DISCRETIZING STOCHASTIC TRACTOGRAPHY: A FAST IMPLEMENTATION

Juan Eugenio Iglesias[*]

Medical Imaging Informatics, UCLA

Paul Thompson, Cheng-Yi Liu, Zhuowen Tu

Laboratory of Neuro Imaging, UCLA

ABSTRACT

Probabilistic tractography has emerged as an alternative to classical deterministic methods to overcome their lack of connectivity information between different brain regions. However, it relies on statistical sampling, which is computationally taxing. In this study, a well-known, random walk based stochastic tractography method is discretized by limiting the set of directions that a sampling particle can follow. This sets up to a framework based on a Markov chain that can accommodate all the desirable features of stochastic tractography, principally trajectory regularization through particle deflection. The system produces results that are comparable to those by the stochastic algorithm it is based on ($\rho = 0.79$), though 60 times faster.

Index Terms— HARDI, stochastic, tractography, fast

1. INTRODUCTION

Diffusion weighted magnetic resonance imaging (DW-MRI) is the only way of imaging axonal fiber tracts in vivo. Comparing the T2 intensities with and without gradients in the magnetic field in different directions, a water diffusion profile can be estimated for each spatial location. From this data, tractography algorithms are able to reconstruct the fibers. Tractography has potential in surgical planning to avoid damaging tracts with important functionality.

The diffusion profile can be reconstructed from discrete measurements in different ways. The simplest is diffusion tensor imaging (DTI) [1, 2], which fits a zero-mean Gaussian probability distribution function (PDF) to the measurements. To overcome the limitation that a Gaussian is a unimodal distribution and can handle only one fiber population per voxel, a high number of angles must be sampled. This is the principle of high angular resolution diffusion imaging (HARDI) [3]. However, the HARDI data cannot be used directly in tractography because the minima of the DW-MRI measurements do not correspond directly to fiber orientations when two or more fibers intersect. To overcome this problem, the orientation distribution function (ODF) is typically used. The ODF is the PDF of the water diffusivity on a spherical shell [4].

Stochastic tractography has emerged as an alternative to classical deterministic methods when the goal is to obtain connectivity information between different areas of the brain rather than following individual tracts. Based on Monte Carlo sampling, probabilistic tractography simulates the trajectory of a high number of particles that stem from a seed region and follow a regularized random walk in the tensor or ODF field. The image that results from computing the fraction of particles that reaches each voxel in the volume is known as a tractogram, and it measures the probability of being connected to the seed. Apart from providing quantitative information on connectivity, probabilistic tractography also has the advantages of being more robust against misplacement of the seed than deterministic methods and inherently accounting for the uncertainty of the fiber orientations. However, these advantages come at the expense of a much higher computational load, because a reliable tractogram typically requires at least one million particles.

This study presents an approach that can quickly compute the tractogram analytically at the expense of limiting the directional resolution of the HARDI data. The diffusion from a seed region is modeled as a discrete Markov chain in which the state variable is a bivector that encodes the probability mass for each position / direction pair. The evolution of the state variable is governed by a transition tensor which is designed to mimic the behavior of particles in a well-known stochastic tractography method presented by Perrin et al. [5]. Their framework was chosen as a model because it combines all the desirable features of stochastic tractography in a simple and clear manner. The main advantage of our approach is its speed: the tractogram for any seed region can be calculated in a few seconds, allowing the user to interactively explore the connectivity of the brain.

There has been previous work on discretizing the movement of particles to the orientations corresponding to neighboring voxels. The image is represented by a graph in which the connectivity between two voxels is estimated by finding the lowest cost path between them. The cost is minimized with dynamic programming in [6], A^* in [7], and a variant of Dijkstra's algorithm in [8]. The weights of the edges depend on the local diffusion properties of the tissue.

The aforementioned methods focus on calculating the connectivity between two user-defined regions, but they do not provide the full brain tractogram from a seed region. This feature was present in a recent study [9] in which the connectivity between two voxels was defined as the lowest affinity along the strongest path between them. Still, none of these methods incorporate all the key elements of stochastic tractography, particularly particle deflection in a high angular resolution ODF field. Moreover, they do not really evaluate the impact of limiting the transitions to grid directions on the final tractogram. Addressing these two shortcomings is the main contribution of this study. First, we demonstrate how the transition tensor can be designed to mimic the behavior of sampling particles in the stochastic method by Perrin et al. Second, the output of our method is quantitatively evaluated by comparing it with the results provided by Perrin et al.'s method in the same dataset.

2. MATERIALS

10 HARDI images acquired at the Center for Magnetic Resonance at the University of Queensland using a 4 Tesla Bruker Medspec scanner with a transverse electromagnetic headcoil were used in

[*]This work was funded by the NSF grant 0844566 and the NIH grant U54 RR021813. The authors would like to thank G. de Zubicaray, K. McMahon and M. Wright from the University of Queensland for acquiring the data. The first author would also like to thank the U.S. Department of State's Fulbright program for the funding.

this study. Diffusion-weighted scans utilized a single-shot echo planar technique with a twice-refocused spin echo sequence to minimize eddy-current induced distortions. The timing of the diffusion sequence was optimized for signal-to-noise ratio (SNR). 94 diffusion-sensitized gradient directions and 11 baseline images with no diffusion-sensitization were obtained for every subject. Imaging parameters were: b-value = 1159 s/mm^2, TE/TR = 92.3/8,259 ms, voxel size = 1.8 mm 1.8 mm 2.0 mm, image size 128 x 128 x 55 voxels. The acquisition time was approximately 15 minutes. The 11 baseline images were merged into a single estimate of the T2 reference volume which was then used to calculate a mask corresponding to the brain using the well-known BET algorithm for skull stripping [10]. The resulting mask was then applied to all the diffusion images.

3. METHODS

The steps to calculate a tractogram from a HARDI volume are the following: first, the ODF field is estimated from the HARDI data. Then it is deconvolved into the fiber ODF (fODF) field, which has been proven to have a positive impact on the tractography results [11]. From the fODF, the transition tensor can be calculated. At this point, the tractogram can be quickly generated from the state bivector corresponding to any user-defined seed region.

3.1. Calculation of the fODF

The fODF (henceforth $\boldsymbol{\psi}_f(\boldsymbol{r})$) is ideally zero everywhere except for the directions where the voxel is crossed by fibers. Assuming that the ODF (henceforth $\boldsymbol{\psi}(\boldsymbol{r})$) is the result of smoothing the fODF with a blurring kernel i.e. $\boldsymbol{\psi}(\boldsymbol{r}) \propto K\boldsymbol{\psi}_f$, it is possible to recover the fODF using deconvolution. If we define $\boldsymbol{g}(\boldsymbol{r})$ as the $D \times 1$ signal vector at location $\boldsymbol{r}$, the fODF can be estimated from the signal using Tikhonov regularization as follows:

$$\boldsymbol{\psi}_f(\boldsymbol{r}) \approx \frac{1}{Z(\boldsymbol{r})}\left[\left(K^tK+\varepsilon I\right)^{-1}K^t\boldsymbol{\psi}\left(\boldsymbol{r}\right)\right]_+ = \ldots$$

$$= \frac{1}{Z'(\boldsymbol{r})}\left[\left(K^tK+\varepsilon I\right)^{-1}K^tA\boldsymbol{g}\left(\boldsymbol{r}\right)\right]_+ = \frac{1}{Z'(\boldsymbol{r})}\left[A'\boldsymbol{g}\left(\boldsymbol{r}\right)\right]_+ \quad (1)$$

where ε is the regularization constant, $Z(\boldsymbol{r})$ and $Z'(\boldsymbol{r})$ are partition functions that ensure $\mathbf{1}^t\boldsymbol{\psi} = \mathbf{1}^t\boldsymbol{\psi}_f = 1$, and A is a matrix that estimates the regularized spherical harmonic coefficients of $\boldsymbol{g}(\boldsymbol{r})$, modifies them to obtain the coefficients of the ODF, and evaluates them at the set of H directions for which the ODF is to be reconstructed [12]. The matrix A' is the final signal-to-fODF matrix. Tikhonov regularization has the advantage that it can be easily integrated in the matrix framework, while performing comparably to more sophisticated, computationally expensive methods [13]. The blurring matrix K is estimated non-parametrically as follows: the N_{an} voxels with the highest anisotropy in the available images are found; their ODFs are calculated and aligned; their azimuthal components are averaged to ensure rotational invariance with respect to z; and the average of the N_{an} calculated profiles is computed. Each of the H columns of K is the result of rotating this estimate to align it with each of the H directions where the ODF is reconstructed.

3.2. The state bivector and the transition tensor

In the random walk approach to stochastic tractography, a set of particles is placed in a seed region and allowed to follow a regularized random walk governed by the ODF at each point until before exiting the brain white matter. Because particles are allowed to move in any direction, they mostly travel along non-grid locations, making it unfeasible to calculate the tractogram analytically. However, if the particles are constrained to grid locations, the stochastic sampling process can be replaced by a bivector-valued Markov chain to obtain the approximate tractogram analytically in a short time.

We define the state bivector $Q = Q^{(t)}_{p,u}$ as the probability mass at spatial location $\boldsymbol{r}_p$ with direction u for discrete time t. The first index ranges from $p = 1, \ldots, N$, the number of voxels, while the second index ranges from $u = 1, \ldots, 98$, corresponding to the unique directions that connect a voxel with the 124 neighbors in the 3-D lattice corresponding to a $5\times5\times5$ cube. Memorizing the direction makes it possible to regularize the trajectory of the probability mass. The 98-neighborhood is a good compromise between the angular resolution, step length variability and computational cost.

The evolution of the state bivector is characterized by T, a tensor of type $(2,2)$ such that $T^{p,u}_{q,v}$ represents the transition probability from position/speed (p,u) to (q,v). Its dimensionality is thus $(N \times 98) \times (N \times 98)$, very large but extremely sparse. The transition from step t to step $t+1$ is a tensor contraction: $Q^{(t+1)}_{q,v} = \sum_{p,u} T^{p,u}_{q,v} Q^{(t)}_{p,u}$. The elements of the transition tensor T can be designed to simulate the behavior of a particle in the stochastic tractography method by Perrin et al. as accurately as possible. In their study, particles move according to the rule:

$$\boldsymbol{d}^{(t)} = \gamma(\boldsymbol{r})\boldsymbol{d}_{fODF}(\boldsymbol{r}) + (1-\gamma(\boldsymbol{r}))\boldsymbol{d}^{(t-1)} \quad (2)$$

where $\boldsymbol{d}_{fODF}(\boldsymbol{r})$ is the result from sampling the fODF at location $\boldsymbol{r}$ restricted to a 30° half cone defined by the incident direction $\boldsymbol{d}^{(t-1)}$. The scalar field $\gamma(\boldsymbol{r})$ corresponds to the standard deviation of the fODF at each voxel normalized to its maximum across the image. Particles move in the ODF field until before exiting the brain white matter.

To mimic this behaviour, the non-zero elements of the transition tensor must satisfy the following conditions: 1) the origin voxel p must the in the brain white matter (a white matter mask is calculated by thresholding $\gamma(\boldsymbol{r})$ and smoothing the result with a closing operator); 2) the origin and destination voxels p and q must be 98-neighbors in the lattice; 3) direction v has to be equal (parallel) to the vector $\boldsymbol{r}_q - \boldsymbol{r}_p$, memorizing the direction of movement; and 4) direction u is deflected into v by a vector $\boldsymbol{d}_{fODF}$ (see equation 2) such that u and $\boldsymbol{d}_{fODF}$ form an angle under a certain threshold (30° in Perrin et al.'s study). The last two conditions regularize the trajectory of the probability mass.

If a transition is compatible with these conditions, then it is assigned a probability proportional to the sum of the values of $\boldsymbol{\psi}_f(\boldsymbol{r}_p)$ in the directions $\boldsymbol{d}_{fODF}$ that deflect direction u into direction v according to equation 2. The probabilities are then scaled to $\sum_q \sum_v T^{p,u}_{q,v} = 1$ (unless the sum is zero, representing a dead end). The algorithm to calculate the elements of T is detailed in Table 1.

3.3. Calculating the tractogram

Once the transition tensor has been computed, the tractogram is generated as follows. First, the seed vector is constructed as:

$$Q^{(0)}_{p,u} = \begin{cases} \frac{1}{|\mathbb{S}|}\sum_{h=1}^{H}\psi_{f,h}(\boldsymbol{r}_p)\delta[\Upsilon(\theta_h,\varphi_h)-u], & \boldsymbol{r}_p \in \mathbb{S} \\ 0, & \boldsymbol{r}_p \notin \mathbb{S} \end{cases}$$

where $\mathbb{S}$ is the set of seed voxels, $\delta[\cdot]$ denotes the Kronecker delta and and $\Upsilon(\theta,\varphi)$ is a function that discretizes an arbitrary angle to

· Calculate and normalize $\gamma(\boldsymbol{r}) \Leftarrow \gamma(\boldsymbol{r})/\max_{\boldsymbol{r}}[\gamma(\boldsymbol{r})]$
· Compute white matter mask: $\gamma(\boldsymbol{r}) > \gamma_{min}$, refine with closing
· Initialize all elements of $T^{p,u}_{q,v}$ to zero
· FOR each position/speed pair (p, u) such that $\boldsymbol{r}_p$ is in the mask
  · SET SUM $\Leftarrow$ 0, PROBS[v] $\Leftarrow$ 0, for $v = 1, \ldots, 98$
  · FOR each direction of the fODF $h = 1, \ldots, H$
    · IF ANGLE(u,h)<ANGLE_MAX
      · SET CDIR $\Leftarrow \gamma(\boldsymbol{r}_p)\boldsymbol{d}_h + (1 - \gamma(\boldsymbol{r}_p))\boldsymbol{d}_u$
      · SET PROBS[$\Upsilon(CDIR)$] += $\boldsymbol{\psi}_h(\boldsymbol{r})$, SUM += $\boldsymbol{\psi}_h(\boldsymbol{r})$
    END
    · IF SUM>0
      · FOR each direction $v = 1, \ldots, 98$
        · SET $q \Leftarrow p + v$ (point+vector=point)
        · SET $T^{p,u}_{q,v} \Leftarrow$ PROBS[v]/SUM
END; END; END; END

Table 1. Pseudocode for calculating the transition tensor. $\Upsilon(\theta, \varphi)$ is a function that downsamples the H directions of the fODF to the 98 directions allowed in our method: $\Upsilon(\theta, \varphi) = u$ when the angle formed by directions (θ, φ) and u is smaller than the angle formed by (θ, φ) and any other direction of the 98.

the 98-neighborhood (see caption of Table 1). Then the distribution at time t is $Q^{(t)}_{q,v} = (T^{p,u}_{q,v})^{(t)} Q^{(0)}_{p,u}$, where $(\cdot)^{(t)}$ denotes t successive contractions. If the contraction is carried out a sufficient number of times t_{max}, the whole probability mass will have exited the volume and the tractogram M can be built by marginalizing $Q^{(t)}_{p,u}$ with respect to direction and time:

$$M(\boldsymbol{r}_p) = \sum_{u=1}^{98} \sum_{t=0}^{t_{max}} Q(t)_{p,u}$$

Because a 5×5×5 neighborhood is used, the probability mass can skip some rows of voxels, causing some striation artifacts. To ameliorate this problem, the final output is calculated by smoothing $M(\boldsymbol{r})$ with a Gaussian kernel and, for visualization purposes, taking the logarithm of the result to compress its range.

4. EXPERIMENTS AND RESULTS

4.1. Setup of experiments

Since the goal of stochastic tractography is to estimate the connectivity between regions in the brain, the following experiment was carried out. First, seeds were placed on well-known tracts in the brain. Target regions of interest (ROI) which are known to be connected to the seeds were then defined. The ROI was a 10mm diameter sphere in all cases. Finally, the connectivities from both methods were obtained by summing the tractograms within the target ROIs and subsequently compared. Four seed / ROI pairs were defined: 1) seed in the genu of the corpus callosum, two target ROIs (left/right) at the front of the brain, near the lower part of the superior front gyrus; 2) seed in the splenium of the corpus callsoum, two target ROIs (left/right) towards the back of the brain, near the most posterior part of the middle occipital gyrus; 3) seed in the corticospinal tract near the hippocampus, target ROI between the leg and hand areas of the motor cortex; and 4) seed in the posterior region of the cingulate gyrus, target ROI in the anterior region.

The system parameters were set as follows. The HARDI data were modeled by a spherical harmonic expansion of order six to compute the ODF analytically [12]. The regularization parameter was set to $\gamma = 0.006$, as suggested in the paper. The ODF was calculated for $H = 321$ directions corresponding to the 7^{th} order, icosahedron-based tessellation of the unit sphere. The ODF was deconvolved into the ODF using $\epsilon = 0.0005$ for the regularization (equation 1). The white matter mask was calculated as $\gamma(\boldsymbol{r}) > \gamma_{min} = 1/3$. The step size used in Perrin et al.'s method was set to 0.5 mm, and 10^6 particles were sampled. Regarding the maximum angle, $ANGLE\_MAX = 30^\circ$ was used for Perrin et al.'s algorithm, as suggested in their paper. For our algorithm, it is convenient to allow a larger angle for the larger step size. $ANGLE\_MAX = 45^\circ$ was found to provide good results in preliminary experiments. Finally, the kernel width for the filter which smooths the tractogram $\sigma_s = 1.5$mm was also determined from preliminary experiments.

4.2. Results: comparison with Perrin et al.'s method

The scatter plot in Figure 1 compares the connectivity values from the two methods for the ten images and the four tracts. The outputs are quite well correlated ($\rho = 0.79$), indicating that our method is a good approximation of the stochastic method it is based on, except when a high precision tractogram is required. Even if that is the case, our method can still be used to quickly find a good initialization for the computationally expensive, sampling-based stochastic tractography. A good seed could be interactively localized with our method and then used as input for the probabilistic algorithm. This way the stochastic tractography will be assuredly successful, as opposed possibly having to modify the seed location and rerun the slow method several times. Figure 2 depicts the renderings of the seeds and corresponding target ROIs for the genu and the splenium of the corpus callosum in one of the brains, as well as the output tractogram.

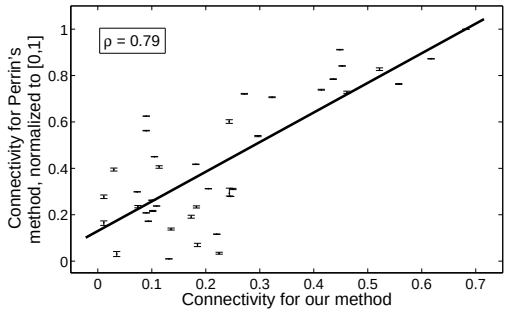

Fig. 1. Scatter plot of the connectivity measures provided by Perrin et al.'s method (executed three times per case, hence the ranges displayed in the figure rather than points) and ours. The connectivity was normalized to the maximum value.

4.3. Results: execution speed

Execution speed is the most noteworthy feature of the method developed in this study. Both Perrin et al.'s method and ours were implemented in Java and tested on a desktop with an Intel i7 CPU. The algorithms were parallelized to take advantage of the multi-core structure of the computer. Table 2 lists the running times of the main steps of the algorithm for the two methods, averaged over ten executions on the same volume. Despite the overhead of calculating the transition matrix, our method is significantly faster than the stochastic approach. The difference between the two methods is especially

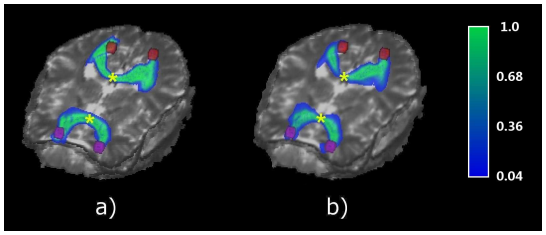

Fig. 2. Rendering of T2 baseline volume, seed (yellow asterisks) and target regions (colored spheres), and resulting tractogram. The intensity of the tractograms is normalized to [0,1] and thresholded at 0.04 for display. a) Genu (purple spheres) and splenium (red spheres) of corpus callosum for Perrin et al.'s method. b) Output for our method.

Calculation	Perrin et al. [5]	FAST
fODF	2.64 ± 0.13 s	2.69 ± 0.10 s
Transition tensor	N/A	44.2 ± 2.1 s
Tracking	428 ± 18 s	6.92 ± 0.28 s

Table 2. Execution times for our method and Perrin et al.'s.

large if the seed region is to be moved due to unsatisfactory results. In this case, recalculating the transition tensor, which is the most computationally expensive step of our algorithm, is unnecessary and the new output can be generated in a few seconds, while the stochastic sampling must be restarted over a period of approximately seven minutes. It is thus feasible to use our method in an interactive fashion, which is not possible with stochastic methods.

5. DISCUSSION

A computationally efficient probabilistic tractography method has been presented in this study. Though the method is not stochastic in nature, it produces a connectivity map from a HARDI volume in the spirit of stochastic tractography. By discretizing the fODF, the tractogram can be quickly calculated analytically. The method uses a bivector to store the position and speed of the probability mass stemming from a seed region. Each step in the random walked is then represented by a contraction operation with a transition tensor, which is designed to mimic the behavior of particle in a probabilistic method by Perrin et al. Our method was compared to an implementation of this algorithm, and the results demonstrate highly correlated outputs, with the difference that our method runs approximately 60 times faster. Once the transition tensor has been computed, interactive exploration of the brain connetivity is possible with our method because the tracking runs in approximately seven second: if the user is not satisfied with the tracking results, he can quickly refine them by shifting the seed region.

Most fiber bundles can be tracked by the proposed method in spite of discretization effects. However, tracking very thin bundles can be difficult if their orientation is not well approximated by any of the 98 allowed directions. In such cases, it would be possible to upscale the image only in the region where the user thinks the tractography might have failed. Generalizing the method to irregular grids is trivial thanks to the graph structure. In this case, reducing $ANGLE\_MAX$ in the ROI might be beneficial to compensate for the shorter step length. Exploring this possibility, as well as automatically detecting the problematic regions, remains as future work.

Finally, it is important to note that our method requires a large amount of memory to store the transition tensor. The Java implementation uses prior knowledge of the tensor structure to minimize memory usage, but still requires 11 GB. If the image size or resolution were smaller or more memory were available, then it would be possible to speed up the method by a factor of n by computing the tensor $\left(T^{p,u}_{q,v}\right)^n$ and then updating the bivector as: $Q^{(t)}_{q,v} = \left(\left(T^{p,u}_{q,v}\right)^n\right)^{(t/n)} Q^{(0)}_{p,u}$. However, this improvement comes at the expense of largely increasing the memory requirements because the sparsity of T decreases cubicly with n.

6. REFERENCES

[1] P.J. Basser, J. Mattiello, D. LeBihan, et al., "Estimation of the Effective Self-Diffusion Tensor from the NMR Spin Echo," *J. Magn. Reson., Ser. B*, vol. 103, pp. 247–247, 1994.

[2] C. Pierpaoli, P. Jezzard, PJ Basser, A. Barnett, and G. Di Chiro, "Diffusion tensor MR imaging of the human brain," *Radiology*, vol. 201, no. 3, pp. 637–648, 1996.

[3] DS Tuch, RM Weisskoff, JW Belliveau, and VJ Wedeen, "High angular resolution diffusion imaging of the human brain," in *ISMRM conf. proc.*, 1999, p. 321.

[4] D.S. Tuch, "Q-Ball Imaging," *Magn. Reson. Med.*, vol. 52, no. 6, pp. 1358–1372, 2004.

[5] M. Perrin, C. Poupon, Y. Cointepas, B. Rieul, N. Golestani, C. Pallier, D. Riviere, A. Constantinesco, D. Le Bihan, and JF Mangin, "Fiber tracking in q-ball fields using regularized particle trajectories," in *Lect. Notes Comput. Sci. (IPMI)*, 2005, vol. 3565, pp. 52–63.

[6] N. Fout, J. Huang, and Z. Ding, "Visualization of neuronal fiber connections from DT-MRI with global optimization," in *ACM Symp. Appl. Comput. conf. proc.*, 2005, pp. 1200–1206.

[7] D. Merhof, M. Richter, F. Enders, P. Hastreiter, O. Ganslandt, M. Buchfelder, C. Nimsky, and G. Greiner, "Fast and accurate connectivity analysis between functional regions based on dt-mri," in *Lect. Notes Comput. Sci. (MICCAI)*, 2006, vol. 4191, pp. 225–233.

[8] Y. Iturria-Medina, EJ Canales-Rodriguez, L. Melie-Garcia, PA Valdes-Hernandez, E. Martinez-Montes, Y. Aleman-Gomez, and JM Sanchez-Bornot, "Characterizing brain anatomical connections using diffusion weighted MRI and graph theory," *Neuroimage*, vol. 36, no. 3, pp. 645–660, 2007.

[9] S.N. Sotiropoulos, L. Bai, and C.R. Tench, "Fuzzy anatomical connectedness of the brain using single and multiple fibre orientations estimated from diffusion mri," Comput. Med. Imaging and Graphics (in press), 2009.

[10] S.M. Smith, "Fast robust automated brain extraction," *Hum. Brain Mapping*, vol. 17, no. 3, pp. 143–155, 2002.

[11] M. Descoteaux, R. Deriche, T.R. Knosche, and A. Anwander, "Deterministic and Probabilistic Tractography Based on Complex Fibre Orientation Distributions," *IEEE Trans. Med. Imaging*, vol. 28, no. 2, pp. 269–286, 2009.

[12] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Regularized, fast, and robust analytical Q-ball imaging," *Magn. Reson. Med.*, vol. 58, no. 3, pp. 497–510, 2007.

[13] B. Jian and BC Vemuri, "A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI," *IEEE Trans. Med. Imaging*, vol. 26, no. 11, pp. 1464–1471, 2007.