

Fast Approximate Stochastic Tractography

Juan Eugenio Iglesias · Paul M. Thompson ·
Cheng-Yi Liu · Zhuowen Tu

Published online: 8 April 2011
© Springer Science+Business Media, LLC 2011

Abstract Many different probabilistic tractography methods have been proposed in the literature to overcome the limitations of classical deterministic tractography: i) lack of quantitative connectivity information; and ii) robustness to noise, partial volume effects and selection of seed region. However, these methods rely on Monte Carlo sampling techniques that are computationally very demanding. This study presents an approximate stochastic tractography algorithm (FAST) that can be used interactively, as opposed to having to wait several minutes to obtain the output after marking a seed region. In FAST, tractography is formulated as a Markov chain that relies on a transition tensor. The tensor is designed to mimic the features of a well-known probabilistic tractography method based on a random walk model and Monte-Carlo sampling, but can also accommodate other propagation rules. Compared to the baseline algorithm, our method circumvents the sampling process and provides a deterministic solution at the expense of partially sacrificing sub-voxel accuracy. Therefore, the method is strictly speaking not stochastic, but provides a probabilistic output in the spirit of stochastic tractography methods. FAST was compared with the random walk model using real data from 10 patients in two different ways: 1. the probability maps produced by the two methods on five well-known fiber tracts were directly compared using metrics from the image registration literature; and 2. the connectivity measurements between different regions of the brain given by the two methods were compared using the correlation coefficient ρ . The results show that the connectivity measures provided by the two algorithms are

well-correlated ($\rho=0.83$), and so are the probability maps (normalized cross correlation 0.818 ± 0.081). The maps are also qualitatively (i.e. visually) very similar. The proposed method achieves a 60x speed-up (7 s vs. 7 min) over the Monte Carlo sampling scheme, therefore enabling interactive probabilistic tractography: the user can quickly modify the seed region if he is not satisfied with the output without having to wait on average 7 min.

Keywords Probabilistic tractography · Diffusion-weighted MRI · Markov chain

Introduction and Background

Diffusion weighted MRI (DW-MRI) is the only way of imaging axonal fiber tracts in the brain in vivo. DW-MRI consists of a reference T2-MRI scan and a number of diffusion-weighted scans, each with a sensitizing magnetic field gradient in a different direction. Comparing the intensities with and without the gradient, the water diffusion for each probed direction can be estimated. DW-MRI makes it thus possible to reconstruct a full diffusion profile at each location in the brain. From these profiles, fiber tracts can be inferred using a tractography method. The reconstructed fiber tracts are widely used in research to study the functionality of the different parts of the brain and to correlate anomalies in the fiber tracts with diverse neurological disorders. In clinical environments, it is used in surgical planning in order to avoid damaging fiber tracts with important functionality.

There are different ways of reconstructing a continuous diffusion profile at each location from a set of discrete measurements. In diffusion tensor imaging (DTI, Basser et al. 1994; Pierpaoli et al. 1996), an ellipsoid is fitted to a

J. E. Iglesias (✉) · P. M. Thompson · C.-Y. Liu · Z. Tu
UCLA,
Los Angeles, CA, USA
e-mail: jeiglesias@ucla.edu

small number of diffusivity measurements (as few as six). To resolve complex fiber geometries (e.g. crossings), it is necessary to acquire a higher number of directions (typically 50 or more). This is known as high angular resolution diffusion imaging (HARDI (Tuch et al. 2002; Frank 2001; Ozarslan and Mareci 2003)). Mathematical entities more complicated than an ellipsoid, such as fourth-order tensors or expansions in spherical harmonics, can then be fitted to the data (Barmpoutis et al. 2007; Ghosh et al. 2008; Frank 2002; Alexander et al. 2002; Alexander et al. 2001; Behrens et al. 2003).

Fiber tracking cannot be performed directly on the diffusion MRI data because, in presence of more than one fiber population, the directions of maximal signal attenuation by diffusion do not correspond to the fiber orientations. For example, in a 90° crossing of two fibers, the maximum attenuation occurs at a 45° angle. Instead, the orientation distribution function (ODF, Tuch 2004) can be used. Given the 3D probability density function (PDF) of water diffusivity i.e. the probability that water diffuses a certain distance in a certain direction, the ODF is obtained as the radial projection of this function. Therefore, it represents how likely it is to observe any diffusion along a given direction, and is defined on a spherical shell. The ODF accounts for the uncertainty in water diffusion and its maxima are assumed to correspond to the orientation of the fibers.

There are two families of tractography methods: deterministic and probabilistic/stochastic. In classic deterministic streamline tractography (Basser et al. 2000; Lazar et al. 2003; Parker and Alexander 2003; Kreher et al. 2005; Wedeen et al. 2008; Descoteaux et al. 2009; Qazi et al. 2009), the field of most likely fiber orientations as given by the ODFs is followed from a seed voxel or region to obtain the trajectory of a fiber tract. The tracking stops when a point with low diffusion anisotropy (i.e. not corresponding to fibers) is reached or when a high-curvature trajectory is required to follow the field. It is possible to regularize the trajectory of the track by imposing constraints on its curvature or by using deflection i.e. combining the local ODF information and the direction from the previous tracking step to calculate the new direction of propagation, rather than following the direction of maximum diffusion directly. For example, in DTI this can be done by multiplying the incoming direction by the local tensor (Lazar et al. 2003). While regularization can be useful in certain scenarios (e.g. to follow relatively straight tracts that go through fiber crossings, especially in DTI), they have the disadvantage that they can lose track of highly curved fiber tracts.

On the other hand, stochastic tractography accounts for the uncertainty of the fiber orientation estimates as captured by the ODF. Rather than using just the directions

corresponding to the most likely orientations, the whole ODF is used in a probabilistic framework. These methods achieve higher robustness to noise, partial volume effects and selection of seed region. Moreover, they have the advantage that they provide quantitative connectivity information between different brain regions. The main idea is to sample the ODF field at each step to decide the direction of the next move (Behrens et al. 2003; Parker and Alexander 2003; Descoteaux et al. 2009; Koch et al. 2002; Friman et al. 2006; Behrens et al. 2007; Perrin et al. 2005; Kaden et al. 2007; Jbabdi et al. 2007; Parker and Alexander 2005), as opposed to updating the direction of the tract deterministically as in streamline tractography. The sampling can be limited to a certain angle within the incident direction to regularize the trajectory. Descoteaux et al. 2009 suggest that the fiber ODF (fODF) should be used instead of the ODF in stochastic tractography to improve the results. The fODF is a PDF of the orientation of the fibers in a given location, and it can be estimated directly from the diffusion signal (Tournier et al. 2004) or from the ODF (Descoteaux et al. 2009) using deconvolution. The fODF is sharper than the ODF; ideally it would be zero everywhere except for the exact directions in which the voxel is crossed by fibers.

If a large amount of tracking instances (particles) are generated from a seed region in a Monte-Carlo scheme, a connection probability map can be obtained by computing the fraction of particles that reaches each voxel in the volume. This map provides a measure of the probability that the voxel is connected to the seed. The computational load is much higher than for deterministic methods, which typically run in less than a second. A high-quality connection probability map requires sampling a large number of particles, up to 10^5 – 10^6 , depending on the seed region size, desired level of granularity in the map, noise level and distance to target region (the longer the tract and the noisier the data, the more particles are lost on the way). This makes interactive use very impractical: the user cannot quickly modify the position of the seed and re-run the algorithm if he is not satisfied with the obtained map.

This paper extends our previous work (Iglesias et al. 2010) on “fast approximate stochastic tractography” (FAST), an approach that can quickly compute the connection probability map analytically at the expense of limiting the directional resolution of the path propagation. Since the algorithm is deterministic, the method is strictly speaking not stochastic, but we call it so because it provides a quantitative estimate of the connectivity in the spirit of stochastic tractography algorithms. The diffusion from a seed region is modeled as a discrete Markov process in which each variable corresponds to a position/direction pair. The probability mass shifts from a seed region to the rest of the brain following a Markov process that is

governed by a transition tensor. This circumvents the time-consuming stochastic Monte-Carlo sampling. The probabilistic connectivity map for any seed region can then be calculated in a few seconds, which enables interactive exploration of the connectivity of the brain.

There are a number of studies in which the movement of particles is discretized to the orientations corresponding to neighboring voxels. The data are seen as a graph in which the voxels are nodes and neighboring nodes are connected by edges. Then, if the user defines a seed and a target region, the strength of the connection can be estimated as the probability of the most likely path between them. The path with minimum cost can be found using Dijkstra-like algorithms. For example, Fout et al. (2005) used a 26-neighborhood to construct the graph and dynamic programming for finding the most likely path. Merhof et al. (2006) used a hexahedral grid (each voxel has 74 neighbors) and A* for the optimization. Iturria-Medina et al. (2007), Lifshits et al. (2009), Sotiropoulos et al. (2010a) and Zalesky (2008) used different variants of Dijkstra's algorithm and neighborhood sizes. Despite being designed to compute the connectivity between two regions, some of these methods (Iturria-Medina et al. 2007; Sotiropoulos et al. 2010a; Zalesky 2008; Lifshits et al. 2009) can be modified to provide a full brain connection probability map from a seed region without the need of sampling trajectories (see also Sotiropoulos et al. (2010b)).

Front propagation techniques also use discrete methods to explicitly generate connection probability maps without sampling. Parker et al. (2002) evolve a front from a seed point in the field of main DTI eigenvectors using a variant of the fast marching method (Sethian 2000). Fast marching is a particular case of the level set method (Osher and Fedwick 2003) in which the sign of the speed function never changes and that can be solved very efficiently. Campbell et al. 2005 use a flow-based surface evolution algorithm very similar to Parker et al. 2002, but making use of the full diffusion profile. Staempfli et al. 2006 proposed a speed function designed to maintain the front evolution direction in crossings and bifurcations.

While most of the discussed discrete propagation methods can generate a full connection probability map, they rely on very simple models of particle propagation which tend to produce excessive branching (Sotiropoulos et al. 2010b). The two main contributions of this study are: 1. a second-order discretization of the particle propagation model which, using a Markovian formulation, is able to accommodate all the elements of sampling-based stochastic tractography, including spatial regularization during tracking in a high angular resolution ODF field; and 2. a proper quantitative evaluation to assess the impact of discrete path propagation. We designed the tensor to mimic a well-known stochastic tractography method (Perrin et al. 2005,

2008, we chose this algorithm because it is general and includes a particle deflection term). This makes it possible to compare the outputs from this algorithm and FAST in order to quantitatively evaluate the impact of the made approximations. Compared with our previous short conference article (Iglesias et al. 2010), this paper includes a more detailed explanation of the methods, experiments with synthetic data, and more experiments with real patient data.

The rest of the paper is organized as follows. “Methods” describes the HARDI data used in this study, as well as the methods to estimate the ODF, deconvolve it into the fODF, and perform tractography. “Experiments and Results” describes experiments on synthetic and real data. Finally, “Discussion” includes conclusions and future work.

Methods

Dataset

Real Data

Ten HARDI volumes from ten different subjects were used in this study. The data were acquired with a 4 Tesla Bruker Medspec scanner. Each volume included 94 diffusion-sensitized gradient directions and 11 baseline images with no diffusion-sensitization (i.e. T2-weighted). The voxel size was 1.8 mm×1.8 mm×2.0 mm, the image size 128×128×55 voxels, and the acquisition parameters TE/TR=92.3/8,259 ms, b-value=1159 s/mm². The acquisition time was approximately 15 min. The 11 baseline images were merged into a single estimate of the T2 reference (Aja-Fernandez et al. 2007). This reference was used to calculate a mask corresponding to the brain using the BET algorithm (Smith 2002). The mask was then applied to all the diffusion images. In the rest of the paper, we use the terms “scan” and “volume” interchangeably to refer to the HARDI data from a subject.

Synthetic Data

Synthetic data simulating crossings and bifurcations at 1 mm×1 mm×1 mm resolution were generated for preliminary experiments. For areas with one or no fiber populations, the attenuation by diffusion is modeled as (Basser 1995):

$$S(\mathbf{r}, \boldsymbol{\beta}) = S_0(\mathbf{r})e^{-b\boldsymbol{\beta}'P(\mathbf{r})\Lambda(\mathbf{r})P(\mathbf{r})'\boldsymbol{\beta}} \quad (1)$$

where \mathbf{r} is the position vector, $\boldsymbol{\beta}$ is the probed diffusion direction and $P=[\mathbf{e}_1|\mathbf{e}_2|\mathbf{e}_3]$, $\Lambda=diag(\lambda_1, \lambda_2, \lambda_3)$ are the matrices with the orthonormal eigenvectors and the eigen-

values of the diffusion tensor (ellipsoid) respectively. For single fiber populations, \mathbf{e}_1 is oriented along the fiber direction and $\lambda_1 \gg \lambda_2 = \lambda_3$. Outside the simulated fibers, isotropic diffusion is assumed: the direction of \mathbf{e}_1 is irrelevant and $\lambda_1 = \lambda_2 = \lambda_3$.

In crossing and bifurcation regions, the following two-tensor model was used (Basser 1995):

$$S(\mathbf{r}, \boldsymbol{\beta}) = S_0(\mathbf{r})w(\mathbf{r})e^{-b\boldsymbol{\beta}'P_1(\mathbf{r})\Lambda_1(\mathbf{r})P_1(\mathbf{r})'\boldsymbol{\beta}} + S_0(\mathbf{r})[1-w(\mathbf{r})]e^{-b\boldsymbol{\beta}'P_2(\mathbf{r})\Lambda_2(\mathbf{r})P_2(\mathbf{r})'\boldsymbol{\beta}} \quad (2)$$

where the weights w and $1-w$ describe the contributions of each fiber population to the MRI signal. In this study, the following values are used: $b=1200\text{s/mm}^2$, $[\lambda_1, \lambda_2, \lambda_3]^t = [2.0, 0.2, 0.2]^t 10^{-3}\text{mm}^2/\text{s}$ for the fibers (these are typical values for the corpus callosum), $\lambda_1 = \lambda_2 = \lambda_3 = 2.5 \cdot 10^{-3}\text{mm}^2/\text{s}$ outside the fibers, and $w=1-w=0.5$ in crossings and bifurcations. Eleven reference T2 volumes and $D=94$ gradient images were generated using the same directions that were available for the real data. Rician noise was added to simulate a signal-to-noise ratio $SNR=20$.

Calculation of the fODF

It has been shown that using the fODF rather than the smoother ODF in stochastic tractography can drastically improve the tracking results (Descoteaux et al. 2009). If $S_i(\mathbf{r})$ is the HARDI data for direction i (θ_i, φ_i) at location \mathbf{r} , we define the signal vector $\mathbf{g}(\mathbf{r}) = [S_1(\mathbf{r}), S_2(\mathbf{r}), \dots, S_D(\mathbf{r})]^t$, where D is the number of directions of the data. Then the ODF can be reconstructed at an arbitrary set of H directions as:

$$\boldsymbol{\psi}(\mathbf{r}) = \frac{1}{Z(\mathbf{r})} A_3 A_2 A_1 \mathbf{g}(\mathbf{r}) = \frac{1}{Z(\mathbf{r})} A \mathbf{g}(\mathbf{r}) \quad (3)$$

where A is a matrix that summarizes the linear operations that estimate the ODF from the signal vector $\mathbf{g}(\mathbf{r})$ (Descoteaux et al. 2007) and $Z(\mathbf{r})$ just ensures that $\boldsymbol{\psi}(\mathbf{r})$ is a PDF i.e. $\mathbf{1}^t \boldsymbol{\psi}(\mathbf{r}) = \mathbf{1}$. $A = A_3 A_2 A_1$ is the product of three matrices. A_1 estimates the spherical harmonic (SH) coefficients of the signal vector. Rather than using direct least squares estimation, a Laplace-Beltrami regularization scheme (Descoteaux et al. 2007) is used. A_2 converts the SH coefficients of the signal vector to the SH coefficients of the ODF using the Funk-Radon transform (again, see Descoteaux et al. 2007). A_3 simply evaluates the SH expansion of the ODF at the H directions at which it is to be reconstructed.

In this study, we assume a blurring model in which the ODF is the convolution of the underlying fODF with a kernel:

$$\boldsymbol{\psi}(\mathbf{r}) \propto K \boldsymbol{\psi}_f(\mathbf{r}) \quad (4)$$

If the mixing matrix K is known, the fODF $\boldsymbol{\psi}_f$ can be recovered with a simple matrix inversion. K is the result of horizontally concatenating the responses to a single fiber orientation for each of the H directions where the ODF is reconstructed. The response for a single orientation can be estimated from the data by averaging the ODFs of the most anisotropic voxels, which must correspond to single fiber populations. This response can then be aligned to each of the H reconstruction directions through rotations in order to obtain the H columns of K .

Finally, once the mixing matrix is known, the fODF can be recovered by premultiplying the ODF by the Tikhonov-regularized inverse of K :

$$\begin{aligned} \boldsymbol{\psi}_f(\mathbf{r}) &\propto \left[(K^t K + \varepsilon I)^{-1} K^t \boldsymbol{\psi}(\mathbf{r}) \right]_+ \\ &\propto \left[(K^t K + \varepsilon I)^{-1} K^t A \mathbf{g}(\mathbf{r}) \right]_+ \propto [A' \mathbf{g}(\mathbf{r})]_+ \\ &\Rightarrow \boldsymbol{\psi}_f(\mathbf{r}) \frac{1}{Z'(\mathbf{r})} [A' \mathbf{g}(\mathbf{r})]_+ \end{aligned} \quad (5)$$

In Eq. 5, $A' = (K^t K + \varepsilon I)^{-1} K^t A$ is the final signal-to-fODF reconstruction matrix. The parameter ε is the regularizer for the fODF estimation. The operator $[\cdot]_+$ denotes that the negative elements of the vector are set to zero, in case the deconvolution produced negative coefficients due to insufficient regularization. Finally, $Z'(\mathbf{r}) = \mathbf{1}^t [A' \mathbf{e}(\mathbf{r})]_+$ is the updated partition function that ensures $\mathbf{1}^t \boldsymbol{\psi}_f(\mathbf{r}) = 1$ after the deconvolution. There are more sophisticated deconvolution strategies in the literature that better deal with the non-negativity constrain, but they are much slower, they cannot be directly integrated into the matrix framework and, as shown in (Jian and Vemuri 2007), their improvement with respect the strategy used in the study is only marginal.

The State Bivector and the Transition Tensor

In FAST, the direction/velocity pairs for a particle are designed to mimic a well-known tractography method (described in [Baseline Sampling-Based Method](#)), but limiting the particles to stay on grid locations. Thanks to this discretization, the sampling process can be replaced by a bivector-valued Markov chain to compute the connection probability map analytically. The time-dependent state bivector ([State Bivector](#)) stores the probability that a particle would be located in a voxel at a given time instant and moving in a certain direction. The bivector evolves through time according to a transition tensor ([Transition Tensor](#)) which simulates the movement of particles from the seed. Once all the probability mass has exited the brain white matter, the connection probability map can be approximated by marginalizing the state bivector along

the time and directional dimensions (“Relationship with the Connection Probability Map”).

Baseline Sampling-Based Method

In this study, Perrin et al.’s algorithm (Perrin et al. 2005; Perrin et al. 2008) is used as a representative example of probabilistic tractography algorithm. In their method, particles follow a regularized random walk according to the following rule. At each discrete step, the particle follows the direction:

$$\mathbf{d}^{(t)} = \gamma(\mathbf{r})\mathbf{d}_{fODF}(\mathbf{r}) + [1 - \gamma(\mathbf{r})]\mathbf{d}^{(t-1)} \tag{6}$$

where $\mathbf{d}_{fODF}(\mathbf{r})$ is the result from sampling the fODF at location \mathbf{r} restricted to a 30° half cone defined by the incident direction $\mathbf{d}^{(t-1)}$. The scalar function $\gamma(\mathbf{r})$ corresponds to the standard deviation of the fODF at each voxel normalized to its maximum across the image, which is a measure of the anisotropy of the diffusion. Therefore, the trajectory is regularized because the direction of the previous step is considered for: 1. giving it a weighted contribution in $\mathbf{d}^{(t)}$; and 2. using it to constrain the sampling of the fODF. Particles keep traveling across the volume until they exit the brain white matter, which is estimated as the set of voxels for which $\gamma(\mathbf{r})$ remains above a certain threshold γ_{min} .

The connection probability map is computed by letting a large number of particles follow the random walk from the seed and assigning to each voxel in the volume the fraction of particles that has reached it. This procedure is slow for two reasons. First, the number of particles that must be sampled to obtain a smooth, reliable map can be large, depending on the distance between the seed and target regions and their sizes. Second, following the trajectory of the particles requires interpolating the HARDI data at non-grid locations. Ideally, the interpolation should be carried out in the corresponding Riemannian manifold (Cheng et al. 2009), which is very slow. Here we use trilinear interpolation instead; this is a common choice in the literature thanks to its simplicity and speed, despite being less exact.

State Bivector

In our discrete formulation, no particle sampling is necessary because the state bivector stores the continuous fraction of mass (equivalent to the fraction of particles in “Baseline sampling-based method”) at each location, time and direction of traveling. Specifically, we define the bivector $Q = Q_{p,u}^{(t)}$ as the probability mass at spatial location \mathbf{r}_p with direction u for discrete time t (Fig. 1a). Compared with the sampling-based method, p and u play the role of \mathbf{r}

and $\mathbf{d}^{(t-1)}$ respectively in Eq. 6. The first index ranges from $p=1$ to $P=N$, the number of voxels in the volume, while the second index ranges from $u=1$ to $u=98$, corresponding to the directions (ignoring colinearities) that connect a voxel with the 124 neighbors in the 3-D lattice corresponding to a $5 \times 5 \times 5$ cube (see Fig. 1b). The 98-neighborhood provides a good compromise between: i) angular resolution: the larger the neighborhood, the better the resolution (98 directions here); ii) computational complexity: the larger the neighborhood, the larger the memory requirements; and iii) step length variability: the larger the neighborhood, the more variable the step length is (between 1 and 3 voxels in our case). Large step variabilities bias the propagation in certain directions. Moreover, large step lengths can lead to artificial jumps over structures outside the white matter mask. For example, in Perrin et al. 2008, the probability mass could jump between neighboring gyri.

The state bivector is illustrated in the upper row of Fig. 2, which shows its distribution in space at different time points without considering the direction of movement i.e. $\sum_{u=1}^{98} Q_{p,u}^{(t)}$. These images are the equivalent of “freezing” the particles in sampling-based methods, therefore ignoring the direction u .

Transition Tensor

The evolution of the state bivector is determined by T , a tensor of type (2, 2). $T_{q,v}^{p,u}$ represents the transition probability from position-direction (p, u) to (q, v) . T is very large (dimensionality $N \times 98 \times N \times 98$) but extremely sparse. The transition from step t to step $t+1$ is mathematically a tensor contraction:

$$Q_{q,v}^{(t+1)} = \sum_{p=1}^N \sum_{u=1}^{98} T_{q,v}^{p,u} Q_{p,u}^{(t)} \tag{7}$$

The elements of the transition tensor T can be designed to mimic the behavior of a particle in the baseline method as accurately as possible. First, all the elements that do not satisfy all of the following four conditions must be set to zero:

1. The origin voxel p must be in the brain white matter. A mask that defines the white matter is calculated by thresholding $\gamma(\mathbf{r})$; the output is refined by applying a morphological closing operator with a spherical kernel in order to smooth it and remove small holes (see Fig. 1c–d).
2. Voxels p and q must be 98-neighbors in the 3D lattice to avoid jumps.
3. Direction v has to be equal (parallel) to the vector $\mathbf{r}_q - \mathbf{r}_p$. This is the way of “memorizing” the direction of the previous movement of the probability mass.

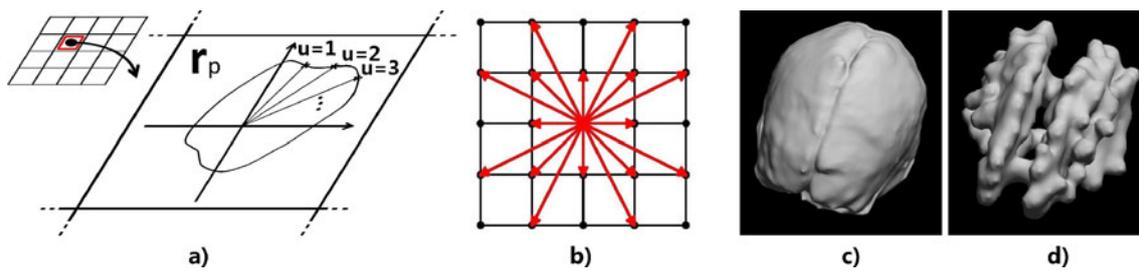


Fig. 1 **a** State bivector at a given time step: at each location \mathbf{r}_p , each orientation $u=1,2,\dots,98$ has a probability that a particle from the seed is traveling in that direction. **b** Unique directions connecting a pixel with its neighbors in a 5×5 square. In three dimensions, there are 98

neighbors defined by unique directions out of the 124 in a $5\times 5\times 5$ cube. **c** Brain extraction provided by BET. **d** White matter mask obtained by thresholding $\gamma(\mathbf{r})$ and refining the output with a closing operator

4. Direction u is deflected into v by a vector \mathbf{d}_{fODF} (see Eq. 6) such that u and \mathbf{d}_{fODF} form an angle under a certain threshold (30° in Perrin et al.'s study). This regularizes the trajectory of the probability mass.

Then, if a transition is compatible with these conditions, it is assigned a probability proportional to the sum of the values of $\psi_f(\mathbf{r}_p)$ in the directions \mathbf{d}_{fODF} that deflect direction u into direction v according to Eq. 6 i.e. the directions that, when substituted in the equation, yield a direction that is closer to v than to any other direction of the 98. To ensure that no probability mass is lost in the Markov chain, the probabilities are scaled so that $\sum_q \sum_v T_{q,v}^{p,u} = 1$, unless the sum is zero, meaning that the voxel is a dead end i.e. it is out of the white matter mask. In that case, all transitions from the voxel are equal to zero, simulating the end of the trajectory of a particle. The algorithm to calculate the elements of T is detailed in Table 1.

Compared with the sampling-based method, u and $v \leftrightarrow \mathbf{r}_q - \mathbf{r}_p$ play the roles of $\mathbf{d}^{(t-1)}$ and $\mathbf{d}^{(t)}$ respectively in Eq. 6. The tensor $T_{q,v}^{p,u}$ controls the evolution of the state bivector the same way as $\mathbf{d}_{fODF}(\mathbf{r})$ controls the movement of the particles in Eq. 6 (see upper row of Fig. 2 for an example of the evolution of the bivector according to the tensor).

Relationship with the Connection Probability Map

The main advantage of FAST is that, once the transition tensor T has been calculated, the connection probability map for any seed voxel or region can be generated very quickly. Given a seed region, one just builds the seed vector $Q_{p,u}^{(0)}$ as:

$$Q_{p,u}^{(0)} = \begin{cases} \frac{1}{|S|} \sum_{h=1}^H \psi_{f,h}(\mathbf{r}_p) \delta[\Upsilon(\theta_h, \varphi_h) - u], & \mathbf{r}_p \in S \\ 0, & \mathbf{r}_p \notin S \end{cases} \quad (8)$$

where S is the set of seed voxels, $\delta[\cdot]$ denotes the Kronecker delta and $\Upsilon(\theta, \varphi)$ is a function that discretizes an arbitrary angle to the 98-neighborhood (see caption of Table 1). In other words, we are spreading the probability mass evenly across the voxels in the seed region ($1/|S|$) and, within each voxel, the distribution of the probability mass is proportional to the fODF. Then the distribution at time t is:

$$Q_{q,v}^{(t)} = \left(T_{q,v}^{p,u} \right)^{(t)} Q_{p,u}^{(0)} \quad (9)$$

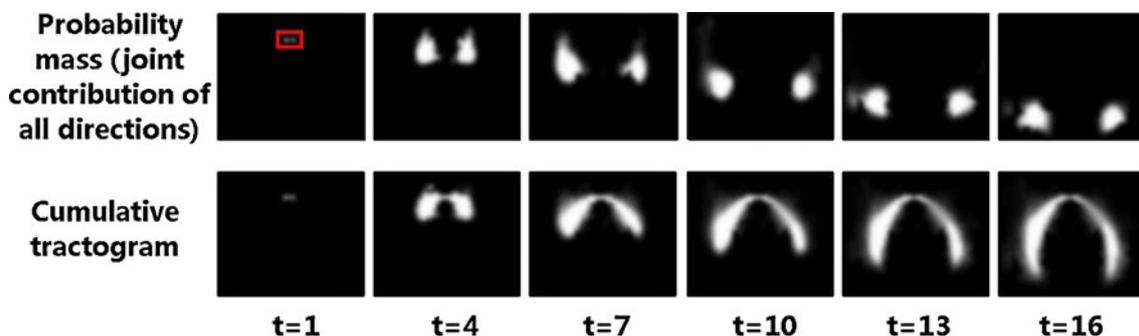


Fig. 2 Evolution of the probability mass and the state bivector throughout the tractography, shown for an axial slice. The seed region is marked in red. Upper row: probability mass resulting from marginalizing the state bivector with respect to the directional variable $\sum_{u=1}^{98} Q(t)_{p,u}$. This is equivalent to freezing the particles in

a sampling-based method. Lower row: cumulative connection probability map i.e. how they would look like if tractography had been terminated early at time t . They are generated as $\sum_{u=1}^{98} \sum_{t'=0}^t Q(t')_{p,u}$. The images are log-transformed for easier visualization

Table 1 Pseudocode for calculating the transition tensor. $\gamma(\mathbf{r})$ corresponds to the standard deviation of the fODF at each voxel normalized to its maximum across the image. $\Upsilon(\theta, \varphi)$ 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

```

· Calculate and normalize  $\gamma(\mathbf{r})$  the standard deviation of the fODF:  $\gamma(\mathbf{r}) \leftarrow \gamma(\mathbf{r}) / \max_{\mathbf{r}}[\gamma(\mathbf{r})]$ 
· Calculate white matter mask as  $\gamma(\mathbf{r}) > \gamma_{min}$  and refine it with a morphological closing operator.
· Initialize all elements of  $T_{q,v}^{p,u}$  to zero
· FOR each position/direction pair  $(p, u)$  such that  $\mathbf{r}_p$  belongs to the white matter mask
· SET PROBS[v] ← 0, for  $v=1, \dots, 98$ 
· SET SUM ← 0
· FOR each direction of the fODF  $\mathbf{d}_h, h=1, \dots, H$ 
· IF ANGLE( $\mathbf{d}_u, \mathbf{d}_h$ ) < ANGLE_MAX
· SET CDIR ←  $\gamma(\mathbf{r}_p)\mathbf{d}_h + (1-\gamma(\mathbf{r}_p))\mathbf{d}_u$ 
· SET PROBS[ $\Upsilon(\text{CDIR})$ ] +=  $\psi_h(\mathbf{r})$ 
· SET SUM +=  $\psi_h(\mathbf{r})$ 
END
· IF SUM > 0
· FOR each direction  $v=1, \dots, 98$ 
· SET  $q \leftarrow p+v$  (point+vector=point)
· SET  $T_{q,v}^{p,u} \leftarrow \text{PROBS}[v]/\text{SUM}$ 
END; END; END; END
    
```

where $(T_{q,v}^{p,u})^{(t)}$ denotes successive application of the contraction operator t times. If the bivector is unfolded into a vector and the transition tensor into a matrix, the contraction can be implemented as a matrix product; this has the advantage that one of the many implementations of sparse matrices which are publicly available can be used (an in-house Java implementation was used in this study).

Each contraction is equivalent to allowing all the particles to take one step simultaneously in the baseline method, so the sampling process is circumvented. The contraction must be applied a sufficient number of times t_{max} so that all the probability mass exits the volume, i.e. until the state bivector is approximately zero. Then, the connection probability map, which is the probability mass that has visited each voxel (fraction of particles in the baseline method) can be approximated by the sum of the state bivector across time and direction:

$$M'(\mathbf{r}_p) = \sum_{u=1}^{98} \sum_{t=0}^{t_{max}} Q(t)_{p,u} \tag{10}$$

The state bivector must be marginalized with respect to both direction and time because, in the baseline method, neither the directions in which particles are traveling nor their time of arrival is considered when computing the connection probability map. When summing across time, we are not considering that some probability mass can be counted twice if there are any loops. However, this effect is negligible because the white matter is well structured. The lower row of Fig. 2 illustrates the process of formation of the connection probability maps by showing how the map would look like if the tractography was stopped early at different time points.

Finally, because a $5 \times 5 \times 5$ 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 $M'(\mathbf{r})$ with a Gaussian kernel. For visualization purposes, the logarithm of the result is taken to compress its range and the output min-max normalized to obtain values between zero and one:

$$M(\mathbf{r}) = \frac{\log[M'(\mathbf{r}) * G(\mathbf{r}, \sigma_s)] - M_{min}}{M_{max} - M_{min}} \tag{11}$$

where $G(\mathbf{r}, \sigma_s)$ is a Gaussian smoothing kernel with standard deviation σ_s and M_{min}/M_{max} are the minimum and maximum of the smoothed connection probability after taking the logarithm.

The tensor contraction runs sufficiently fast so that the user can explore the connectivity of different regions of the brain interactively. Once a volume is loaded in memory and the transition tensor has been precomputed, which is done in a relatively short time (under a minute), the T2 baseline or the anisotropy image $\gamma(\mathbf{r})$ can be displayed on the screen so that the user can click on different points of the volume and get their connectivity maps very quickly.

Experiments and Results

Experimental Setup

First, the synthetic data was used in preliminary experiments to tune the system parameters. Then, the real data was used for formal evaluation through two experimental studies. In the first experiment, the connection probability

maps are compared directly as a whole using well-known similarity metrics from the image registration literature: the mean square error and the normalized cross-correlation (see Table 2 for the definitions). Since the goal of stochastic tractography is to assess the connectivity between regions in the brain, a second experiment comparing the connectivities was set up as follows. First, seeds were placed on well-known tracts in the brain. Then, target regions of interest (ROI) which are known to be connected to the seeds were defined. Finally, the connectivities from both methods were obtained by summing the connection probability maps within the target ROIs and then compared.

Five well-known tracts were used in the experiments. The target ROI was in all cases a 10 mm diameter sphere. The seed/ROI pairs were the following:

1. Seed in the genu of the corpus callosum, two target ROIs (left/right) in the frontal lobe near the lower part of the superior front gyrus.
2. Seed in the splenium of the corpus callosum, two target ROIs (left/right) in the occipital lobe, 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.
4. Seed in the central region of the cingulate gyrus, above the corpus callosum, and target ROI near the parahippocampal gyrus.
5. Seed in the central region of the corpus callosum, two target ROIs (left/right) at the superior region of the frontal cortex near the central fissure.

The system parameters were set as follows. The HARDI data were modeled by a spherical harmonic expansion of order six to compute the signal to ODF matrix (Descoteaux et al. 2007), setting the Laplace-Beltrami regularizer to 0.006, as suggested in that study (this is reflected in matrix A_1 in Eq. 3). The ODF was calculated for $H=321$ directions corresponding to the 7th order, icosahedron-based tessellation of the unit sphere. The ODF was deconvolved into the fODF using $\varepsilon=0.0005$ for the Tikhonov regularization (Eq. 5). The white matter mask was calculated as $\gamma(\mathbf{r}) > \gamma_{min} = 1/3$. The step size used in Perrin et al.'s method was made equal to 0.5 mm, and the number of particles was set to 5×10^4 for the synthetic data and 10^6 for the real images.

Table 2 Metrics used for directly comparing tractograms. A_i and B_i are the i^{th} voxels of the volumes to compare A and B respectively

Metric	Definition
Mean squares	$MS = \frac{1}{N} \sum_{i=1}^N (A_i - B_i)^2$
Normalized cross-correlation	$NCC = \frac{\sum_{i=1}^N A_i B_i}{\sqrt{\sum_{i=1}^N A_i^2 \sum_{i=1}^N B_i^2}}$

Regarding the maximum angle, $ANGLE_MAX=30^\circ$ was used for Perrin et al.'s algorithm, as suggested in their paper. However, it is convenient to allow a larger angle in our method in order to compensate for the larger step size. This parameter was determined from the experiments using synthetic data (see "Results on Synthetic Data"). Finally, For the mixing matrix K in Eqs. 4 and 5, the 10,000 voxels with highest anisotropy were extracted from the data, their ODFs computed, their directions of maximal diffusion aligned with the z axis and the average taken.

Results on Synthetic Data

The synthetic data was used to tune the maximum turning angle $ANGLE_MAX$ and the width of the smoothing kernel σ_s (Eq. 11). Several values of the parameters were tested, and $ANGLE_MAX = 45^\circ$ and $\sigma_s = 1.5$ mm were selected based on similarity with the output from Perrin et al.'s method and knowledge of ground truth.

Tracking results with these parameter values are shown in Fig. 3 for crossings at different angles and a smooth bifurcation. The connectivity maps are stretched to the interval [0,1], thresholded at 0.1 and log-transformed for easier visualization. Even though the maps are already normalized fractions, the ratio of probability mass that reaches the target ROI is usually small and normalizing the scores makes comparisons easier. Thresholding eliminates artifacts of the random walk that hinder the visualization. The middle row shows the output provided by Perrin et al.'s method. When the angle is 30° , the detected tract bifurcates into the crossing fiber in the orientation that is closest to the direction of movement of the particles traveling from the seed, while it suffers from minimal leaking in the opposite direction. At 60° and 90° , there is also some leaking along the crossing fiber.

The lower row in the figure displays the results from FAST. Compared with the output from the stochastic tractography, FAST provides very similar results for both the bifurcation and the crossings. At the tuned value of $ANGLE_MAX$, even the leakage patterns produced by FAST are very similar to those from the output of the method by Perrin et al.

Results on Real Data

Direct Comparison of Connection Probability Maps

In a first experiment, the connection probability maps from FAST and Perrin et al.'s method are compared directly using the metrics from Table 2. The metrics were computed on the unthresholded maps. The average results across the ten subjects and the standard deviations are displayed in Table 3. The highest similarity is achieved when the seeded

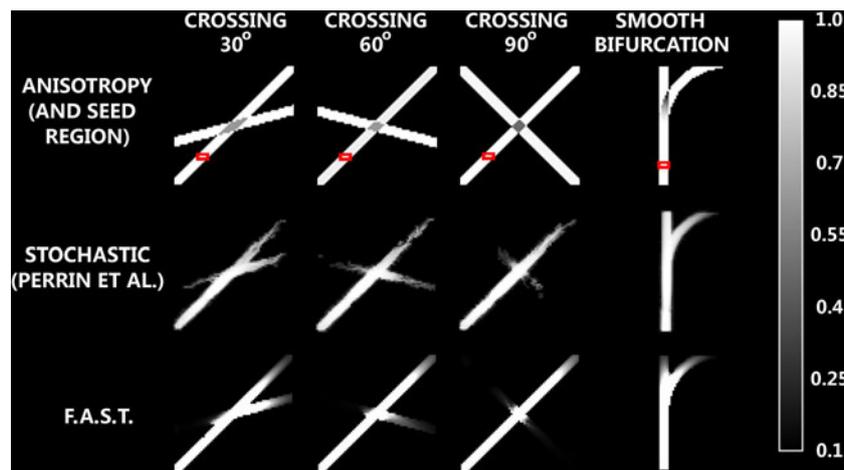


Fig. 3 Axial slice of synthetic fiber crossings at different angles and a smooth bifurcation: anisotropy (*top row*), connection probability map from method by Perrin et al. (*middle*) and probability map from FAST (*bottom*). The red box in the fractional anisotropy images marks the

seed region. The intensity of the connectivity map is stretched to [0,1], thresholded at 0.1 and log-transformed for display. The threshold, which removes artifacts of the random walk, was determined by visual inspection

region is located at the genu of the corpus callosum, but the metrics for the other locations are also satisfactory and quite consistent with one another. Even the cingulum, which is the most curved of the evaluated tracts, displays good correlation. Figure 4 compares the maximum intensity projection of the unthresholded outputs for the different tracts in a sample volume. Again, the visual appearance of the connection probability maps for the two methods is very similar.

Comparison of Connectivity Measurements

Since the goal of stochastic tractography is to estimate how strongly connected different parts of the brain are, it is interesting to compute the connectivity between regions with FAST and Perrin et al.’s algorithm and compare the scores. Given a seed region and a target ROI, the connectivity is defined as the sum of the connection probability map (again, unthresholded) within the target ROI. For these experiment, the five seed/ROI pairs defined in “*Experimental Setup*” were used.

Figure 5 shows renderings of the seeds, target ROIs and output connection probability maps for the five tracts. The maps are very similar, and the connectivity values provided

by FAST and Perrin et al.’s method are well correlated ($\rho=0.83$). Figure 6 shows a scatter plot for the connection probabilities given by the two methods for the five seed-target regions and the ten subjects (50 points in total). Because the output of the stochastic method is not deterministic, the range of the results for three executions is displayed in the plot. These results indicate that FAST is a good approximation of the stochastic method it is based on (even for the highly curved cingulum bundle), except when a very high-precision probability map is required. In that case, it can still be beneficial to use FAST to interactively find a good location for the seed regions, and then run the sampling-based method. This way the output of the stochastic method will be assuredly satisfactory, as opposed to possibly having to modify the seed location and rerun the slow probabilistic method multiple times.

Execution Speed

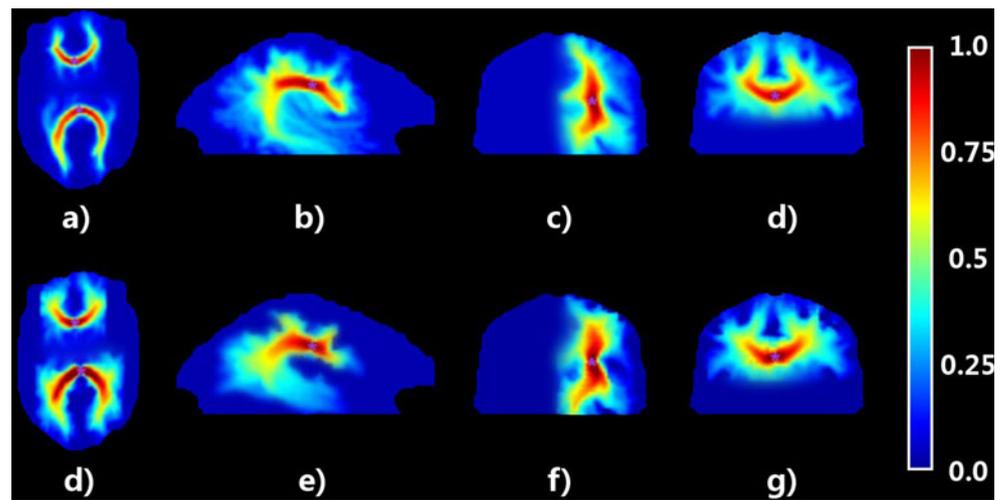
The most noteworthy feature of the proposed method is its speed. Both approaches were implemented in Java and tested on a desktop with an Intel i7 CPU. In the implementation of the stochastic method, one million particles

Table 3 Similarity metrics comparing the normalized tractograms produced by FAST and by Perrin et al’s method. For each seed region and metric, the mean and standard deviation across the ten subjects is

displayed. MS stands for mean squares, and NCC for normalized cross-correlation

Metric	Seed Region					
	Genu of corpus callosum	Splenium of corpus callosum	Cingulum	Pyramidal tract	Corpus callosum (left/right)	All five combined
MS ($\times 10^{-3}$)	0.101±0.132	0.482±0.284	0.265±0.166	0.363±0.272	0.468±0.255	0.336±0.269
NCC	0.869±0.071	0.778±0.079	0.752±0.040	0.824±0.088	0.867±0.064	0.818±0.081

Fig. 4 Maximum intensity projections of the connection probability map for one of the ten subjects. The seeds are marked with purple asterisks. The probability maps have been normalized to [0,1] and log-transformed for display, but not thresholded. **a** Splenium and genu of corpus callosum, Perrin et al.'s method, axial projection. **b** Cingulum, sagittal projection. **c** Pyramidal tract, coronal projection. **d** Left/right connections of corpus callosum, coronal projection. **e–h** Outputs from FAST



were sampled in order to obtain a high-quality connection probability map. The algorithms were parallelized to take advantage of the multi-core structure of the computer. Parallelization is simple in both cases: for the stochastic method, each particle is independent of the others. For each time step in FAST, different threads can work on different parts of the state bivector corresponding to different image regions. The preprocessing can also be parallelized easily, because both the estimation of the fODF and the calculation of the transition tensor can be divided into regions which are handled by different threads.

The running times of the different steps of the two methods, averaged over ten tractographies of the pyramidal tract on the volume in Fig. 6, are displayed in Table 4. The speed-up is especially large if the seed region is to be moved due to unsatisfactory results or if a connection

probability map from a different brain structure is to be computed. In this case, recalculating the transition tensor (the most computationally expensive step of our algorithm) is not necessary and the new output can be generated in a few seconds, whereas the stochastic sampling would take approximately 7 min. This makes it feasible to use FAST in an interactive fashion, which is not possible with stochastic methods.

Discussion

A very fast probabilistic tractography method has been presented in this study. Given HARDI data, the algorithm produces a connectivity map for a seed region in the brain. The method relies on discretizing the directions of

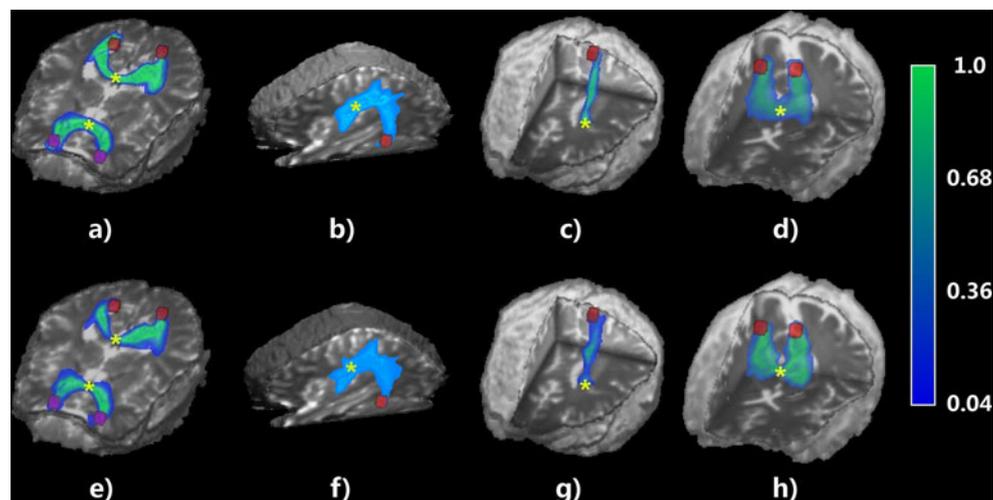


Fig. 5 rendering of T2 baseline volume, seed and target regions, and resulting connection probability map. The seeds are marked with yellow asterisks, whereas the target ROIs and marked with colored spheres. The intensity of the probability maps is normalized to [0,1], thresholded at 0.04 and log-transformed for display. Once more, the

threshold was determined by visual inspection. **a** Genu (*purple spheres*) and splenium (*red spheres*) of corpus callosum for Perrin et al.'s method. **b** Cingulum tract **c** Corticospinal tract. **d** Left/right connections of corpus callosum. **e–h** Output provided by FAST

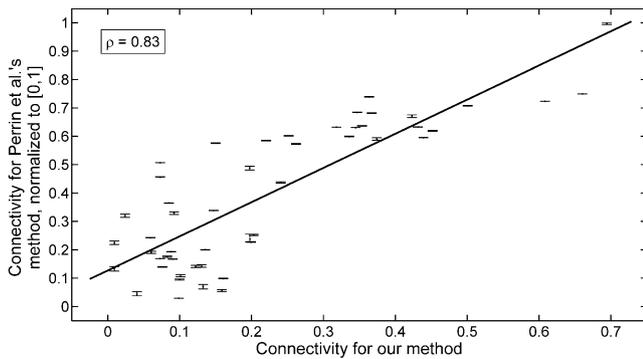


Fig. 6 scatter plot of the connectivity measures provided by our method and Perrin et al.'s. Each data point shows the range over three executions of the latter. The connectivity was normalized to the maximum value. The correlation coefficient is $\rho=0.83$

propagation to calculate the connection probability map in a deterministic fashion in a very short time. Probability mass represented by a state bivector flows from the seed region following a transition tensor in a similar way in which particles do in sampling-based stochastic tractography methods. The bivector and the tensor include information not only on the position but also on the direction of movement, which allows the proposed framework to accommodate second-order models of particle propagation such as trajectory regularization through particle deflection.

The main difference between FAST and other fast discrete propagation methods is this second-order model. Neither graph-based techniques nor front propagation algorithms can simulate the propagation of a particle with inertia in the white matter, which can be desirable in some tracking applications. On the other hand, all the discrete propagation methods discussed in the introduction (Fout et al. 2005; Merhof et al. 2006; Iturria-Medina et al. 2007; Sotiropoulos et al. 2010a; Parker et al. 2002; Zalesky 2008; Lifshits et al. 2009; Sotiropoulos et al. 2010b; Staempfli et al. 2006) can be formulated in our framework, since:

- It supports all the affinity functions used in those studies. We used a function of the direction of arrival and the ODF in order to replicate the results from Perrin et al., but it would be immediate to replace it by an arbitrary function (for example including the ODF of the proposed destination voxel).
- It supports any type of neighborhood.
- It supports shortest path search using Dijkstra's algorithm, even though this aspect has not been analyzed in this study.

Another contribution of this paper is to quantitatively assess the impact of the discretization on the tractography results. We have compared FAST to an imple-

mentation of Perrin et al.'s stochastic tractography method. The transition tensor was designed to mimic the behaviour of particles in the baseline method as closely as possible. The results on synthetic and real data show that FAST produces connectivity measures that correlate well with the output of the baseline stochastic method.

The best feature of FAST is that, once the transition tensor has been computed, it is approximately 60 times faster than the stochastic method it is based on. Loading the volume, computing the fODF and the transition tensor takes approximately 45 s. We did not use correction for solid angle (Aganj et al. 2010) when computing the ODF nor constrained optimization methods when deconvolving it into the fODF (Jian and Vemuri 2007), which would make these initial computations slower. Once the tensor is ready, the tracking part of the algorithm runs in approximately 7 s. This way, a user can utilize FAST to explore the connectivity of different parts of the brain, quickly refining the position and size of the seed regions if he is not satisfied with the results.

Though 98 orientations typically suffice to capture a wide range of fiber bundles, the limited directional resolution of the method can still be a disadvantage when tracking very thin and/or curved fiber tracts: if the fibers happen to lie along a direction that is relatively far from the 98 allowed orientations, the algorithm might still have problems following the bundle. Another minor disadvantage of the method is that, as explained above, it requires precomputing the transition matrix before starting the tracking.

Finally, it is important to note that FAST requires a certain amount of memory to store the transition tensor. The implementation takes advantage of the fact that the majority of the elements of the tensor are zero elements because they correspond to illegal transitions, but it still requires 11 GB of RAM. However, there are different ways in which the memory requirements could be lowered. First, it could be possible to replace our in-house implementation of sparse matrices by a more efficient one. Another possibility would be to use 16-bit floating point data to represent the elements of the transition tensor, instead of the 32-bit representation used

Table 4 Execution times for the different computations required by FAST and the method by Perrin et al.

Computation	Method	
	Perrin et al.	FAST
Fiber ODF (fODF)	2.58±0.14 s	2.70±0.12 s
Transition tensor	N/A	43.9±2.1 s
Fiber tracking	442±18 s	7.28±0.33 s

in the current version. The impact on the results should be minimal, whereas the memory requirements would be reduced by 50%. Finally, it would also be possible to use memory compression techniques, but that would make the method slower.

Information Sharing Statement

The data used in this study belong to an ongoing study and cannot be shared at the moment. However, they will eventually be released and posted at the website of the Laboratory of Neuro Imaging (<http://www.loni.ucla.edu/Research/Databases/>). At that time, the JAVA code (both source and compiled) will be made available at the first author's website (<http://www.jeiglesias.com>) so that the results in this study can be reproduced.

Acknowledgements This work was funded by the National Science Foundation (grant no. 0844566), the Office of Naval Research (grant no. N000140910099), and the National Institutes of Health through grants U54 RR021813 (Center for Computational Biology), IU01MH093765 (The Human Connectome Project) and 5P41RR013642 (Computational Anatomy and Multidimensional Modeling). The authors would like to thank Greig de Zubicaray, Kathie McMahon, Margaret Wright from the Center for Magnetic Resonance at 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.

References

- Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., & Harel, N. (2010). Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. *Magnetic Resonance in Medicine*, *64*(2), 554–566.
- Aja-Fernandez, S., Alberola-Lopez, C., & Westin, C. (2007). Signal LMMSE estimation from multiple samples in MRI and DT-MRI. *Lecture Notes in Computer Science (MICCAI 2007)*, *4792*, 368–375.
- Alexander, A. L., Hasan, K. M., Lazar, M., Tsuruda, J. S., & Parker, D. L. (2001). Analysis of partial volume effects in diffusion-tensor MRI. *Magnetic Resonance in Medicine*, *45*(5), 770–780.
- Alexander, D. C., Barker, G. J., & Arridge, S. R. (2002). Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, *48*(2), 331–340.
- Barmpoutis, A., Jian, B., Vemuri, B. C., & Shepherd, T. M. (2007). Symmetric positive 4th order tensors & their estimation from diffusion weighted MRI. *Lecture notes in computer science (IPMI 2007)*, *4584*, 308–319.
- Basser, P. J. (1995). Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. *Nuclear Magnetic Resonance in Biomedicine*, *8*(7), 333–344.
- Basser, P. J., Mattiello, J., LeBihan, D., et al. (1994). Estimation of the effective self-diffusion tensor from the NMR spin echo. *Journal of Magnetic Resonance Series B*, *103*, 247.
- Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., & Aldroubi, A. (2000). In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, *44*(4), 625–632.
- Behrens, T. E. J., Woolrich, M. W., Jenkinson, M., Johansen-Berg, H., Nunes, R. G., Clare, S., et al. (2003). Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, *50*(5), 1077–1088.
- Behrens, T. E. J., Berg, H. J., Jbabdi, S., Rushworth, M. F. S., & Woolrich, M. W. (2007). Probabilistic diffusion tractography with multiple fiber orientations: what can we gain? *Neuroimage*, *34*(1), 144–155.
- Campbell, J. S. W., Siddiqi, K., Rymar, V. V., Sadikot, A. F., & Pike, G. B. (2005). Flow-based fiber tracking with diffusion tensor and q-ball data: validation and comparison to principal diffusion direction techniques. *NeuroImage*, *27*(4), 725–736.
- Cheng, J., Ghosh, A., Jiang, T., & Deriche, R. (2009). A Riemannian framework for orientation distribution function computing. *Lecture Notes in Computer Science (MICCAI 2009)*, *5761*, 911–918.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., & Deriche, R. (2007). Regularized, fast, and robust analytical Q-ball imaging. *Magnetic Resonance in Medicine*, *58*(3), 497–510.
- Descoteaux, M., Deriche, R., Knosche, T. R., & Anwander, A. (2009). Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Transactions on Medical Imaging*, *28*(2), 269–286.
- Fout, N., Huang, J., & Ding, Z. (2005). Visualization of neuronal fiber connections from DT-MRI with global optimization. In *Proceedings of the ACM symposium on Applied computing 2005*, pages 1200–1206. NY: ACM.
- Frank, L. R. (2001). Anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, *45*(6), 935–939.
- Frank, L. R. (2002). Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magnetic Resonance in Medicine*, *47*(6), 1083–1099.
- Friman, O., Farnback, G., & Westin, C. (2006). A Bayesian approach for stochastic white matter tractography. *IEEE Transactions on Medical Imaging*, *25*(8), 965–978.
- Ghosh, A., Descoteaux, M., & Deriche, R. (2008). Riemannian framework for estimating symmetric positive definite 4th order diffusion tensors. *Lecture notes in computer science (MICCAI 2008)*, *5241*, 858–865. Springer.
- Iglesias, J. E., Thompson, P. M., Liu, C. Y., & Tu, Z. (2010). Discretizing stochastic tractography: a fast implementation. *Proceedings of IEEE ISBI 2010*, 1381–1384.
- Iturria-Medina, Y., Canales-Rodriguez, E. J., Melie-Garcia, L., Valdes-Hernandez, P. A., Martinez-Montes, E., Aleman-Gomez, Y., et al. (2007). Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *Neuroimage*, *36*(3), 645–660.
- Jbabdi, S., Woolrich, M. W., Andersson, J. L. R., & Behrens, T. E. J. (2007). A Bayesian framework for global tractography. *Neuroimage*, *37*(1), 116–129.
- Jian, B., & Vemuri, B. C. (2007). A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Transactions on Medical Imaging*, *26*(11), 1464–1471.
- Kaden, E., Knosche, T. R., & Anwander, A. (2007). Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. *NeuroImage*, *37*(2), 474–488.
- Koch, M. A., Norris, D. G., & Hund-Georgiadis, M. (2002). An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*, *16*(1), 241–250.
- Kreher, B. W., Schneider, J. F., Mader, I., Martin, E., Hennig, J., & Ilyasov, K. A. (2005). Multitensor approach for analysis and tracking of complex fiber configurations. *Magnetic resonance in medicine*, *54*(5), 1216–1225.

- Lazar, M., Weinstein, D. M., Tsuruda, J. S., Hasan, K. M., Arfanakis, K., Meyerand, M. E., et al. (2003). White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18(4), 306–321.
- Lifshits, S., Tamir, A., & Assaf, Y. (2009). Combinatorial fiber-tracking of the human brain. *NeuroImage*, 48(3), 532–540.
- Merhof, D., Richter, M., Enders, F., Hastreiter, P., Ganslandt, O., Buchfelder, M., et al. (2006). Fast and accurate connectivity analysis between functional regions based on dt-mri. *Lecture notes in computer science (MICCAI 2006)*, 4191, 225–233. Springer.
- Osher S., & Fedwik, R. P. (2003). Level set methods and dynamic implicit surfaces. Springer.
- Ozarslan, E., & Mareci, T. H. (2003). Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 50(5), 955–965.
- Parker, G. J. M., & Alexander, D. C. (2003). Probabilistic Monte Carlo based mapping of cerebral connections utilizing whole-brain crossing fiber information. *Lecture Notes in Computer Science (IPMI 2003)*, 732, 684–696.
- Parker, G. J. M., & Alexander, D. C. (2005). Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions B*, 360(1457), 893–902.
- Parker, G. J. M., Wheeler-Kingshott, C. A. M., & Barker, G. J. (2002). Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 21(5), 505–512.
- Perrin, M., Poupon, C., Cointepas, Y., Rieul, B., Golestani, N., Pallier, C., et al. (2005). Fiber tracking in q-ball fields using regularized particle trajectories. *Lecture notes in computer science (IPMI 2005)*, 3565, 52–63. Springer.
- Perrin, M., Cointepas, Y., Cachia, A., Poupon, C., Thirion, B., Riviere, D., Cathier, P., El Kouby, V., Constantinesco, A., & Le Bihan, D. Mangin, J. F. (2008). Connectivity-based parcellation of the cortical mantle using q-ball diffusion imaging. *Journal of Biomedical Imaging*, 1–10. January
- Pierpaoli, C., Jezzard, P., Basser, P. J., Barnett, A., & Di Chiro, G. (1996). Diffusion tensor MR imaging of the human brain. *Radiology*, 201(3), 637–648.
- Qazi, A. A., Radmanesh, A., O'Donnell, L., Kindlmann, G., Peled, S., Whalen, S., et al. (2009). Resolving crossings in the corticospinal tract by two-tensor streamline tractography: Method and clinical assessment using fMRI. *Neuroimage*, 47, 98–106.
- Sethian, J. A. (2000). *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Cambridge University Press.
- Smith, S. M. (2002). Fast robust automated brain extraction. *Human Brain Mapping*, 17(3), 143–155.
- Sotiropoulos, S. N., Bai, L., & Tench, C. R. (2010a). Fuzzy anatomical connectedness of the brain using single and multiple fiber orientations estimated from diffusion MRI. *Computerized Medical Imaging and Graphics*, 24(6), 504–513.
- Sotiropoulos, S. N., Bai, L., Morgan, P. S., Constantinescu, C. S., & Tench, C. R. (2010b). Brain tractography using Q-ball imaging and graph theory: improved connectivities through fibre crossings via a model-based approach. *NeuroImage*, 49(3), 2444–2456.
- Staempfli, P., Jaermann, T., Crelier, G. R., Kollias, S., Valavanis, A., & Boesiger, P. (2006). Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. *Neuroimage*, 30(1), 110–120.
- Tournier, J. D., Calamante, F., Gaidan, D. G., & Connelly, A. (2004). Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *Neuroimage*, 23(3), 1176–1185.
- Tuch, D. S. (2004). Q-Ball Imaging. *Magnetic Resonance in Medicine*, 52(6), 1358–1372.
- Tuch, D. S., Reese, T. G., Wiegell, M. R., Makris, N., Belliveau, J. W., & Wedeen, V. J. (2002). High angular resolution diffusion imaging reveals intravoxel white matter heterogeneity. *Magnetic Resonance in Medicine*, 48(4), 577–582.
- Wedeen, V. J., Wang, R. P., Schmahmann, J. D., Benner, T., Tseng, W. Y. I., Dai, G., et al. (2008). Diffusion spectrum magnetic resonance imaging (DSI) tractography of crossing fibers. *NeuroImage*, 41(4), 1267–1277.
- Zalesky, A. (2008). DT-MRI fiber tracking: a shortest paths approach. *IEEE Transactions on Medical Imaging*, 27(10), 1458–1471.