivity computation. This is achieved by considering the minimum *average* cost of a path rather than the minimum value of the accumulated cost itself (see Fig. 2).

$$\hat{C}^{\boldsymbol{out}}_{current} = \min_{\substack{n=neighbor \\ \boldsymbol{in}=direction}} \frac{C^{\boldsymbol{in}}_n + cost(\boldsymbol{in}, \boldsymbol{out}, \boldsymbol{d_n}, ODF_{current}, ODF_n)}{L^{\boldsymbol{in}}_n + |\boldsymbol{d_n}|} \quad (2)$$

where $L^{\boldsymbol{in}}_n$ represents the length of the current best path to neighbor n at direction $\boldsymbol{in}$, which is the sum of the lengths of each edge along the path from the seed to n. The length of a graph edge will be equal to 1 for an edge between voxels with a common face, $\sqrt{2}$ for an edge between voxels with only one common edge, and $\sqrt{3}$ for an edge between voxels with only one common vertex.

With this modification, the cost of a node will only be updated if the necessary edge traversal cost is less than the average cost of the path, thus reducing the average cost of the path. This allows for longer paths to remain possible even if less anatomically accurate "shortcuts" exist.

2.5 Cost of traversing a graph edge

The cost function used in Eqn. 2 is based on the diffusion in the current and neighboring voxel as well as the consistency of the incoming direction *in*, outgoing direction *out* and the direction and distance towards the neighbor n:

$$cost(\boldsymbol{in}, \boldsymbol{out}, \boldsymbol{d_n}, ODF_{current}, ODF_n) = |\boldsymbol{d_n}| * \Big[\alpha * [f(ODF_{current}, \boldsymbol{out}) \quad (3)$$
$$+ f(ODF_n, \boldsymbol{in})] + [Penalty(\boldsymbol{in}, \boldsymbol{out}) + Penalty(\boldsymbol{in}, \boldsymbol{d_n}) + Penalty(\boldsymbol{out}, \boldsymbol{d_n})]\Big]$$

where *Penalty* is an angle penalty function between two given directions and f is a diffusion weighting function, both of which are further discussed below. α is a weight that can be set based on the chosen f and the parameters of the *Penalty* function. In most of our experiments, we have simply opted for equal weighting of the two terms, i.e. $\alpha = 1$ which we empirically found to yield satisfactory results. A more thorough empirical evaluation of this parameter was conducted for the synthetic datasets where ground truth was known. As can be expected from studying the above equation, we found that lower values of α rely heavily on path consistency and may be appropriate for processing noisy images where this can serve as a regularization term. Higher values of α, on the other hand, assign more importance to the raw data than to path consistency; this may be an appropriate setting for images with higher SNR ratio and/or when curvier fiber bundles are being investigated.

Angle penalty function. We use an angle penalty function to allow the propagation of the diffusion only between two directions that are close to each other, that is to say, that have a small angle between them. This is to prevent the propagation from going backward and less likely along orthogonal paths. This function

penalizes paths that are highly curved and reduces the likelihood of 'switching' paths at crossing fiber voxels. This penalty metric is based on the angle between the two directions and has a value between 0 (no loss in propagation) and 1 (no propagation allowed):

$$Penalty(\boldsymbol{v_1}, \boldsymbol{v_2}) = \frac{1}{2} * (1 - erf(\frac{(\boldsymbol{v_1} \cdot \boldsymbol{v_2}) - \mu}{\sigma})) \tag{4}$$

The μ and σ parameters determine the mean and spread of the error function. We choose $\mu = 0$ since we want the penalty function to be centered around a 90 degree angle, with going "backward" having the highest penalty. Figure 1C illustrates the angle penalty for $\sigma = 1$, as used for all results presented in the next section. The angle penalty function may allow for a tighter or a wider propagation along crossing fibers based on σ. If the function allows for a too wide angle of propagation, there will be almost no loss of connectivity when crossing over to other fibers. Our tests so far showed that the results are rather stable with respect to the σ parameter. Further testing and comparison with alternative penalty functions remain for future work.

Diffusion weighting function $f(ODF)$. Since we want the cost at a voxel in a given direction to be inversely proportional to the amount of diffusion at that voxel along that direction, we need to define a diffusion weighting function that will transform the input ODF data. Given the ODF ϕ, and the direction d, we have initially tried the simplest possible choice by using $f(\phi, \boldsymbol{d}) = 1 - \phi[\boldsymbol{d}]$, which has, as expected, yielded great computational efficiency. However, we found that this diffusion weighting function was too simplistic and did not accurately capture the underlying data. In particular, consider two ODF's, one with two peaks, O_1, and one with a single peak, O_2 (on the unit hemisphere). The above weighting function would assign the same value to the direction along one of the peaks in O_1 and a direction that is relatively separated from the single peak in O_2. The desired behavior, clearly, is to assign a smaller cost to directions that are the preferred directions of diffusion according to the ODF. For this reason, we have used a Finsler-type weighting function modeled after Melonakos et al. [10]. Given the total number of sampled directions $nDirs$, we define the following weighting function:

$$f(\phi, \boldsymbol{d}) = 1 - \frac{nDirs * \phi[\boldsymbol{d}]}{\sum_{dir} \phi[\boldsymbol{dir}] \cdot \boldsymbol{d}_\perp} \tag{5}$$

where $\boldsymbol{d}_\perp$ represents the normal to the ODF surface at $\boldsymbol{d}$. This function assigns weights based on the ratio of diffusion along the current direction $\boldsymbol{d}$ to diffusion along all other sampled directions projected to the plane whose normal $\boldsymbol{d}$ represents.

3 Results

Simple Synthetic Data. We first evaluate our algorithm on three synthetic 3D examples (Fig. 2): a) a single fiber tract, b) two crossing fiber tracts with

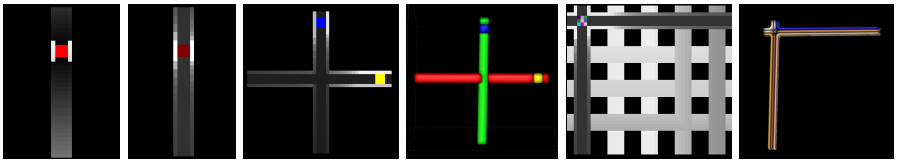

Fig. 2. *Single fiber*: minimum cost (a), average cost (b). Minimum cost increases away from the source, while the average cost remains constant. *Crossing fibers* with two seed regions. The average cost map(c) and the volumetric rendering of the paths originating from each seed region (d) clearly demonstrate the algorithm can resolve the crossing fibers. Note that in (b) and (c), the high cost of taking the expensive low-diffusion path is distributed over longer fibers, which explains the gradient along the outer band. *Grid dataset*. The average cost map (e) shows only the fibers that go through the source can be cheaply reached; the other fibers, although accessible, require multiple sharp turns that do not agree with the diffusion data. The 3D rendering of the origins of the "cheapest" paths (f) shows each path can be traced back to the originating voxel.

the source in each one of the two tracts and c) a grid of fiber tracts with a single source. Computation time for all synthetic datasets was under 2 minutes on a standard linux 64 workstation for 92 sampled directions for the ODF.

For the single source and tract settings, the diffusion propagates along the synthetic fibers; while the minimum cost of a path increases with distance to the source, the average cost remains the same. The crossing fiber synthetic example illustrates how well our method handles/propagates along crossing fibers as well as the ease of obtaining volumetric renderings of the paths as a side product of the algorithm. Here, we thresholded the average cost map to create a "propagation map". Fig. 2d shows a volumetric rendering of the seed labels associated with each voxel inside this propagation map. Clearly, the algorithm was able to successfully resolve the crossing fibers, while keeping track of the origin of each track, its length, as well as its average cost, all in a very efficient manner.

The grid example tests a multiple crossing fibers situation (Fig. 2e and 2f). The 3D rendering clearly shows that the algorithm was able to resolve the crossing fibers successfully, as well as distinguish between each voxel in the source. The average cost map shows that the cost propagation was effectively carried out to the entire grid, even though the cost was high to reach the regions where multiple sharp turns that deviate from the local diffusion model are necessary.

Numerical Fiber Generator. While the simple synthetic data presented above illustrates the basic function of our algorithm and is useful for demonstrating the key concepts, it is much too simplistic to be used for validation purposes. In real life, fibers seldom align perfectly with grid axis and images have considerable noise. For these reasons, we have tested our algorithm using a publicly available synthetic dataset created by a software, Numerical Fiber Generator (NFG)[3] for generating random numerical structures consisting of densely packed bundles of fibers, including kissing and crossing fibers, fibers of

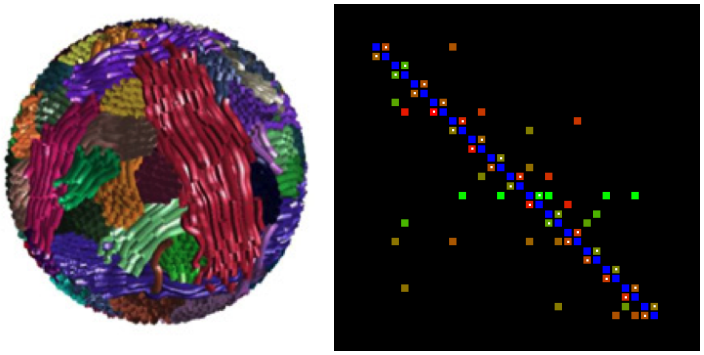

Fig. 3. NFG data. Left, a representation of the phantom A (figure courtesy of Close et al.) Right, connectivity matrix between the 30 seed regions. The color of each square represents the computed cost (red=high, green=low); only connections cheaper than the true connection are shown for each seed. Blue squares represent self-connections (0 cost); squares marked with white dots are true connections.

various radii and realistic noise. For our experiment, we have used the phantom A distributed with the NFG software, which consists of a matrix size $20 \times 20 \times 20$, packing 15 fiber bundles in various configurations (Fig. 3).

In this experiment, we have computed connectivity from each one of the 30 seed regions using 160 sampled directions. To allow for the recovery of the curvier paths, we set $\alpha = 8$. The computation time was under 5 minutes. We have also computed streamline tractography (using Slicer3, www.slicer.org) for illustrating the 'true' paths. We have filtered the computed streamlines to only show the strands that pass through the target region. For fair comparison, we have added the two cost maps starting from the source and target regions. Figure 4 shows the volume rendering of the cost maps overlaid with the streamline results for all 15 fiber bundles. The cost maps were thresholded at 8 percent of the maximum cost value, an empirically set value that needs to be further investigated. We further constructed a connectivity matrix by computing connectivity cost from each source region and averaging the cost at each target region. The resulting connectivity matrix (Fig. 3 right), indicates that all true connections are recovered, with very few "false positives".

Primate Data. We applied our method to *in-vivo* primate data as a first application on real datasets. Ten rhesus monkeys of age 12 months were scanned on a 3T Siemens Trio scanner with 8-channel phase array trans-receiving volume coil. DWI were acquired along 60 directions with voxel size of $1.3 \times 1.3 \times 1.3 mm^3$, with a matrix size of $128 \times 128 \times 80$. A population average was computed via the DWI atlas building method described in [11]. The ODF image was computed on this DWI atlas. Source regions were manually placed in the internal capsule (3 consecutive axial slices), in the genu of the corpus callosum (CC) (single sagittal slice) and in the whole length of the CC (single sagittal slice). Fig. 5 shows the

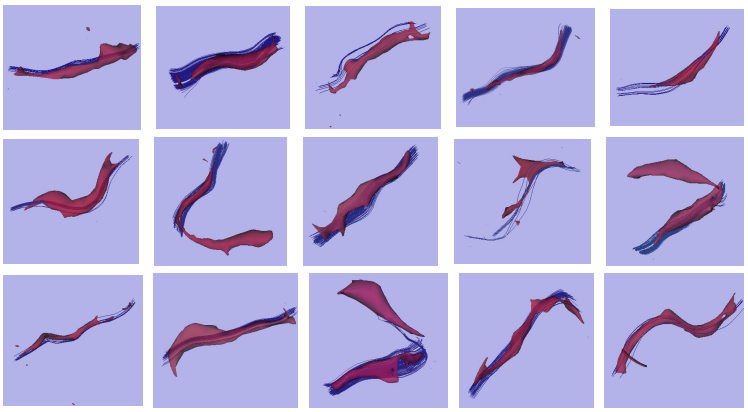

Fig. 4. Individual fiber bundle paths recovered from the NFG phantom. The cost maps were thresholded and rendered using marching cubes (red). For each seed region, we also show the streamline tractography results for illustration purposes (blue).

computed average costs, using 40 sampled directions on the unit sphere. These computations took approximately 80 minutes. Note that all the primate cost maps were computed with the same parameters and resulted in comparable cost values despite the differences in ROI size as well as tract length and shape. It should be noted that for this dataset, the atlas building process was done using co-registration based on a single tensor representation. This leads to a rather strong smoothing of the ODF data and the derived connectivity maps.

Importantly, note the lateral projections of the genu into the frontal lobe; this clearly demonstrates that our method is able to successfully resolve *crossing fibers*. Similarly, the internal capsule connectivity maps show large lateral sections of the cerebrospinal tract, which are only accessible via crossing fibers.

Mouse Data. We next applied our connectivity method to *ex-vivo* mouse data to test its performance at very high resolution settings. The DWI were acquired along 42 directions with a voxel size of $59 \times 71 \times 75$ microns, leading to a matrix size of $159 \times 288 \times 188$. A seed region was placed on a single coronal slice for the fornix region. We computed connectivity using 20 sampled directions, which took approximately 4 hours. Comparison of this with the performance on the monkey dataset shows that the algorithm's run time is mainly driven by the number of sampled directions (approximately quadratic), followed by spatial resolution (approximately linear), which is consistent with the theoretical analysis. Fig. 6 shows the computed average cost, overlayed with the streamline tractography results for comparison purposes. Note the left-right asymmetry in the streamline tractography, which is accurately captured by the cost map.

Human data. The final dataset consists of a single healthy adult volunteer scanned with the same protocol at 3 different b-values (1000, 2000, 3000) with

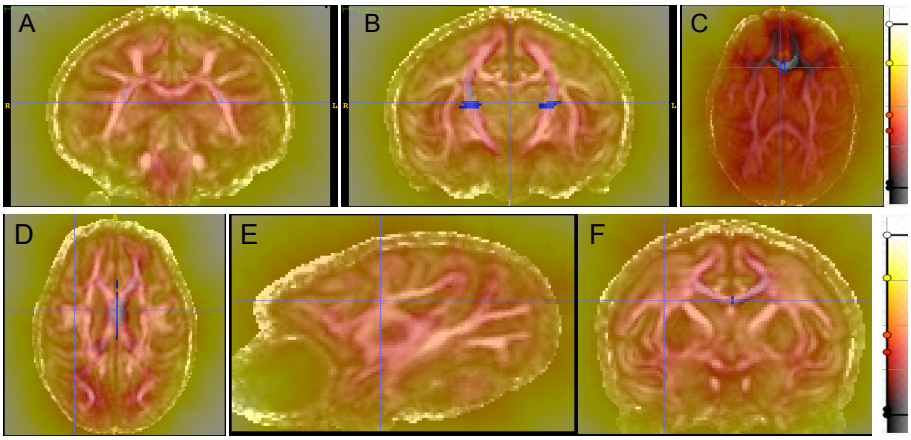

Fig. 5. Average cost map computed on the rhesus monkey DWI atlas with the source located in the internal capsule (A, B); the genu of the CC (C); and the full length of the CC (D, E, F). The grayscale overlays show the FA map, the heat map shows average cost. The source regions are shown in blue.

a Siemens Tim Trio 3T scanner. DWI were acquired along 64 directions with an isotropic voxel size of 2mm, with a matrix size of $106 \times 106 \times 76$. After rigidly registering the three scans, a manually placed ROI in the genu of the corpus callosum was used as source. The average cost maps, computed using the exact same parameters (20 sampled directions on the unit sphere, computation time 10-12 minutes) as well as the same ROI, are visualized in Figure 7 using the same color map. As expected, the connectivity maps appear less smoothed and we find sharper connectivity as the b-value increases. Additionally, this dataset further illustrates both the robustness of our algorithm to noise and the comparability of our connectivity metric across different scan parameters, ROI and brain sizes, tract properties, and species.

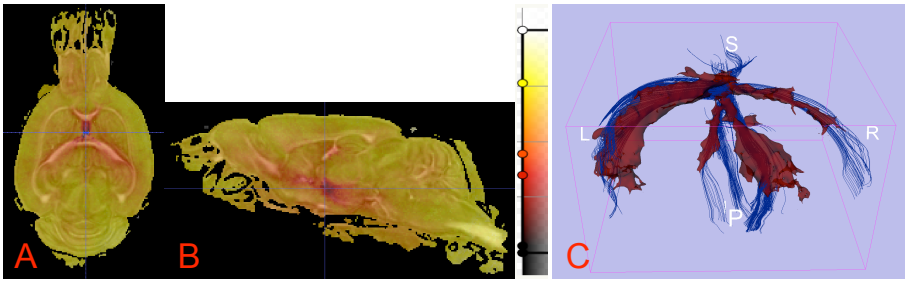

Fig. 6. Average cost map for fornix on a mouse. (A) axial view, (B) sagittal view (C) 3D rendering of the thresholded cost map (red) and streamline tractography (blue). The grayscale overlays show the FA map, the heat map shows the average cost.

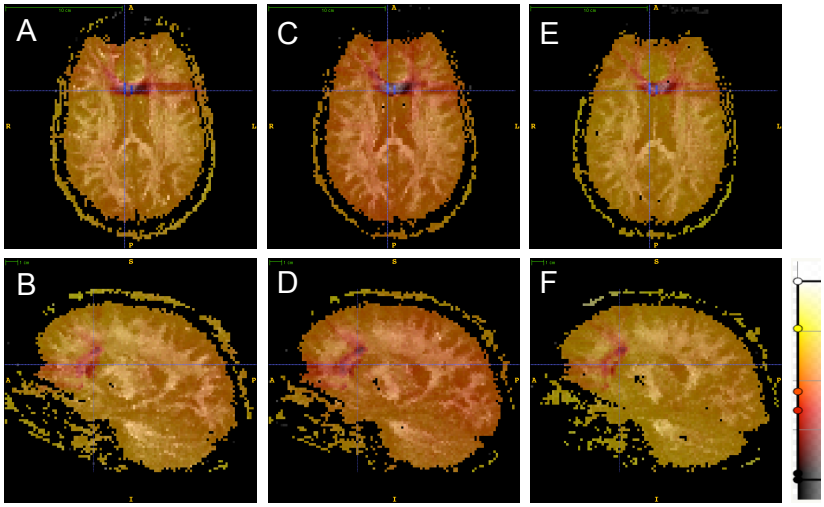

Fig. 7. The average cost maps for the genu tract on the same human subject for different b-values: 1000, 2000, 3000 (left to right; top: axial view, bottom: sagittal view). Note that the connectivity map becomes sharper with higher b-values, in addition to more robust tracking.

4 Discussion

We present a new method for the computation of diffusion imaging based white matter connectivity. This method is efficient and resilient to noise, handles multiple seed regions straightforwardly and works well in presence of crossing fiber tracts. The proposed method is generic and could be easily applied to non-diffusion data as long as local directional data can be derived.

A significant contribution of this method is the ability to recover longer fiber tracts, which are typically biased against due to accumulated cost or loss of probability. Our average cost metric overcomes these difficulties. While in theory the average cost metric will in fact favor longer tracts (if the propagation cost is infinitesimally small while the tract length increases, the average cost will be reduced by following a long tract), these situations will only occur in synthetic data since the propagation cost is not negligible in real data where perfect grids of tracts colinear with the diffusion axes do not exist. Therefore, our method in practice offers a length-independent connectivity metric as demonstrated by the presented datasets. Such a metric is far more suitable than length-dependent metrics for subject-specific studies of neuropathology as well as network studies where the connectivity strengths of various ROI's need to be compared. Additionally, it should be noted that the connectivity maps, rather than thresholded binary fiber reconstructions (which are used here for visualization purposes), are the primary product of this method targeting network studies.

Currently we are performing more thorough evaluation studies of the method using human, non-human primate as well as rodent imaging data. Furthermore,

we are working on the empirical evaluation of parameters and the improvement of the ODF weight function and the angle penalty function.

Acknowledgments. This work is funded by the following funding sources: R41 NS059095, R21 MH084132, P50 MH078105, P50 MH078105-01A2S1, RC1-AA019211, P01-DA022446, U01-AA020022, NAMIC U54 EB005149, R01 MH091645, ACE R01 HD055741, IDDRC P30 HD03110.

References

1. Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Bray, J., Smith, S.: Characterization and propagation of uncertainty in DW-MRI. Magnetic Resonance in Medicine 50(5), 1077–88 (Jan 2003)
2. Boucharin, A., Oguz, I., Vachet, C., Shi, Y., Sanchez, M., Styner, M.: Efficient, graph-based white matter connectivity from orientation distribution functions via multi-directional graph propagation. Proc. SPIE 7962, 79620S (2011)
3. Close, T.G., Tournier, J.D., Calamante, F., Johnston, L.A., Mareels, I., Connelly, A.: A software tool to generate simulated white matter structures for the assessment of fibre-tracking algorithms. NeuroImage 47(4), 1288–300 (Oct 2009)
4. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical Q-ball imaging. Magn Reson Med 58(3), 497–510 (Sep 2007)
5. Draganski, B., Kherif, F., Klöppel, Cook, P., Alexander, D., Parker, G., Deichmann, Ashburner, Frackowiak, R.: Evidence for segregated and integrative connectivity patterns in the human basal ganglia. J Neuroscience 28, 7143–52 (2008)
6. Fischler, M., Tenenbaum, J., Wolf, H.: Detection of roads and linear structures in low-resolution aerial imagery using a multisource knowledge integration technique. Readings in computer vision: issues, problems, principles, and paradigms (1987)
7. Fletcher, T., Tao, R., Jeong, W., Whitaker, R.: A volumetric approach to quantifying region-to-region WM connectivity in DT-MRI. IPMI pp. 346–58 (2007)
8. Iturria-Medina, Y., Canales-Rodríguez, E.J., Melie-García, L., Valdés-Hernández, P.A., Martínez-Montes, E., Alemán-Gómez, Y., Sánchez-Bornot, J.M.: Characterizing brain anatomical connections using DWMRI and graph theory. NeuroImage 36(3), 645–60 (2007)
9. Jones, D.K., Horsfield, M.A., Simmons, A.: Optimal strategies for measuring diffusion in anisotropic systems by MRI. Magn Reson Med 42(3), 515–25 (1999)
10. Melonakos, J., Mohan, V., Niethammer, M.: Finsler tractography for white matter connectivity analysis of the cingulum bundle. MICCAI (Jan 2007)
11. Niethammer, M., Shi, Y., Benzaid, S., Sanchez, M., Styner, M.: Robust model-based transformation and averaging of DWI applied to DW atlas construction. MICCAI-CDMRI (2010)
12. Pichon, E., Westin, C.F., Tannenbaum, A.R.: A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography. MICCAI 8, 180–7 (2005)
13. Sotiropoulos, S.N., Bai, L., Tench, C.R.: Fuzzy anatomical connectedness of the brain using single and multiple fibre orientations estimated from diffusion MRI. Comput Med Imaging Graph 34(6), 504–13 (Sep 2010)
14. Styner, M., Coradi, T., Gerig, G.: Brain morphometry by distance measurement in a non-Euclidean, curvilinear space. IPMI (Jun 1999)
15. Zalesky, A.: DT-MRI fiber tracking: a shortest paths approach. IEEE Trans Med Imaging 27(10), 1458–71 (Oct 2008)