Published in final edited form as:

Proc SPIE. 2011 ; 7962: 79620S. doi:10.1117/12.878291.

Efficient, Graph-based White Matter Connectivity from Orientation Distribution Functions via Multi-directional Graph Propagation

Alexis Boucharin^a, Ipek Oguz^a, Clement Vachet^a, Yundi Shi^a, Mar Sanchez^c, and Martin Styner^{a,b}

^aDepartment of Psychiatry, University of North Carolina, Chapel Hill, USA

^bDepartment of Computer Science, University of North Carolina, Chapel Hill, USA

^cDepartment of Psychiatry, Emory University School of Medicine, Atlanta, USA

Abstract

The use of regional connectivity measurements derived from diffusion imaging datasets has become of considerable interest in the neuroimaging community in order to better understand cortical and subcortical white matter connectivity. Current connectivity assessment methods are based on streamline fiber tractography, usually applied in a Monte-Carlo fashion. In this work we present a novel, graph-based method that performs a fully deterministic, efficient and stable connectivity computation. The method handles crossing fibers and deals well with multiple seed regions. The computation is based on a multi-directional graph propagation method applied to sampled orientation distribution function (ODF), which can be computed directly from the original diffusion imaging data. We show early results of our method on synthetic and real datasets. The results illustrate the potential of our method towards subject-specific connectivity measurements that are performed in an efficient, stable and reproducible manner. Such individual connectivity measurements would be well suited for application in population studies of neuropathology, such as Autism, Huntington's Disease, Multiple Sclerosis or leukodystrophies. The proposed method is generic and could easily be applied to non-diffusion data as long as local directional data can be derived.

1. INTRODUCTION

The use of regional connectivity measurements derived from diffusion imaging datasets has become of considerable interest in the neuroimaging community in order 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.

Stochastic tractography methods overcome some of these problems by repeatedly applying streamline fiber tractography in a Monte-Carlo fashion.^{1, 2} 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.^{3–6} 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.³ These methods are inherently more robust to noise in diffusion data than tractography methods and they 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. Furthermore, these methods do not take into account the consistency of path orientation along the minimal path.

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. Such individual connectivity measurements would be well suited for application in population studies of neuropathology, such as Autism, Huntington's Disease, Multiple Sclerosis or leukodystrophies.

2. METHODS

Motivated by the well-known F^* (pronounced “f-star”) graph traversal algorithm⁷ and previous work using F^* to compute non-linear distances within the brain,⁸ we extended our own implementation of F^* to incorporate multiple incoming and outgoing directions for computing the overall probability or cost of a voxel. While we implemented two formulations of this algorithm, one for propagating costs and one for propagating probabilities, in this paper we will focus purely on the latter.

2.1 Scalar F^* graph traversal

The basic F^* algorithm⁷ 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 reduce 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.

The F^* algorithm keeps track of the traversal probability of each edge and the current optimal probability for each graph node (voxel). Note that this basic F^* algorithm is only suitable for graph traversal problems where the probability associated with traversing an edge is just a 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.⁹ 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. This ODF image is sampled at each voxel into a spherical sampled distribution using a standard icosahedron subdivision scheme.¹⁰ Each vertex corresponds to one direction that can be defined by two angles on the sphere (θ and ϕ). Such a subdivision allows a high number of directions N in each voxel based on the desired level of accuracy, depending on the level of subdivision

chosen SL , such that $N=12+30*SL+20*\frac{(SL-1)*SL}{2}$.

All voxels included in the source will be represented as an isotropic/uniform ODF (same diffusion value along each direction).

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 graph node. Each node (voxel) visit consists of updating the probability for each sampled direction of the current voxel based on the probability of its neighbors. To this end, for each sampled direction *out*, for each neighbor n , the probability of arriving to n in direction *in* and traversing an additional edge along *out* is computed. The probability thus obtained is:

$$P_{current}^{out} = ODF_{current}^{out} * \sum_{n=neighbor} \sum_{in=direction} P_n^{in} * prob(\vec{in}, \vec{out}, \vec{dn}) \quad (1)$$

where $\vec{dn}$ represents the vector from the voxel-center of *current* to the voxel-center of n . The $ODF_{current}^{out}$ terms refer to the value of the input orientation distribution function sampled at direction $\vec{out}$ for location *current*. Figure 2 illustrates each of the voxels and vectors that factor into this probability function.

In contrast with the scalar F* algorithm, the multi-directional F* algorithm keeps track of the traversal probability for each edge *along each sampled direction* and the current probability for each graph node *along each sampled direction* at each node.

2.3.1 Probability of traversing a graph edge—The probability function used in Eqn. 1 is based on the consistency of the incoming direction *in*, outgoing direction *out* and the direction and distance towards the neighbor n :

$$prob(\vec{in}, \vec{out}, \vec{dn}) = |\vec{dn}| * Similarity(\vec{in}, \vec{out}) * Similarity(\vec{in}, \vec{dn}) \quad (2)$$

where *Similarity* is an angle similarity function between two given directions discussed in detail in the next subsection.

2.3.2 Angle similarity function—We use an angle similarity 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 from doing infinite loops (situations where a voxel would participate in its own diffusion via a few intermediate voxels). This function penalizes paths that are highly

curved and reduces the likelihood of ‘switching’ paths at crossing fiber voxels. This similarity metric is computed via an error function (*erf*) applied to the cosine of the angle between the two directions. The similarity metric has a value between 0 (no propagation allowed) and 1 (no loss in propagation):

$$\text{Similarity}(\vec{v}_1, \vec{v}_2) = \frac{1}{2} * (1 - \text{erf}(\vec{v}_1 \cdot \vec{v}_2)) \quad (3)$$

2.4 Normalization

In order to ensure no diffusion is lost when traversing a voxel, we need to normalize the probabilities at each iteration. This is accomplished by computing the sum over the incoming diffusion probabilities and the outgoing probabilities, and scaling each probability by the ratio of these two.

$$P_{current}^{\vec{out}} = P_{current}^{\vec{out}} * \frac{\sum P_{incoming}}{\sum P_{outgoing}} \quad (4)$$

2.5 Computational improvements

The final piece of the algorithm is a computation speed-up obtained by sharpening of the ODF functions. In comparison with synthetic datasets, ODFs from real datasets are often smooth and do not possess very sharp diffusion distributions. Thus, before we applied our method to real datasets, the ODFs were sharpened (Fig. 3). This ODF sharpening procedure first normalizes the diffusion probability values so that they sum to 1 on each hemisphere and raise each value to the power of x ($x = 6$ in this case). To improve the sharpening we further apply a threshold to the diffusion probability values. The computation thus takes in account only non-thresholded directions of propagation at each voxel.

3. RESULTS

3.1 Synthetic Data

In order to evaluate the proposed algorithm, we first applied it to several synthetic examples: a) a single strong fiber tract, b) two strong crossing fiber tracts with the source in one of the two tracts and c) a grid of fiber tracts. These examples were not created through simulated diffusion images but rather as directly as sharp (but still somewhat smooth) probability distribution images.

From a single source, the diffusion propagates along the synthetic fibers and fades further from the source (see Fig 4).

The crossing fiber synthetic example illustrates how well our method handles/propagates along crossing fibers. Since we use a soft angle similarity function, the propagation is allowed to even cross over into the crossing fiber, but those values are significantly smaller than those on the main fiber.

The last example, the grid situation is made to test an possible endless loop situation, where the propagation potentially could enter an endless loop of ever changing connectivity. The choice of the angle similarity function allows for a tighter or a wider propagation along crossing fibers in the grid. If the function allows for a too wide diffusion, there will be almost no loss of connectivity when crossing over to other fibers and it will keep modifying

the local connectivity values. Our tests showed that the proposed angle similarity function is rather stable and loop issues did not arise in our tests.

3.2 Real Data

In comparison with synthetic datasets, ODFs from real datasets are often smooth and do not possess very sharp diffusion distributions. Thus, before we applied our method, the ODFs were sharpened. This ODF sharpening procedure first normalizes the diffusion probability values so that they sum to 1 on each hemisphere and raise each value to the power of x ($x = 6$ in this case). To improve the sharpening we further apply a threshold to the diffusion probability values. The computation will take in account only non-thresholded directions of propagation at each voxel. This threshold based sharpening will also solve the problem of encountering uniform distribution in the propagation, as voxels with uniform or close to uniform diffusion distribution will be fully thresholded away. These fully thresholded voxels form a threshold mask that also considerably improves computational efficiency (see Fig. 6).

We applied our method to non-human primate data as a first application on real datasets. The data was acquired from a 6 month old rhesus macaque on a Siemens Trio 3T scanner at the Yerkes Primate Center at Emory University with 60 diffusion weighted scans ($b=1000$) at unique gradient directions and 2 non-diffusion ($b=0$) scans. Figures 5–7 show the computed white matter mask and the corresponding result of the probabilities. As expected, we can see that the computed connectivity map follows the threshold mask (6), diffuses along the white matter regions and fades the further it is from the source. The results display the expected result and illustrate the potential of the method. Computation time was 3.18 hours on a standard 64 bit linux workstation.

4. CONCLUSION

In this work, we present a new method for the computation of diffusion imaging based white matter connectivity. This method allows for efficient computation, resilience 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.

Currently we are further improving the efficiency of the method and perform more thorough evaluation studies of the method using human, non-human primate as well as rodent imaging data.

Acknowledgments

This work was funded by the NADIA U01-AA020022 Scientific Core, the NIH Program Project IP01DA022446-02 and the National Alliance for Medical Image Computing (NA-MIC) grant U54-EB005149. We would like to further acknowledge the financial support of the NIH grant P50 MH078105-01A2S1, NIH STTR grant R41 NS059095, the UNC Neurodevelopmental Disorders Research Center HD 03110, the NIH grant RC1AA019211. We are grateful to Francois Budin for insightful discussions and technical support.

REFERENCES

1. Behrens TEJ, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, Matthews PM, Brady JM, Smith SM. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn Reson Med*. 2003 Nov;50:1077–1088. [PubMed: 14587019]
2. Draganski B, Kherif F, Klöppel S, Cook PA, Alexander DC, Parker GJM, Deichmann R, Ashburner J, Frackowiak RSJ. Evidence for segregated and integrative connectivity patterns in the human basal ganglia. *Journal of Neuroscience*. 2008 Jul;28:7143–7152. [PubMed: 18614684]

3. Fletcher PT, Tao R, Jeong W-K, Whitaker RT. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. *Inf Process Med Imaging*. 2007 Jan. 20:346–358. [PubMed: 17633712]
4. O'Donnell L, Haker S, Westin C. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. *Lect Notes Comput Sc*. 2002 Jan.2488:459–466.
5. Pichon E, Westin C-F, Tannenbaum AR. A hamilton-jacobi-bellman approach to high angular resolution diffusion tractography. *Med Image Comput Comput Assist Interv*. 2005 Jan.8:180–187. [PubMed: 16685844]
6. Melonakos J, Mohan V, Niethammer M. Finsler tractography for white matter connectivity analysis of the cingulum bundle. *MICCAI*. 2007 Jan.
7. 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 Jan.
8. Styner M, Coradi T, Gerig G. Brain morphometry by distance measurement in a non-euclidean, curvilinear space. *IPMI '99: Proceedings of the 16th International Conference on Information Processing in Medical Imaging*. 1999 Jun.
9. Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. *Magn Reson Med*. 2007 Sep.58:497–510. [PubMed: 17763358]
10. Styner M, Oguz I, S X, C B, D P, J J L, M E S, G G. Framework for the statistical shape analysis of brain structures using SPHARM-PDM. *Insight Journal*. 2006

Figure 1.
An Orientation Distribution Function (ODF) representation at a single voxel (left). Spherical distribution (middle) thanks to an icosahedron subdivision (right).

Figure 2.

This diagram shows the variables that affect the probability of traversing an edge in 2D. For the voxel *Current*, for each direction $\vec{out}$, the algorithm considers all neighboring voxels (shown in green) and all the possible directions (shown in red) at these voxels in order to identify the neighbor *n*, which lies in the direction $\vec{dn}$ from the *Current* voxel, and the direction $\vec{in}$ at the neighbor *n*. The probability associated with voxel *Current* for direction $\vec{out}$ is a weighted sum of the probabilities associated with all these possible paths, where the weight is proportional to the likelihood of the path continuing in that specific direction.

Figure 3.
Smooth ODF (left) and the sharpened ODF (right) at one voxel.

Figure 4. Connectivity maps computed on the synthetic examples: unique fiber (left), crossing fibers example (middle), fiber grid example (right). The isotropic source is displayed in red.

Figure 5.
Axial slice of the macaque brain dataset with source region (red) overlaid on the Fractional Anisotropy (FA) image.

Figure 6.
Computed threshold mask for the macaque brain dataset as shown in three orthogonal slices in Insight-SNAP.

Figure 7. Connectivity map for the macaque brain dataset as shown in 3D Slicer. Connectivity map is shown overlaid on the FA image (yellow to red = high to low connectivity).