



Probabilistic white matter fiber tracking using particle filtering and von Mises–Fisher sampling

Fan Zhang^{a,*}, Edwin R. Hancock^b, Casey Goodlett^b, Guido Gerig^b

^a Department of Computer Science, University of York, York YO10 5DD, UK

^b Scientific Computing and Imaging Institute, University of Utah, UT 84112, USA

ARTICLE INFO

Article history:

Received 28 September 2007

Received in revised form 12 May 2008

Accepted 20 May 2008

Available online 5 June 2008

Keywords:

Diffusion tensor MRI

Tractography

Probabilistic fiber tracking

Particle filtering

von Mises–Fisher sampling

ABSTRACT

Standard particle filtering technique have previously been applied to the problem of fiber tracking by Brun et al. [Brun, A., Bjornemo, M., Kikinis, R., Westin, C.F., 2002. White matter tractography using sequential importance sampling. In: Proceedings of the ISMRM Annual Meeting, p. 1131] and Bjornemo et al. [Bjornemo, M., Brun, A., Kikinis, R., Westin, C.F., 2002. Regularized stochastic white matter tractography using diffusion tensor MRI, In: Proc. MICCAI, pp. 435–442]. However, these previous attempts have not utilised the full power of the technique, and as a result the fiber paths were tracked in a goal directed way. In this paper, we provide an advanced technique by presenting a fast and novel probabilistic method for white matter fiber tracking in diffusion weighted MRI (DWI), which takes advantage of the weighting and resampling mechanism of particle filtering. We formulate fiber tracking using a non-linear state space model which captures both smoothness regularity of the fibers and the uncertainties in the local fiber orientations due to noise and partial volume effects. Global fiber tracking is then posed as a problem of particle filtering. To model the posterior distribution, we classify voxels of the white matter as either prolate or oblate tensors. We then construct the orientation distributions for prolate and oblate tensors separately. Finally, the importance density function for particle filtering is modeled using the von Mises–Fisher distribution on a unit sphere. Fast and efficient sampling is achieved using Ulrich–Wood’s simulation algorithm. Given a seed point, the method is able to rapidly locate the globally optimal fiber and also provides a probability map for potential connections. The proposed method is validated and compared to alternative methods both on synthetic data and real-world brain MRI datasets.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Diffusion tensor MRI (DTI) has become a popular tool for non-invasive exploration of the anatomical structure of the white matter in vivo (Basser et al., 1994). It endows each voxel with a 3×3 symmetric positive-definite matrix, which characterises the local water diffusion process. It is based on a local Gaussianity assumption concerning the probability of water molecule motion in a defined time period. White matter fiber tracking or “tractography” estimates likely fiber paths by tracing the local tensor orientations (Mori and van Zijl, 2002; Parker, 2004). In this paper, we present a new and fast probabilistic fiber tracking algorithm which utilises the particle filtering technique and von Mises–Fisher sampling.

1.1. Related literature

1.1.1. Fiber tracking

Broadly, the fiber tracking methods described in the literature can be classified as belonging to two groups. The first group of methods is based on local line propagation techniques (or streamline techniques) (Basser et al., 2000; Lazar et al., 2003; Mori et al., 1999). Step-by-step, they integrate a fiber pathway from a predefined seed point along the principal diffusion directions, which correspond to the principal eigenvectors of the diffusion tensors. The main difference among the methods in this group is the way in which local information is incorporated to locate smooth fiber paths. For instance, Lazar et al. (2003) use the entire diffusion tensor to deflect the estimated fiber trajectory in the desired directions. The main drawback of line propagation methods is that errors accumulate as the propagation takes place over a long distance.

The second group of methods are based on global optimisation techniques (Gossel et al., 2002; Parker et al., 2002; Pichon et al., 2005; Staempfli et al., 2006). Starting from a seed point, they

* Corresponding author. Tel.: +441904432709; fax: +441904432767.

E-mail addresses: zfan@cs.york.ac.uk (F. Zhang), erh@cs.york.ac.uk (E.R. Hancock), gerig@sci.utah.edu (C. Goodlett), gcasey@sci.utah.edu (G. Gerig).

attempt to locate an improved estimate of the true fiber pathway using energy minimisation techniques. For instance, Parker et al. (2002) apply the fast marching technique to propagate connection paths determined by the principal eigenvectors of the tensors. Gossel et al. (2002) apply Kalman filtering to track globally optimal paths, according to a fiber smoothness criterion. Pichon et al. (2005) determine the optimal path between two voxels by solving the Hamilton–Jacobi–Bellman equation using dynamic programming. Prados et al. (2006) use Riemannian geometry and control theory to trace the neural fiber bundles by computing the geodesic distances between seed and end point locations. More recently, Fletcher et al. (2007) develop a volumetric approach for quantitatively studying region-to-region white matter connectivity from diffusion tensor MRI. They use the Hamilton–Jacobi equation to formulate the minimal path problem between two regions.

One common feature of the above methods is that local fiber orientations are determined in a purely deterministic way. However, due to both noise (Macovski, 1996) and ambiguities for voxels where multiple fibers cross or branch (partial volume effects (Alexander et al., 2001)), the local fiber orientations measured by DTI are not completely reliable. For instance, in the case of partial volume effects, the diffusion process within voxels is no longer Gaussian. As a result, the diffusion tensor is an incomplete model of the diffusion signal. A natural way to deal with this problem is to measure uncertainty of the local fiber orientation measurement probabilistically at each voxel. Jones (2003) quantifies the uncertainties in the fiber orientations using the bootstrap method and then visualises the orientational field using uncertainty cones.

1.1.2. Probabilistic fiber tracking

As a result, probabilistic fiber tracking methods have received considerable interest recently as a means of incorporating orientational uncertainties (Brun et al., 2002; Bjornemo et al., 2002; Behrens et al., 2003, 2007; Friman et al., 2006; Lazar and Alexander, 2005; Parker et al., 2003; Parker and Alexander, 2003, 2005). Instead of reconstructing fiber pathways deterministically, they aim to measure the probability of connectivity between brain regions.

These methods have two stages. In the first stage, they model the uncertainty in DTI measurements at each voxel using a probability density function (PDF) for the fiber orientations (Behrens et al., 2003; Friman et al., 2006). Behrens et al. (2003) were the first to formalise the PDF for local fiber orientations using a Bayesian framework. They present three models for describing the local diffusion process of water molecules, and their models use different levels of complexity. The parameters of the models are estimated using a Gibbs sampler. They applied a Markov Chain Monte Carlo (MCMC) procedure to sample from a single fiber orientation distribution. They then used the set of samples to model the uncertainty in orientation. Friman et al. (2006) proposed an alternative Bayesian method based on a simplified and more tractable diffusion tensor model. They replace the continuous PDF of fiber orientation by a set of uniform samples on a unit sphere. However, their PDF does not consider the uncertainties due to partial volume effects. Broadly speaking, the diffusion tensor allows the reliable definition of the principle diffusion direction, which corresponds to the local fiber orientation. However, in voxels with more complex configurations such as fiber crossings, the orientation conveyed by a diffusion tensor is ambiguous. To overcome this difficulty, a multi-tensor (or multi-compartment) extension of the single diffusion tensor model has been proposed (Tuch et al., 2002; Frank, 2001; Behrens et al., 2007). The basic idea here is to approximate the diffusion process using a mixture of Gaussian densities (Frank, 2001). Behrens et al. (2007) determine the complexity of the fiber structure at each voxel using automatic relevance determination. They then extend their previous work (Behrens et al., 2003) to deal with

multi-orientations within a voxel by using a multi-compartment description of the diffusion process. More sophisticated methods for modeling PDFs of multi-fiber orientations at a voxel have also been developed for characterising the observed diffusion. For instance, Tuch (2004) describes the observed diffusion for high angular resolution diffusion imaging (HARDI) using the q -space framework, which is able to resolve fiber crossings. Tournier et al. (2004) express the diffusion weighted signal as a spherical convolution of the response function. The approximate PDF for fiber orientations is then recovered using spherical deconvolution. McGraw et al. (2006) model the PDF for fiber orientations using a mixture of von Mises–Fisher distributions. They use this model for segmenting high angular resolution diffusion MRI. Bhalerao and Westin (2007) build the PDFs from HARDI using a hyperspherical von Mises–Fisher mixture model. They map the fiber orientation samples to a 5D representation which can avoid the ambiguities associated with the sign flips of directions in 3D. They then fit their model using Maximum Likelihood estimation. Although these sophisticated models can better represent diffusion weighted signals, new imaging methods require much finer angular resolution and more scanning time than DTI (often 50 or more gradient directions are required). Therefore, in this paper, we focus on estimating the PDF of fiber orientations using the widely adopted diffusion tensor imaging model.

In the second stage, probabilistic tracking algorithms simply repeat a streamline propagation process (typically 1000–10000 times) with propagation directions randomly sampled from the PDF for fiber orientation. The fraction of the streamlines that pass through a voxel provide an index of the strength of connectivity between that voxel and the seed point. By contrast with the first stage, few methods have been developed to efficiently sample fiber paths from orientation PDF. The main difference between existing methods for probabilistic fiber tracking is found in the way that the PDF for fiber orientation is modeled. Most methods estimate the connectivity map by sampling directly from the PDF for fiber orientation. For instance, Parker et al. (2003) modeled the uncertainty of the fiber orientations using the normal distribution, and the parameters of the distribution is controlled in a heuristic way. They then sequentially sample fiber paths from the normal distribution. Parker and Alexander (2003) have further developed their previous work to deal with multi-fiber crossings. They distinguish between prolate and oblate tensors using the spherical harmonic parameters of the diffusion weighted signals. They fit the DWIs using both single and multi-tensor models. The normal distribution is used to control the uncertainties and sample fiber paths in a efficient way. Cook et al. (2004) further improved the method by capturing the orientation errors using the Watson distribution. The advantage of these methods is that they are able to efficiently sample fiber paths in a probabilistic way. Although the normal distributions are easily sample from, the choice is motivated by computational expediency and it is not clear that it accurately models the sources of uncertainty present. The main differences between these sequential sampling methods of Parker et al. (2003), Parker and Alexander (2003), Cook et al. (2004) and our method are the weighting and resampling stages of particle filtering. On the other hand, the uncertainties of the fiber orientations have been modeled in a theoretically principled Bayesian framework by Behrens et al. (2003, 2007) and Friman et al. (2006). However, the fiber orientation distributions are difficult to simulate directly. As a result, the sampling process is not easily tractable. Thus it is necessary to resort to Markov Chain Monte Carlo methods (Behrens et al., 2003) or to evaluate the PDF discretely with low angular resolution. One drawback of these tracking methods is their computational complexity (often more than several hours on a high-end PC (Behrens et al., 2003; Friman et al., 2006)), and this is unacceptable in practice.

1.1.3. Particle filtering for fiber tracking

Particle filtering was originally investigated for probabilistic fiber tracking by Brun et al. (2002) and Bjornemo et al. (2002). In a short abstract paper, Brun et al. (2002) were the first to sketch out the idea of applying the sequential importance sampling and resampling mechanism (also known as particle filtering) for tractography. Bjornemo et al. (2002) further detailed and developed the ideas in Brun et al. (2002). Specifically, they discussed in detail the conceptual model of importance sampling and resampling. They then construct the uncertainties of fiber orientations using a Gaussian noise model. Stochastic fiber paths are generated from the noisy distribution in a goal directed way. Finally, The connections between the proposed tracking model and the importance sampling and resampling framework are discussed. Since the weights of the sampled paths are constant throughout the propagation process, the weighting and resampling technique are not incorporated into their tracking model.

1.2. Contribution

Developing on the work of Brun et al. (2002) and Bjornemo et al. (2002), in this paper we present a fast and novel probabilistic method for white matter fiber tracking based on particle filtering and von Mises–Fisher sampling. We first formulate fiber tracking using a non-linear state space model and recursively compute the posterior distribution using particle filtering (Doucet et al., 2000, 2001; Gordon et al., 1993). This technique has been successfully used in computer vision for visual tracking (Isard and Blake, 1998) and contour extraction (Perez et al., 2001). It provides a sound statistical framework for propagating a sample-based approximation of the posterior distribution. There are almost no restrictions on the type of model that can be used. As a result, we can model the fiber orientation distributions for prolate tensors and oblate tensors separately. The proposed tracking model can capture both smoothness regularity of the fibers and the uncertainties of the local fiber orientations due to both noise and partial volume effects. Since samples from the posterior path distribution are maintained at each propagation step, different decision criteria can be used to identify the optimal fiber. This procedure is similar to the active testing and tree pruning method for maximum a posteriori (MAP) road tracking developed by Geman and Jedynak (1996). Given a seed point, our method is able to rapidly locate the global optimal fiber path and also provide a probabilistic connectivity map between the seed point and all other voxel locations.

Compared to the methods proposed by Brun et al. (2002) and Bjornemo et al. (2002), we make the following two novel contributions.

First, our proposed method generates reliable fiber paths by fully utilising the weighting and resampling technique of particle filtering. We model the probability density for fiber orientations in a theoretically justified way. In order to overcome the difficulty of sampling the theoretical fiber orientation distribution, we simulate paths using a simpler approximating distribution (namely the importance density function) for fiber orientations which is efficient to sample from. We then sequentially evaluate and adjust the sampled paths according to the true orientation distribution. In this way, the method achieves both efficiency and accuracy. Second, to implement an effective particle filtering algorithm for fiber tracking, we first apply the von Mises–Fisher distribution to model the prior and importance density of our tracking process. Fast sampling is realised on the unit sphere and fiber paths are efficiently generated using the simulation algorithm for the von Mises–Fisher distribution developed by Ulrich (1984) and Wood (1994).

A preliminary version of the algorithm reported here was first described in a conference paper (Zhang et al., 2008). However, here we consolidate the work and expand the description of the meth-

od. Moreover, we provide additional qualitative and quantitative results to validate the method. The outline of the paper is as follows. In Section 2, we first formally develop the global fiber tracking model based on a non-linear state space model, and we also show how particle filtering can be used to recursively compute the posterior distribution, so as to compute optimal fiber paths and a probabilistic connectivity map. Section 3 provides the model ingredients necessary to implement the global tracking technique described in Section 2. Specifically, we describe how to construct the observation density, the prior density, the importance density function and the simulation for von Mises–Fisher distribution. In Section 4, we evaluate both qualitatively and quantitatively the performance of the algorithm on synthetic data and real-world diffusion MRI datasets. We also compare the results of our method with those obtained using alternative methods described elsewhere in the literature. Section 5 concludes the paper and discusses directions for future research.

2. Tracking algorithm

The problem of fiber tracking from a 3D diffusion MRI volume is to extract the most likely fiber pathway from a predefined seed point. Contrary to the standard tracking problem where local information is gathered as time progresses, we collect the whole set of data before tracking begins. Thus, at each step of propagation, we set the observation set as the data visible only from the current position. In this sense, the fiber tracking problem is similar to the contour extraction (Perez et al., 2001) and road tracking (Geman and Jedynak, 1996) in computer vision. However, it is a more challenging problem because there are numerous fiber paths in white matter that exhibit crossings and dispersion.

We formulate the fiber tracking problem using a state space model. Given the prior probability densities that characterise the properties of the expected fiber paths and the observation densities that characterise the uncertainty of local fiber orientations, a posterior distribution of the target fiber can be estimated. Because of the complex geometry of the fiber paths and the various uncertainties of the orientation measurements, both the prior density and the observation density are non-Gaussian. Thus, standard linear state space techniques such as the Kalman filter are inappropriate here, and a non-linear filter is necessary. In contrast to the work of Gossel et al. (2002) which used the Kalman filter to locate the optimal path with regard to smoothness constraint for the fibers, our method deals with both smoothness regularity and uncertainties of fiber orientations induced by noise and partial volume effects.

2.1. Global tracking model

A white matter fiber path P can be modeled as a sequence of points in the image space $\Omega \subset R^3$, i.e. $P_{n+1} = (x_0, x_1, \dots, x_{n+1})$, as shown in Fig. 1. Thus, commencing from a seed point x_0 , the progressive growth of a path in discrete time can be described as

$$x_{i+1} = x_i + \rho_i \hat{v}_i, \quad (1)$$




Fig. 1. Path representation.

where ρ_i and \hat{v}_i are, respectively, the step size and the direction of propagation (unit vector) at step i . As with most previous methods, we set the step size to be a constant, i.e. $\rho_i = \rho, i = 0, \dots, n$. Thus, the dynamics of path growth is only determined by the propagation directions \hat{v}_i . In the following, we denote a path as a sequence of unit vectors $P_{n+1} = \hat{v}_{0:n} = \{\hat{v}_0, \dots, \hat{v}_n\}$. Let \mathcal{Y} be the set of observations or image data of a 3D diffusion weighted imaging (DWI) volume. The image data observed at the location indexed i , \hat{v}_i , is $y_i = \mathcal{Y}(\hat{v}_i) = \mathcal{Y}(x_i)$, i.e. $\mathcal{Y}(\hat{v}_i)$ refers to the DWI information observed at the starting point of unit vector \hat{v}_i connecting x_i and x_{i+1} (see Fig. 1). Our goal is to propagate a sequence of unit vectors $\{\hat{v}_i, i = 0, 1, \dots, n\}$ that best estimate the true fiber path based on the conditional prior density $p(\hat{v}_{i+1}|\hat{v}_{0:i})$ and the conditional observation density $p(\mathcal{Y}|\hat{v}_{0:i})$.

We assume that the tracking dynamics forms a Markov chain, so that

$$p(\hat{v}_{i+1}|\hat{v}_{0:i}) = p(\hat{v}_{i+1}|\hat{v}_i). \quad (2)$$

This means the new state is conditioned directly only on the immediately preceding state, and is independent of the past. Thus, the prior for the fiber path is

$$p(\hat{v}_{0:n}) = p(\hat{v}_0) \prod_{i=1}^n p(\hat{v}_i|\hat{v}_{i-1}). \quad (3)$$

We also assume that the observations or diffusion measurements (DWI) are conditionally independent given $\hat{v}_{0:n}$, i.e.

$$p(\mathcal{Y}|\hat{v}_{0:n}) = \prod_{r \in \Omega} p(\mathcal{Y}(r)|\hat{v}_{0:n}), \quad (4)$$

where Ω is the set of voxels of the image volume. Additionally, if we assume that the diffusion measurement at a point does not depend on any points in the history of the path, then $p(y_i|\hat{v}_{0:i}) = p(y_i|\hat{v}_i)$. Based on Eq. (3) and the assumptions, the posterior distribution $p(\hat{v}_{0:n}|\mathcal{Y})$ can be expanded as

$$p(\hat{v}_{0:n}|\mathcal{Y}) = p(\hat{v}_0|\mathcal{Y}) \prod_{i=1}^n p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y}). \quad (5)$$

Applying Bayes theorem, we have

$$p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y}) = \frac{p(y_i|\hat{v}_i)p(\hat{v}_i|\hat{v}_{i-1})}{p(y_i)}. \quad (6)$$

In our tracking model we do not assume any prior information about the diffusion measurements (DWI), thus we can simply consider $p(y_i)$ as a fixed regularity factor of the system. As a result, we can write

$$p(y_i) = \int p(y_i|\hat{v}_i)p(\hat{v}_i|\hat{v}_{i-1})d\hat{v}_i. \quad (7)$$

Most previous methods reported in the literature (Behrens et al., 2003, 2007; Parker et al., 2003; Parker and Alexander, 2003) for probabilistic fiber tracking estimate the posterior $p(\hat{v}_{0:n}|\mathcal{Y})$ by sampling 1000–10000 streamline paths from the conditional observation density $p(y_i|\hat{v}_i)$. However, the density $p(y_i|\hat{v}_i)$ is often complicated, thus the sampling is difficult and time consuming (Behrens et al., 2003, 2007). Moreover, these methods do not take into account the smoothness constraint for fibers. On the other hand, Friman et al. (2006) estimate the posterior by sampling from $p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y})$. This sampling is again difficult and requires the computation of the integral in Eq. (7) over the new state. To avoid these difficulties, Friman et al. discretise the problem using a finite set of several thousand directions for propagating paths from \hat{v}_{i-1} to \hat{v}_i . Thus, sampling the discretised version of $p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y})$ becomes feasible, and the integral becomes a sum. In addition to introducing errors, this discretised sampling is still time consuming, since each discretised direction must be evaluated at all locations in the volume. More-

over, simple sequential sampling methods may degenerate as the number of propagation steps becomes large (Doucet et al., 2001).

2.2. Recursive posterior using particle filtering

We wish to estimate the posterior distribution iteratively in time. By inserting Eq. (6) into Eq. (5), we have

$$p(\hat{v}_{0:n}|\mathcal{Y}) = p(\hat{v}_0|\mathcal{Y}) \prod_{i=1}^n p(\hat{v}_i|\hat{v}_{i-1}) \prod_{i=1}^n \frac{p(y_i|\hat{v}_i)}{p(y_i)}, \quad (8)$$

where $p(\hat{v}_0|\mathcal{Y})$ is predefined. The modeling of the transition probability $p(\hat{v}_i|\hat{v}_{i-1})$ and the distribution $p(y_i|\hat{v}_i)$ will be detailed in the next section of this paper. We recast the problem of tracking the expected fiber path as that of approximating the maximum a Posteriori (MAP) path from the posterior distribution.

It is straightforward to obtain the following recursive formula for the posterior from Eq. (8)

$$p(\hat{v}_{0:i+1}|\mathcal{Y}) = p(\hat{v}_{0:i}|\mathcal{Y}) \frac{p(\hat{v}_{i+1}|\hat{v}_i)p(y_{i+1}|\hat{v}_{i+1})}{p(y_{i+1})}. \quad (9)$$

Since the denominator of this expression requires the evaluation of a complex high-dimensional integral, it is infeasible to locate the maximum likelihood path analytically. In conjunction with the methods described above, we evaluate the posterior using a large number of samples which efficiently characterise the required posterior. Thus, statistical quantities, such as the mean, variance and maximum likelihood, can be estimated using the sample set. Since it is seldom possible to obtain samples from the posterior directly, we use particle filtering (sequential Monte Carlo technique) to recursively compute a finite set of sample paths from the posterior based on Eq. (9).

To sample a set of K paths, we place K particles at the starting point of the path and allow them to propagate as time progresses. Given the states of the set of particles $\{\hat{v}_{0:i}^{(k)}, k = 1, \dots, K\}$ at time i , the process of sequentially propagating the particles to the next time step $i + 1$ can be described in three stages. These are referred to as *prediction*, *weighting* and *selection* and described in detail in the following paragraphs. Let $\pi(\hat{v}_{0:i}|\mathcal{Y})$ be a so-called importance function which has a support including that of the posterior $p(\hat{v}_{0:i}|\mathcal{Y})$. For our sequential importance sampling, suppose that we choose an importance function of the form (Doucet et al., 2000)

$$\pi(\hat{v}_{0:n}|\mathcal{Y}) = \pi(\hat{v}_0|\mathcal{Y}) \prod_{i=1}^n \pi(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y}). \quad (10)$$

In the first *prediction* stage, the simulated path $\hat{v}_{0:i}^{(k)}$ with index k is grown by one step to be $\hat{v}_{0:i+1}^{(k)}$ through sampling from the importance function $\pi(\hat{v}_{i+1}^{(k)}|\hat{v}_i^{(k)}, \mathcal{Y})$. The new set of paths is generally not an efficient approximation of the posterior distribution at time $i + 1$. Thus, in the second or *weighting* stage, we measure the reliability of the approximation using a ratio, referred to as the importance weight, between the true posterior and its approximation. The importance weight is given by

$$w_{i+1}^{(k)} = \frac{p(\hat{v}_{0:i+1}^{(k)}|\mathcal{Y})}{\pi(\hat{v}_{0:i}^{(k)}|\mathcal{Y})\pi(\hat{v}_{i+1}^{(k)}|\hat{v}_i^{(k)}, \mathcal{Y})}. \quad (11)$$

We are more interested in the normalised importance weights, given by

$$\tilde{w}_{i+1}^{(k)} = \frac{w_{i+1}^{(k)}}{\sum_{l=1}^K w_{i+1}^{(l)}}. \quad (12)$$

Inserting Eqs. (9) and (11) into the above expression, we have the proportionality

$$\tilde{w}_{i+1}^{(k)} \propto \tilde{w}_i^{(k)} \frac{p(\hat{v}_{i+1}^{(k)} | \hat{v}_i^{(k)}) p(y_{i+1} | \hat{v}_{i+1}^{(k)})}{\pi(\hat{v}_{i+1}^{(k)} | \hat{v}_i^{(k)}, \mathcal{Y})}. \quad (13)$$

The choice of importance function plays an important role in determining the performance of particle filtering. This choice will be detailed in the next section of this paper. At this point the resulting weighted set of paths provides an approximation of the target posterior. However, the distribution of the weights $\tilde{w}_{i+1}^{(k)}$ may become skewed as time increases. The purpose of the final or *selection* stage is to avoid this degeneracy. We measure the degeneracy of the algorithm using the effective sample size N_{eff} introduced in Liu (1996),

$$N_{\text{eff}} = \frac{1}{\sum_{k=1}^K (\tilde{w}_{i+1}^{(k)})^2}. \quad (14)$$

When N_{eff} is below a fixed threshold N_s , then resampling procedure is used (Doucet et al., 2001; Gordon et al., 1993). The key idea here is to eliminate the paths or particles with low weights $\tilde{w}_{i+1}^{(k)}$ and to multiply offspring particles with high weights. We obtain the surviving particles by resampling K times from the discrete approximating distribution according to the importance weight set $\{\tilde{w}_{i+1}^{(k)}, k = 1, \dots, K\}$.

Both fiber reconstruction and connectivity map computation can be easily accomplished using the discrete distribution of the posterior conveyed by the importance weight set. The MAP estimate of the true fiber path from starting point x_0 is the path with the maximal importance weight. In order to compute the connectivity map, the algorithm records the full tracking history of all the particles at each time step. The probability of connectivity between x_0 and a specific voxel is computed as the fraction of particles that pass through that voxel.

2.2.1. Relationship to alternative tracking methods

Here, we analyse in detail the relationship between our method and the alternative method proposed by Brun et al. (2002) and Bjornemo et al. (2002). Generally speaking, the two methods are developed in rather different settings. Our method is formulated using a non-linear state space model for tractography, and uses particle filtering as a natural tool for finding the posterior distribution of the fiber paths. By contrast, the method in Brun et al. (2002) and Bjornemo et al. (2002) is built on the vanilla particle filtering technique given in Liu et al. (2001). They propose a so-called rough tracking model using the particle filtering technique. The weakness of their method is that they do not incorporate the weighting and resampling steps of particle filtering in their tracking method and the proposed tracking method works in a goal directed way.

More specifically, Bjornemo et al. (2002) proposed the following weighting scheme for sampling

$$w_{i+1} = w_i \frac{p_i(x_{i-1}) p_i(x_i | x_{i-1})}{p_{i-1}(x_{i-1}) \pi_i(x_i | x_{i-1})}. \quad (15)$$

The model assumed for $p_i(x_i | x_{i-1})$ is based on Gaussian noise. Additionally, it is assumed that $\pi_i(x_i | x_{i-1}) = p_i(x_i | x_{i-1})$ and $p_i(x_{i-1}) = p_{i-1}(x_{i-1})$. As a result, $w_{i+1} = w_i$ for all the propagation steps and resampling therefore becomes unnecessary.

Compared to the method in Brun et al. (2002) and Bjornemo et al. (2002), we use the weighting scheme given in Eq. (13). We sample from the importance density $\pi(\hat{v}_{i+1}^{(k)} | \hat{v}_i^{(k)}, \mathcal{Y})$ which is constructed using a von Mises–Fisher model. The importance density is fitted in a simple way to estimate the observation density. As we will show in the following, our observation $p(y_i | \hat{v}_i)$ is based on a theoretically justified noise model. Unfortunately, it is intractable to make samples directly. Although we make samples from a simplified and more tractable importance distribution, we evaluate and recursively adjust the samples according to their true distribution using the weighting and resampling mechanism of particle

filtering. As a result, the sampled fiber paths are simulated so as to reliably reflect the theoretical posterior $p(\hat{v}_{0:i+1} | \mathcal{Y})$ given in Eq. (9).

3. Algorithm ingredients

In this section, we give the details of the local ingredients of the global tracking model using particle filtering.

3.1. Observation density

We commence by showing how to model the observation density $p(y_i | \hat{v}_i)$, which encodes the uncertainty in local fiber orientation due to both noise and partial volume effects. Our observation density function is constructed using a single diffusion tensor model. Despite its weakness in capturing complex fiber structures, DTI is still the most widely used diffusion MRI modality. Formally, the diffusion weighted intensity s_j is related to the diffusion tensor D by the Stejskal–Tanner equation (Basser et al., 1994)

$$s_j = s_0 \exp(-b_j \hat{g}_j^T D \hat{g}_j), \quad (16)$$

where gradient direction \hat{g}_j and the b -value b_j are the scanner parameters for data acquisition, and, s_0 is the intensity with no diffusion gradients applied.

Let $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ be the decreasing eigenvalues of diffusion tensor D and $\hat{e}_1, \hat{e}_2, \hat{e}_3$ be the corresponding normalized eigenvectors. The degree of anisotropy of water diffusion at a voxel can be characterised using the fractional anisotropy (FA) (Basser and Pierpaoli, 1996) of the diffusion tensor D . We can classify the diffusion tensors in white matter into two groups. For prolate tensors, $\lambda_1 > (\lambda_2 \approx \lambda_3)$, and, for oblate tensors, $(\lambda_1 \approx \lambda_2) < \lambda_3$. Tensor classification can be estimated using the prolate shape metric proposed by Westin et al. (2002), i.e.

$$c_l = \frac{\lambda_1 - \lambda_2}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}}. \quad (17)$$

In this paper, we distinguish prolate tensors and oblate tensors by using a threshold $\tau = 0.25$.

In the case of prolate tensors ($c_l > \tau$), it can be assumed that the single dominant diffusion direction, \hat{e}_1 , is in the direction of the underlying fiber orientation. The modeling of our prior distribution for prolate tensors is based on the work of Anderson (2005), Alexander (2005) and Friman et al. (2006). Borrowing ideas from Anderson (2005), we assume that the prolate diffusion tensor within a voxel is a single axially symmetric tensor. Let us set up a local coordinate system for the diffusion tensor D so that it can be written as a diagonal matrix $\text{diag}(\lambda_{\perp}, \lambda_{\perp}, \lambda_{\parallel})$ where the axis corresponding to λ_{\parallel} is associated with the fiber orientation. In spherical coordinates, let θ be the polar angle from the λ_{\parallel} -axis and ψ be the azimuth angle from one of λ_{\perp} -axis. Then, a gradient direction \hat{g}_j with (θ, ψ) has local coordinates $g(\theta, \psi) = [\sin \theta \cos \psi, \sin \theta \sin \psi, \cos \theta]$. Thus, diffusion along \hat{g}_j is

$$\hat{g}_j^T D \hat{g}_j = g(\theta, \psi) \cdot \text{diag}(\lambda_{\perp}, \lambda_{\perp}, \lambda_{\parallel}) \cdot g(\theta, \psi)^T = \lambda_{\perp} + \cos^2 \theta \cdot (\lambda_{\parallel} - \lambda_{\perp}). \quad (18)$$

Suppose the trace of the diffusion tensor $\text{tr}(D)$ is known and varies insignificantly over the white matter. The mean of the three main diffusivities is $\bar{\lambda} = \text{tr}(D)/3$, and as a result $\lambda_{\parallel} = 3\bar{\lambda} - 2\lambda_{\perp}$. Let \hat{v} be the true fiber orientation, then $\cos \theta = \hat{v} \cdot \hat{g}_j$. Therefore, Eq. (18) can be written as

$$\hat{g}_j^T D \hat{g}_j = \lambda_{\perp} + 3(\hat{v} \cdot \hat{g}_j)^2 (\bar{\lambda} - \lambda_{\perp}), \quad (19)$$

which is equivalent to the constraint tensor model of Alexander (2005) and Friman et al. (2006). By inserting Eq. (19) into Eq. (16), we have

$$s_j = s_0 \exp[-b_j(\lambda_{\perp} + 3(\hat{v} \cdot \hat{g}_j)^2(\bar{\lambda} - \lambda_{\perp}))]. \quad (20)$$

Here, the intensity measured along any gradient direction is subject to two unknown parameters, i.e. \hat{v} and λ_{\perp} .

Due to noise, the intensity u_j in the DWI measured by the scanner is a noisy observation of the true signal s_j . It is well known that the noise in MRI can be described accurately by a Rician distribution Salvador et al. (2005). Salvador et al. (2005) showed that the error distribution conforms closely to a normal distribution with a zero-mean and standard deviation equal to the inverse of the signal-to-noise ratio (SNR), i.e. $\epsilon_j = \log(u_j) - \log(s_j) \sim N(0, \varrho_j^{-1})$, where $\varrho_j = s_j/\sigma_j$ is the SNR. Let the intensities observed at voxel i be $y_i = \{u_0, u_1, \dots, u_M\}$ where M is the number of gradient directions used in data acquisition. Following the idea of Friman et al. (2006), we setup the observation density at voxel i for prolate tensors by multiplying the error distribution over all gradient directions of the DWIs, i.e.

$$\begin{aligned} p(y_i|\hat{v}_i) &= \prod_{j=1}^M N(0, \varrho_j^{-1}) \\ &= \prod_{j=1}^M \frac{\varrho_j}{\sqrt{2\pi}} \exp\left[-\frac{\varrho_j^2(\log u_j - \log s_j)^2}{2}\right], \end{aligned} \quad (21)$$

where s_j is given by Eq. (20). To find the observation density of the variable \hat{v}_i for fiber orientations, we need to solve for three unknown parameters, i.e. $\bar{\lambda} = \text{tr}(D)/3$, λ_{\perp} and ϱ_j , in Eq. (21). The value of $\text{tr}(D)$ and λ_{\perp} can be estimated using the method in (Anderson, 2005). Here, we simply set $\text{tr}(D)$ to be the trace of the diffusion tensor D estimated using the linear least squares estimation (Basser et al., 1994), and set $\lambda_{\perp} = (\lambda_2 + \lambda_3)/2$. The SNR is estimated using the weighted least squares method in (Salvador et al., 2005). Panels (a) and (c) of Fig. 2 show two examples of the fiber orientation distribution calculated using Eq. (21). The figure shows that the orientation distribution of a prolate tensor is concentrated when its FA and c_1 are both relatively large.

In the case of oblate tensors ($c_1 \leq \tau$), the dominant direction of diffusion is ambiguous and Eq. (21) is inappropriate. It is possible that diffusion in the plane defined by \hat{e}_1 and \hat{e}_2 contains two or more significant non-collinear diffusion directions, each corresponding to a separate fiber tract. This situation may be indicative of fiber crossings and branchings, a well-known challenge in the analysis of DTI. In this case, we set up a local coordinate system with \hat{e}_3 as the Z axis and represent the fiber orientation \hat{v} in spherical coordinates. Let θ' be the polar angle from the \hat{e}_3 -axis, i.e. $\theta' = \arccos(\hat{v} \cdot \hat{e}_3)$, and ψ' be the azimuth angle (relative to an arbitrary reference direction in the plane spanned by \hat{e}_1 and \hat{e}_2). The vector \hat{v} is predominantly distributed on the plane spanned by \hat{e}_1 and \hat{e}_2 . Hence, we choose the distribution of the polar angle θ' to be normal with mean $\pi/2$ and standard deviation σ . The azimuth ψ' is assumed to have a uniform distribution over the interval

$[0, 2\pi]$. Thus, our fiber orientation distribution for oblate tensors is given by

$$p(y_i|\hat{v}) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{(\arccos(\hat{v} \cdot \hat{e}_3) - \pi/2)^2}{2\sigma^2}\right] \cdot \frac{1}{2\pi}. \quad (22)$$

Here, \hat{e}_3 is the eigenvector of the diffusion tensor D estimated using linear least squares. Panel (c) of Fig. 2 shows an example of the observation density of the fiber orientation for an oblate tensor in white matter.

3.2. Prior density

The state transition probability $p(\hat{v}_{i+1}|\hat{v}_i)$ specifies a prior distribution for the change in fiber direction between two successive steps. Here, we adopt a model of the prior density based on the von Mises–Fisher (vMF) distribution over a unit sphere. This is one of the simplest parametric distribution for directional data (Mardia and Jupp, 2000).

For a d -dimensional unit random vector \mathbf{x} , the probability density function for the vMF distribution is given by

$$f_d(\mathbf{x}; \mu, \kappa) = \frac{\kappa^{d/2-1}}{(2\pi)^{d/2} I_{d/2-1}(\kappa)} \exp(\kappa \mu^T \mathbf{x}), \quad (23)$$

where $\kappa \geq 0$, $\|\mu\| = 1$, and $I_{d/2-1}(\cdot)$ denotes the modified Bessel function of the first kind and order $d/2 - 1$. The density $f_d(\mathbf{x}; \mu, \kappa)$ is parameterised by the mean direction vector μ and the concentration parameter κ . The greater the value of κ the higher the concentration of the distribution around the mean direction μ . In particular, when $\kappa = 0$, the distribution is uniform over the sphere, and as $\kappa \rightarrow \infty$, the distribution tends to a point density. The distribution is rotationally symmetric around the mean μ , and is unimodal for $\kappa > 0$.

In our case, the directions are defined on a two directional unit sphere in R^3 , i.e. $d = 3$. We choose to model the prior state transition probability using the vMF distribution with mean \hat{v}_i and concentration parameter κ , i.e.

$$p(\hat{v}_{i+1}|\hat{v}_i) = f_3(\hat{v}_{i+1}; \hat{v}_i, \kappa). \quad (24)$$

The value of the concentration parameter κ controls the smoothness regularity of the tracked paths. The value of κ is set manually to optimally balance the prior constraints on smoothness against the evidence of \hat{v}_{i+1} observed from the image data. In our real-world experiments, we set $\kappa = 30$.

3.3. Importance density function

As pointed out by Doucet et al. (2000), the optimal importance density, which minimises the variance of the importance weight \tilde{w}_{i+1} conditional upon \hat{v}_{i+1} and \mathcal{Y} , is $p(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y})$. However, as noted




Fig. 2. Example of the observation density of three tensors from the brain white matter. (a) a prolate tensor with $FA = 0.9299$, $c_1 = 0.9193$, (b) a prolate tensor with $FA = 0.3737$, $c_1 = 0.3297$, (c) an oblate tensor $FA = 0.7115$, $c_1 = 0.2157$. For voxels with prolate tensors, the larger the value of FA and c_1 , the more focused the fiber orientation distribution. For voxels with oblate tensors, fiber orientations are focused on the plane spanned by the two leading eigenvector.

above, the optimal density suffers from two major drawbacks. In our case it is both difficult to sample from $p(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y})$ and to evaluate the integral analytically over the new state. Thus, our aim is to devise a suboptimal importance function that best represents $p(\mathcal{Y}_{i+1}|\hat{v}_{i+1})p(\hat{v}_{i+1}|\hat{v}_i)$ subject to the constraint that it can be sampled from in a straightforward manner.

A popular choice is to use the prior distribution as the importance function, i.e

$$\pi(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y}) = f_3(\hat{v}_{i+1}; \hat{v}_i, \kappa). \quad (25)$$

The von Mises–Fisher distribution in Eq. (24) can be efficiently sampled from using the simulation algorithm developed by Wood (1994). In this case, \hat{v}_{i+1} is predicted from \hat{v}_i and the importance weight is updated using

$$\tilde{w}_{t+1} = \tilde{w}_t p(\mathcal{Y}_{i+1}|\hat{v}_{i+1}). \quad (26)$$

However, the prior importance function is not particularly efficient. Since no observational information is used, the generated particles are often outliers of the posterior distribution. As a result, the weights may exhibit large variations and the results of estimation may be poor. Indeed, if the diffusion tensor at \hat{v}_i is prolate, then the movement to the state v_{i+1} is mainly attributable to the fiber orientation distribution. Thus, the posterior distribution is more strongly influenced by the observation density. For prolate tensors, we believe that the observation density in Eq. (21) is a good choice for the importance function. However, it is difficult to sample from such a density function.

To overcome this problem, we model the observation density in Eq. (21) using the von Mises–Fisher distribution. Since we use an axially symmetric tensor model, the distribution of fiber orientations in Eq. (21) is also rotationally symmetric around the direction of largest probability, as shown in Fig. 2. We therefore use the leading eigenvector, \hat{e}_1^i , of the diffusion tensor D_i at \hat{v}_i estimated using the linear least squares as the mean direction of the fiber orientation distribution. We have found experimentally that the leading eigenvector \hat{e}_1^i of D_i is almost identical to the direction of maximum probability for the distribution in Eq. (21). This is based on a test of 1000 prolate tensors from the brain MRI dataset described later in Section 4.2. The average difference between the two directions is less than 2° . Another issue we need to address is to select the concentration parameter v_i at each state \hat{v}_i . An accurate solution is to fit the von Mises–Fisher distribution to the observation distribution in Eq. (21) using the algorithm described in (Hill, 1981). However, this will significantly increase the computational complexity of the algorithm, which is one of the advantage of our algorithm over other methods such as those described in (Behrens et al., 2003) and (Friman et al., 2006). To overcome this problem, we sample a number of prolate tensors with different FA values from the MRI dataset. We then fit the concentration parameter v_i to the observation density in Eq. (21) for each of these tensors using Hill's algorithm (Hill, 1981). Fig. 3 shows the concentration parameter v_i as a function of the FA of prolate tensors. The figure reveals that the relationship between the concentration parameter and the tensor FA is empirically well described by using an exponential function, i.e.

$$v_i = \alpha + \exp\left(\frac{FA(D_i)^2}{\gamma^2}\right), \quad (27)$$

where α and γ are estimated from the above sampled fittings. Moreover, for particle filtering it is not necessary that the importance density is identical to the observation density. Therefore, for prolate tensors we set the importance density as

$$\pi(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y}) = f_3(\hat{v}_{i+1}; \hat{e}_1^i, v_i), \quad (28)$$

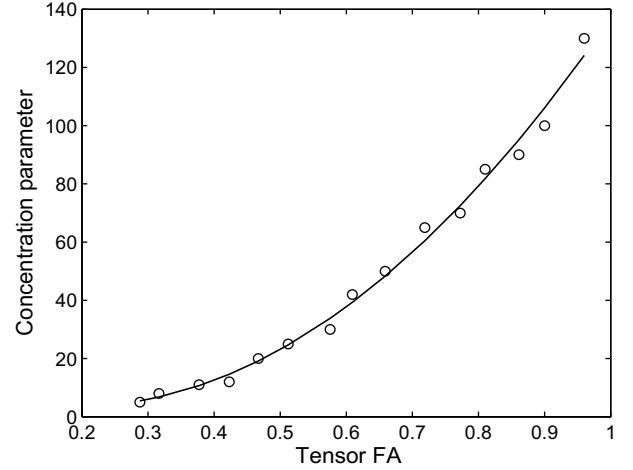


Fig. 3. Fitting concentration parameter as a function of FA value of tensors.

where v_i is estimated from FA using Eq. (27). For oblate tensors, since the observation density in Eq. (22) is wide, in this case we still use the prior given in Eq. (25) as the importance density.

3.4. Simulation of von Mises–Fisher distribution

The von Mises–Fisher distribution can be efficiently sampled from using the simulation algorithm developed by Ulrich (1984), which is further improved by Wood (1994). Ulrich's algorithm is designed for a general group of distributions on unit d -spheres, including the von Mises–Fisher distribution. Here, we discuss the ideas of Ulrich (1984) and Wood (1994) which are necessary for implementing our proposed tracking method. Ulrich (1984) observed that if we sample a unit vector V from the $(d-1)$ -dimensional sphere and sample a scalar random variable from the density

$$g(x) = \frac{(1-w^2)^{(m-3)/2} \exp(\kappa w)}{\sqrt{\pi(\kappa/2)^{d-1}}} I_{d/2-1} \Gamma\left(\frac{d-1}{2}\right), \quad (29)$$

then the concatenated unit vector $X = ((1-W)^{1/2}V^T, W)^T$ has a von Mises–Fisher distribution. Therefore, simulating the von Mises–Fisher distribution can be solved by simulating the density $g(x)$. However, it is still intractable to sample from $g(x)$ directly. To overcome this problem, the acceptance–rejection technique is used to develop an algorithm that makes samples from Eq. (29) by using the following envelop density

$$e(x, q) = \frac{2q^{(d-1)/2}}{B((d-1)/2, (d-1)/2)} \cdot \frac{(1-x^2)^{(d-3)/2}}{[(1+q) - (1-q)x]^{d-1}}, \quad (30)$$

where $B(\frac{d-1}{2}, \frac{d-1}{2}) = \frac{\Gamma(\frac{d-1}{2})}{\Gamma(\frac{d-1}{2})\Gamma(\frac{d-1}{2})}$ is the so-called Beta function. To maximise the acceptance ratio $g(x)/e(x, q)$, $q = (-2\kappa + \sqrt{4\kappa^2 + (d-1)^2})/(d-1)$. As a result, simulation of von Mises–Fisher distribution is achieved by sampling from $e(x, q)$ following the standard acceptance–rejection method. According to the Monte Carlo evaluation performed by Ulrich (1984), if generating a sample from a normal distribution takes CPU time t_0 , then in our case $d=3$ it requires CPU time $4t_0 \sim 7t_0$ to sample a unit vector from the vMF distribution. This is sufficiently efficient for us to sequentially sample fiber paths from the von Mises–Fisher distribution.

3.5. Algorithm outline

To summarise, the iteration steps of the algorithm are as follows:

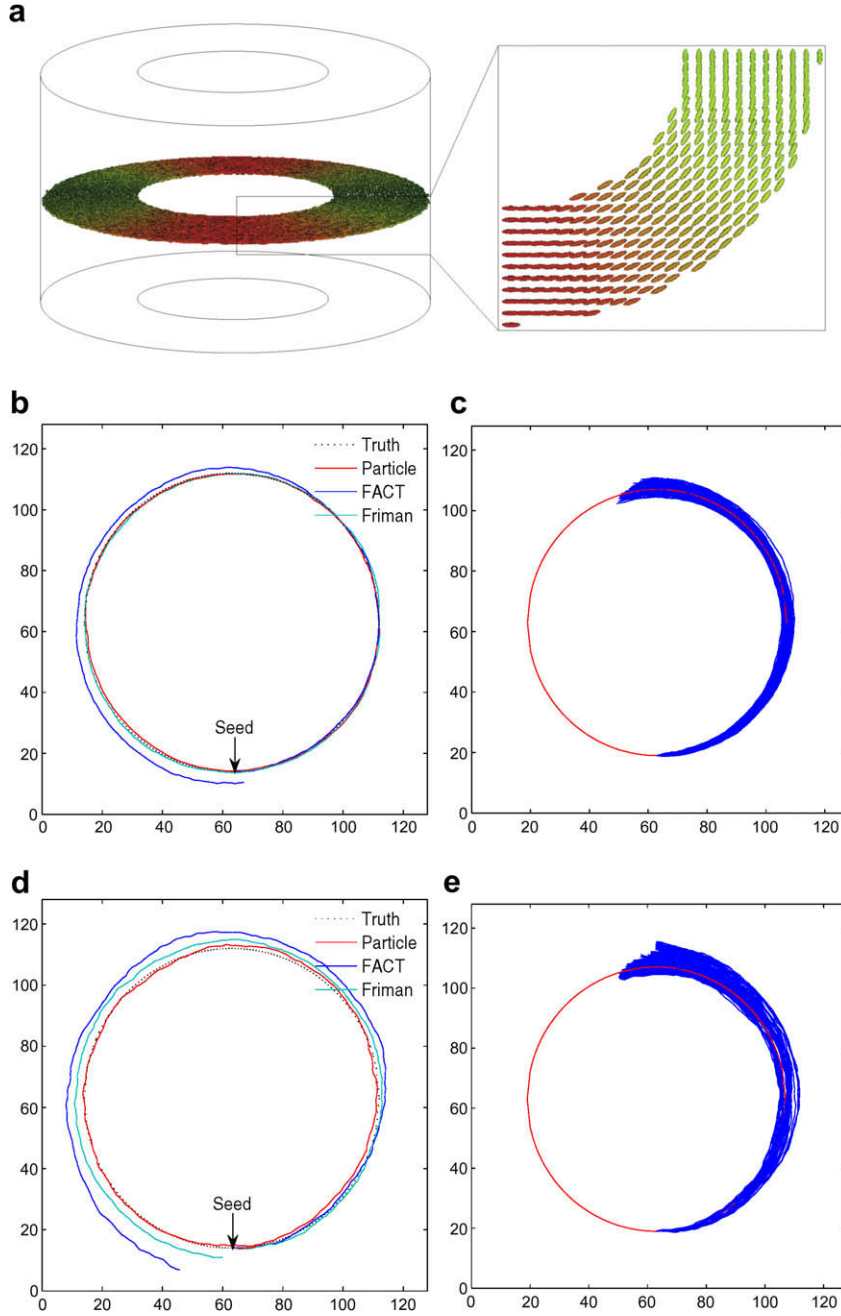


Fig. 4. (a) Synthetic data consisting of a cylinder and a sample slice with a zoomed view. (b) Tracked results of our method (MAP path), streamline method (FACT) (Mori et al., 1999) and the Friman method (Friman et al., 2006) under 10% noise. (c) A snapshot of the sampled paths of our method at propagation step 300 of case in (b). (d) Results under 25% noise. (e) A snapshot of the sampled paths of our method at propagation step 300 of case in (d).

- given K particles at step i : $\hat{v}_{i+1}^{(k)}, k = 1, \dots, K$
- compute diffusion tensor $D_i^{(k)}$ for each particle k using linear least square fitting
- **Prediction:** for $k = 1, \dots, K$
 - if $D_i^{(k)}$ is a prolate tensor, sample $\hat{v}_{i+1}^{(k)}$ according to Eq. (28)
 - if $D_i^{(k)}$ is an oblate tensor, sample $\hat{v}_{i+1}^{(k)}$ according to Eq. (25)
- **Weighting:** for $k = 1, \dots, K$
 - if prolate tensor, compute $\tilde{w}_{i+1}^{(k)}$ from Eq. (13) using Eqs. (21), (24) and (28)
 - if oblate tensor, compute $\tilde{w}_{i+1}^{(k)}$ from Eq. (13) using Eqs. (22), (24) and (25)
 - normalise all these weights
- **Selection:** evaluate N_{eff} using Eq. (14).
 - if $N_{\text{eff}} \geq N_s$, then for $k = 1, \dots, K$, $\hat{v}_{i+1}^{(k)} = \hat{v}_{i+1}^{(k)}$
 - if $N_{\text{eff}} < N_s$, then for $k = 1, \dots, K$, sample an index $z(k)$ from discrete distribution $\{\tilde{w}_{i+1}^{(k)}\}_{k=1, \dots, K}$, and set $\hat{v}_{i+1}^{(k)} = \hat{v}_{i+1}^{(z(k))}$, $\tilde{w}_{i+1}^{(k)} = 1/N$

4. Experimental results

We have evaluated our algorithm both on synthetic tensor fields and real-world MRI brain datasets. We have also qualitatively and quantitatively compared the results of our method with those obtained using the streamline method (Mori et al., 1999) and the probabilistic tracking method of Friman et al. (2006). Since our particles propagate in a continuous domain, an interpolation issue arises for diffusion data that is acquired only on a discrete grid. Here, we use the trilinear interpolation method introduced in (Zhukov and Barr, 2002). This method is computationally inexpensive and can preserve the positive-definiteness on the diffusion tensors.

4.1. Synthetic dataset

We commence by evaluating the performance of the algorithm on synthetic tensor fields. Each of the datasets used in this section contain $128 \times 128 \times 40$ voxels, have an in-plane resolution of 2×2 mm and a slice thickness of 2 mm. The procedure for generating synthetical additive noise is as follows. Suppose that the minimum and maximum image scalar values are u_l and u_h . We refer to the noise level as being $r\%$, if the standard deviation of the distribution from which the noise is sampled is $\sigma_r = \frac{r}{100}(u_h - u_l)$. For

each voxel, a noise value is sampled from an appropriate distribution (Rician in our case) with zero-mean and variance σ_r^2 , and then added to the intensity value of that voxel.

Our first example aims to qualitatively demonstrate the robustness of the algorithm under the influence of noise. To do this, we first generate a noise-free synthetic tensor field. The data contains a single cylinder, and the principal diffusion directions of the voxels within the cylinder form a concentric vector field, as shown in panel (a) of Fig. 4. Each voxel is visualised by an ellipsoid whose principal axes are the three orthogonal eigenvectors of the tensor,

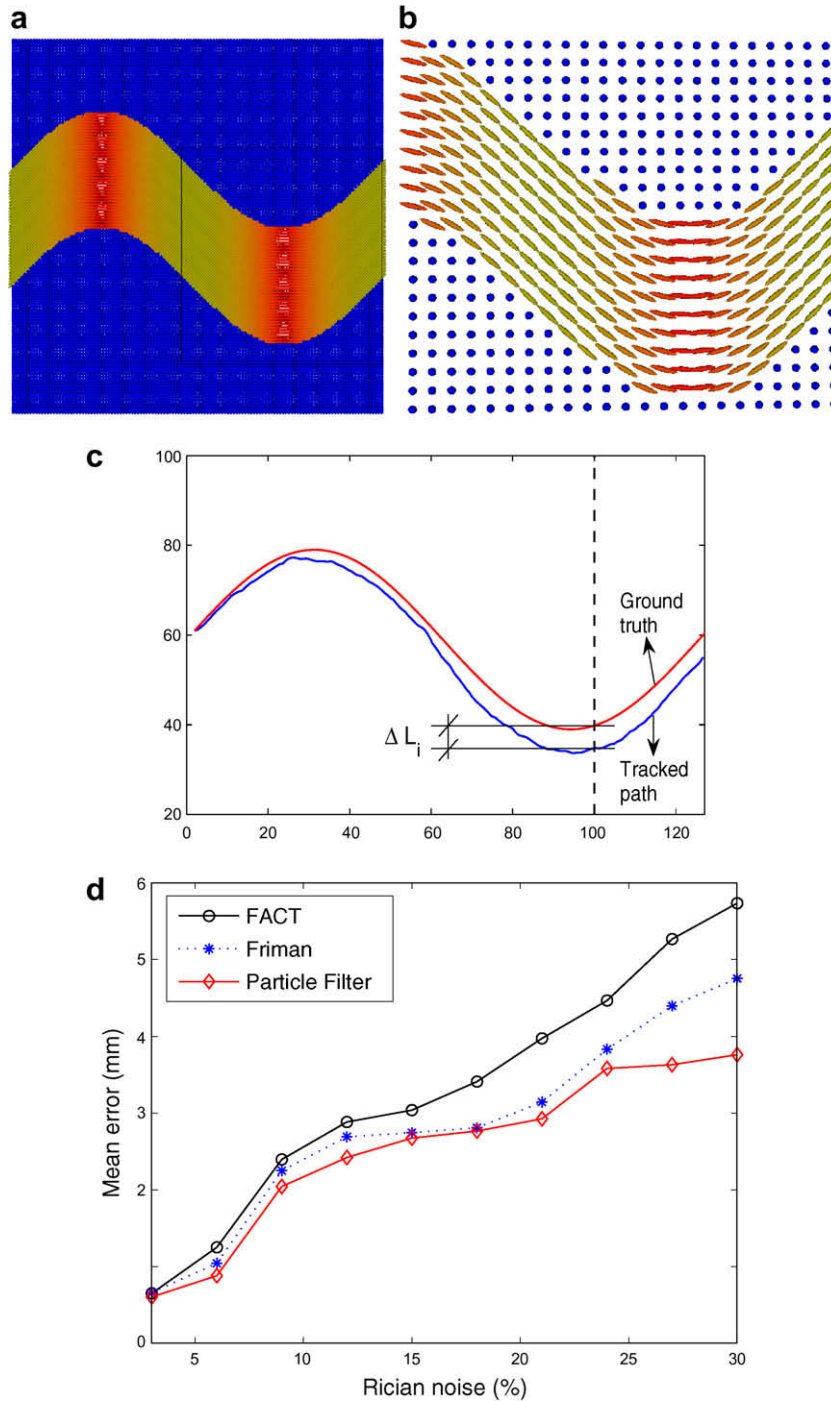


Fig. 5. (a) Synthetic ground truth fiber bundle. (b) A zoomed portion of the image in (a). (c) Illustration of error measurement between the ground truth fiber path and the reconstructed fiber path by fiber tracking algorithms. (d) Mean error comparison of the results of our method with those of FACT algorithm (Mori et al., 1999) and Friman method (Friman et al., 2006) under different levels of Rician noise.

and the radii of the ellipsoid along the axes are determined by the magnitude of the corresponding eigenvalues. Then, we add different levels of noise to the tensor field. The proposed particle filter algorithm is then used to track the global optimal fiber (MAP path) from a seed point using 1000 particles for 650 propagation steps with step size 1 mm. Our result is compared with that obtained using the standard local streamline method (FACT) (Mori et al., 1999) and the Bayesian method of Friman et al. (2006). For Friman's method, we sample 1000 paths commencing from the seed point using the reported discrete sampling technique with 2562 predefined directions on the unit sphere, and select the path with maximal probability as the optimal fiber. The resulting optimal fiber path for each of the three methods together with the ground truth path are shown in subfigures (b) and (d) of Fig. 4. Subfigures (c) and (e) show the trajectories of the particles obtained using our method at propagation step 300 for the results in subfigures (b) and (d), respectively. The figure shows that under relatively mild levels of noise (10% noise) both our method and Friman's method reconstruct the true fiber path quite well. However, our method runs significantly faster than Friman method. For instance, the MATLAB implementation of our method takes less than 100 seconds for 1000 samples to propagate for 100 steps on a PC with P4 CPU. The Friman method requires at least three times more to sample 1000 paths with the same length. Additionally, the MCMC method of Behrens et al. runs significantly more slowly according to their evaluation (Behrens et al., 2003). On the other hand, when the level of noise is large (25% noise), our method performs better than Friman's method as shown in Fig. 4c. This demonstrates that our algorithm samples paths more effectively due to the continuous simulation of the von Mises–Fisher distribution and the resampling step of particle filtering. The results also reveal that the streamline method FACT is sensitive to noise, and that it performs less accurately compared both to our method and Friman's method under low and high levels of noise.

In Fig. 5, we have quantitatively compared the performance of our method with that of two alternative methods (Mori et al., 1999; Friman et al., 2006) on an artificial fiber bundle. Here, we first construct a synthetic ground truth tensor field containing a curved fiber bundle with large curvatures at some locations, as shown in Fig. 5a. Fig. 5b shows a zoomed region of Fig. 5a to better visualise the details of the synthetic fibers. Here, we add noise to the synthetic image in the following way. We first generate a baseline image and six synthetic noise-free diffusion weighted images from the tensor image in Fig. 5a with gradient directions $[1, 0, 1]$, $[\bar{1}, 0, 1]$, $[0, 1, 1]$, $[0, 1, \bar{1}]$, $[1, 1, 0]$ and $[\bar{1}, 1, 0]$. Next, we add increasing levels of Rician noise to the baseline image and each of the six DWIs. We then apply our method, the streamline method (FACT) (Mori et al., 1999) and the Friman method (Friman et al., 2006) to the noisy datasets to reconstruct the fibers. We estimate the tensor image from the noisy DWIs using least squares fitting (Basser et al., 1994). For each method, we obtain the best possible results (in our case MAP path) by manually adjusting the relevant parameters. To evaluate the results, we sample a number of points on the ground truth fiber path, and compute the mean of distance error between the sampled ground truth points and the corresponding points on the reconstructed paths for each of the three methods. Fig. 5c illustrates how the error distance is measured between a sample point on the ground truth path and the corresponding point on the tracked path. Fig. 5d plots the mean distance error for each method as a function of the level of Rician noise. The figure reveals that our method achieves the smallest error at all levels of noise. Moreover, our method also exhibits a more reproducible behavior when the level of noise is severe. The figure also reveals that both our method and the Friman method (Friman et al., 2006) (the probabilistic methods) perform better than the streamline method (Mori et al., 1999). This is most evident at high levels of noise.

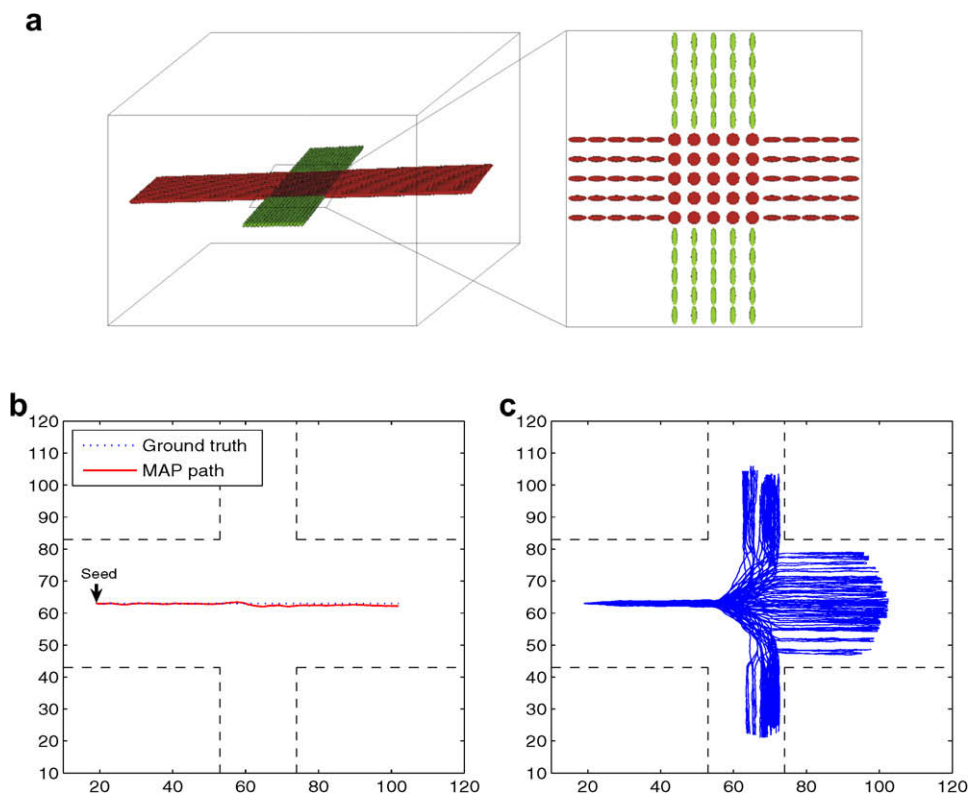


Fig. 6. Top: synthetic data with fiber crossing. Bottom: tracking result of our method.

In Fig. 6, we show the behavior of the algorithm under both noise and fiber crossings. We again generate a noise-free synthetic tensor field. In this case a right angle crossing between horizontal and vertical fibers is synthesized (as shown in the top row of Fig. 6). Additionally, we also add 5% Rician noise to the data set. In the crossing region, the first and second eigenvalues of the tensor are assumed to be equal. As a consequence, the diffusion tensors here are oblate ellipsoids. By contrast, the prolate tensors in regions without fiber crossing are elongated ellipsoids. We then apply our method to track the fiber from a seed by propagating 1000 particles for 200 steps. The globally optimal MAP paths for the particle trajectories are computed and visually compared with the ground truth fiber path in Fig. 6b. Although the principal eigenvector of the oblate tensors are not aligned with the fiber orientations, the result shows that our method still works fairly well under fiber

crossing. The algorithm can interpolate over gaps in the transition region, and allows the prior density to predominantly control the propagation of the particles in crossing regions. Subfigure (c) of Fig. 6 shows the trajectories for each of the particles at the final stage of propagation. The figure further reveals that our method is able to deal with multi-fiber crossings by propagating a proportion of the particles along each of the possible fiber branches. Since the aforementioned methods have not been specifically considered in the modeling of oblate tensors, we do not compare our result with the alternatives.

4.2. Brain diffusion MRI

Real-world diffusion MRI data was acquired from a healthy adult volunteer using a Siemens Allegra 3T head-only scanner. A

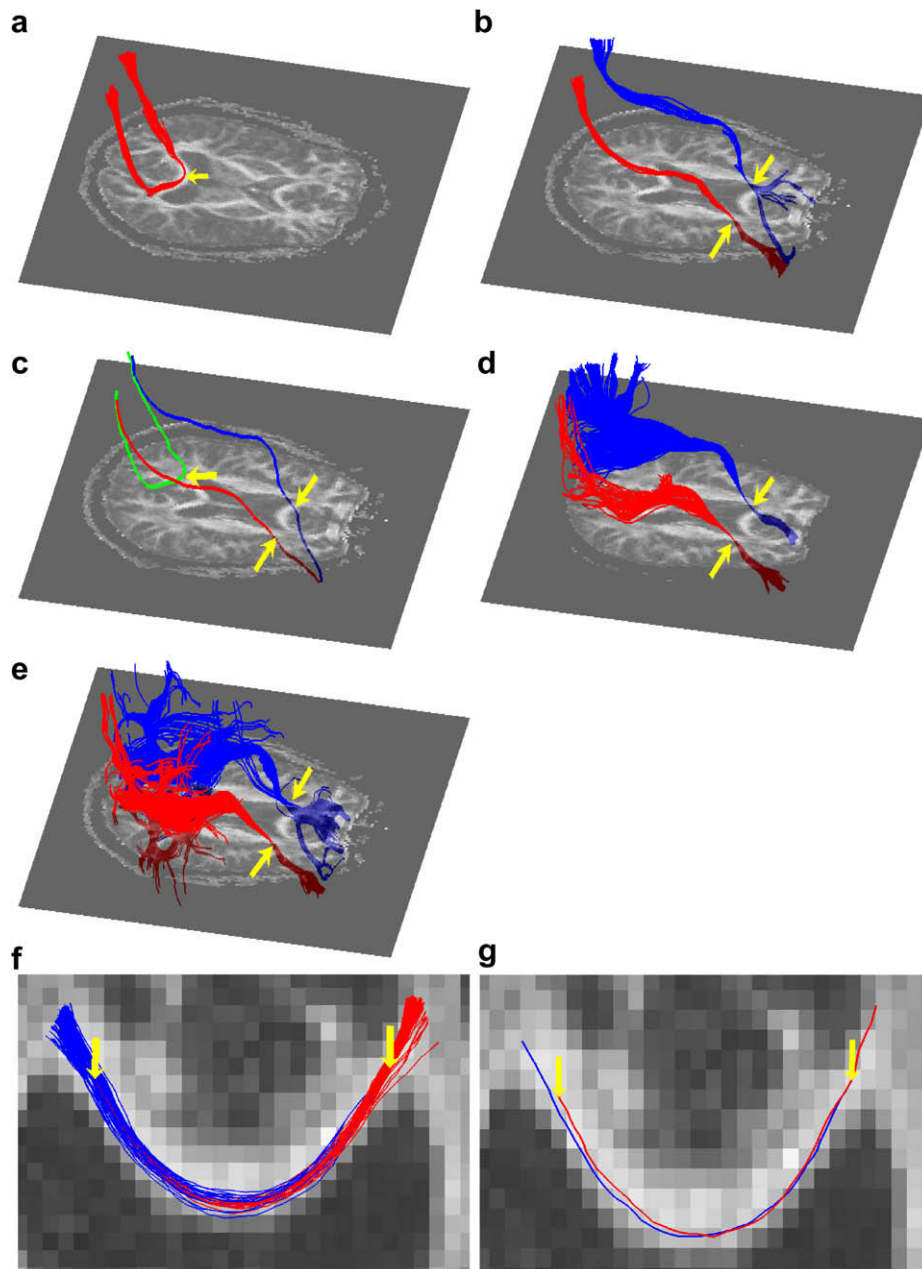


Fig. 7. (a) 1000 particle traces from a seed point in Corpus callosum, (b) from two seed points in the superior longitudinal fasciculus, (c) optimal MAP paths of (a) and (b), (d) 1000 path samples using the method in (Bjornemo et al., 2002) from the same seed points as in (b) with parameters $\alpha = 0.001$, $\beta = 80$, (e) 1000 path samples using Friman's method (Friman et al., 2006) from the same seed points as in (b), (f) zoomed particle traces of two seed points from the MAP path of example (a), (g) optimal MAP paths of (f).

$128 \times 128 \times 58$ volume image was acquired with $2 \times 2 \times 2$ mm voxel resolution. A six direction gradient scheme was used with 10 repetitions per-image, $b = 1000$ s/mm² for the gradient directions, and $b = 0$ s/mm² for the baseline image. Repetitions were aligned via rigid registration of the baseline images. A step length of 1 mm and 5000 particles were used for all examples. The propagation of a particle was halted when it exits white matter, characterised by a low FA value ($FA < 0.2$).

Fig. 7a shows the trajectories of 1000 particles seeded from a point in the Corpus callosum. Although we use 5000 particles for all examples, we show the trajectories of only 1000 particles in the results due to visualisation problems and the limited processing capability of our PC. However, the configurations of the trajectories of all 5000 particles are almost identical to those of the 1000 particles shown in the figures. The reason for this is that because of the re-sampling process used, there are many paths overlapping each other in the final step of tracking (this observation is based on our empirical evaluations of the results). Hence, we discard the repeated paths for visualisation purpose. Fig. 7a shows that the sampled paths provide a robust delineation of the expected fiber bundle. Fig. 7b shows an additional example with two seed points in the superior longitudinal fasciculus. This example reveals how our probabilistic algorithm is able to handle splitting fibers and ambiguous neighborhoods. Fig. 7c shows the global optimal MAP paths of the examples in Fig. 7a and b. We also compared our result with that obtained using our implementation of the method in (Bjornemo et al., 2002) and Friman's method based on the same seed points as shown in Fig. 7b. The distribution of the sampled paths of the method in (Bjornemo et al., 2002) is controlled by a regularisation parameter α and stochastic parameter β . The larger the stochastic parameter, the more dispersed the resulting sample paths. Fig. 7d shows 1000 paths sampled using

the method of (Bjornemo et al., 2002). One problem with the method is that there is no prior distribution for the local fiber orientations and the profile of the sampled paths are empirically controlled by a stochastic parameter. By contrast, both our method and Friman's method aim to locate the true posterior distribution of the fibers from reliable prior distributions of the local fiber orientations. In our method, particles with low probability of existence are eliminated during the resampling stage, and the sampled paths are most concentrated around the final optimal fiber. Fig. 7e shows 1000 sample paths using Friman's method. The figure shows that the sampled paths from Friman's method are more dispersed, with a number of paths which have low probabilities. Moreover, our algorithm runs much faster than Friman's algorithm. To further evaluate the algorithm, we select two seed points from the MAP path of the example in Fig. 7a and let the algorithm track from one to the other. Fig. 7e shows 1000 sample paths from each seed point. The figure shows that the sampled paths from the two seed points are almost overlapping with each other. Fig. 7f shows their two optimal MAP paths, which are very close to each other. Thus, the second seed point can successfully return to the first one along the MAP path. This example shows that the performance of our algorithm is robust and stable.

On the other hand, based on the particle traces, we can calculate the probability of connection between the seed voxel and a specific voxel by computing the fraction of particles passing through that voxel. We can thus produce a probability map of connections between the seed and all remaining voxels. In Fig. 8a, we show the probability map from a seed point in the Corpus callosum. The coloring shows the likelihood of paths (connecting the seed voxel and each of the remaining voxels) generated by our algorithm. Fig. 8b gives a probability map of longer fiber tracts of two seed points shown in Fig. 8b. The result here is compared to that of the method




Fig. 8. Probability map of our algorithm from (a) a seed point in the Corpus callosum and (b) from two seed points in the superior longitudinal fasciculus. (c) Probability map of the method in (Bjornemo et al., 2002) from the same seed points as in (b). (d) Probability map of Friman's method from the same seed points as in (b).

in (Bjornemo et al., 2002), as shown in Fig. 8c, and Friman's method, as shown in Fig. 8d, which gives a wider distribution.

5. Conclusion

We have presented a new method for probabilistic white matter fiber tracking. The global tracking model is formulated using a state space framework, which is implemented by applying particle filtering to recursively estimate the posterior distribution of fibers and to locate the global optimal fiber path. Each ingredient of the tracking algorithm is detailed. For modeling the fiber orientation distribution, we classify voxels of the white matter as either prolate or oblate tensors. For prolate tensors, the orientation distribution is theoretically formulated by combining the axially symmetric tensor model with a noise model for DWI. For oblate tensors, the orientation distribution is computed using a normal distribution of angles between fiber orientations and the smallest eigenvectors of the tensors. Fast and efficient sampling is realised using the von Mises–Fisher distribution. As a consequence, there is no need to apply MCMC sampling (Behrens et al., 2003) or to discretise the state space (Friman et al., 2006) to sample paths from the fiber orientation distribution.

Based on our experimental evaluations, the advantages of the proposed algorithm are threefold. First, unlike previous methods which are computationally very expensive, our method shows improved computational efficiency and is able to rapidly locate the global optimal fiber and to compute the connectivity map for the seed point. Second, our method can more accurately reconstruct the true fiber paths in very noisy images. It gives smaller errors between the reconstructed path and the true fiber path. Third, the proposed method is able to deal with fiber crossings because we separately model the orientation distributions for different shapes of diffusion tensors.

However, there are several ways in which the method can be further improved. For instance, our model of the orientation distribution for oblate tensors is not accurate enough to capture the complex configurations at fiber crossings. Thus, there are large uncertainties at such voxels. More sophisticated methods will have to be developed for dealing with fiber crossings in our future work. One possibility is the use of mixtures of von Mises–Fisher distributions (McGraw et al., 2006). To validate the method in a quantitative way, real-world datasets with annotated ground truth are needed. Developing simulated and real DTI sets with ground truth for validation and across-site comparison is an active area of imaging research, but to our knowledge these are not yet available to the community.

Acknowledgements

The co-authors Casey Goodlet and Guido Gerig are co-investigators of the National Alliance for Medical Image Computing (NA-MIC), funded by the National Institutes of Health through Grant U54 EB005149. Information on the National Centers for Biomedical Computing can be obtained from <http://nihroadmap.nih.gov/bioinformatics>. We also acknowledge support from the NIMH Silvio Conte Center for Neuroscience of Mental Disorders MH064065.

References

Alexander, D., 2005. An introduction to computational diffusion MRI: the diffusion tensor and beyond. In: Weickert, J., Hagen, H. (Eds.), *Visualization and Processing of Tensor Fields*. Springer.

Alexander, A., Hasan, K., Lazar, M., Tsuruda, J., Parker, D., 2001. Analysis of partial volume effects in diffusion-tensor MRI. *Magn. Reson. Med.* 45, 770–780.

Anderson, A., 2005. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magn. Reson. Med.* 54, 1194–1206.

Basser, P., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson., Ser. B* 111, 209–219.

Basser, P., Mattiello, J., LeBihan, D., 1994. MR diffusion tensor spectroscopy and imaging. *Biophys. J.* 66, 259–267.

Basser, P., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. *Magn. Reson. Med.* 44, 625–632.

Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Brady, J., Smith, S., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.* 50, 1077–1088.

Behrens, T., Berg, H.J., Jbabdi, S., Rushworth, M., Woolrich, M., 2007. Probabilistic diffusion tractography with multiple fibre orientations: what can we gain? *NeuroImage* 34, 144–155.

Bhalerao, A., Westin, C., 2007. Hyperspherical von Mises–Fisher Mixture (HvMF) modelling of high angular resolution diffusion MRI. In: *Proc. MICCAI*, pp. 236–243.

Bjornemo, M., Brun, A., Kikinis, R., Westin, C.F., 2002. Regularized stochastic white matter tractography using diffusion tensor MRI. In: *Proc. MICCAI*, pp. 435–442.

Brun, A., Bjornemo, M., Kikinis, R., Westin, C.F., 2002. White matter tractography using sequential importance sampling. In: *Proc. ISMRM Annual Meeting*, pp. 1131.

Cook, P., Alexander, D., Parker, G., 2004. Modelling noise-induced fibre-orientation error in diffusion-tensor MRI. In: *Proc. IEEE International Symposium on Biomedical Imaging: Nano to Macro*, pp. 332–335.

Doucet, A., Godsill, S., Andrieu, C., 2000. On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.* 10, 197–208.

Doucet, A., de Freitas, N., Gordon, N., 2001. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag.

Fletcher, P., Tao, R., Jeong, W., Whitaker, R.T., 2007. A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In: *Proc. Information Processing in Medical Imaging*, pp. 346–358.

Frank, L., 2001. Anisotropy in high angular resolution diffusion-weighted MRI. *Magn. Reson. Med.* 45, 935–939.

Friman, O., Farneback, G., Westin, C., 2006. A bayesian approach for stochastic white matter tractography. *IEEE Trans. Med. Imag.* 25 (8), 965–978.

Geman, D., Jedynak, B., 1996. An active testing model for tracking roads in satellite images. *IEEE Trans. Pattern Anal. Mach. Intell.* 18 (1), 1–14.

Gordon, N., Salmond, D., Smith, A., 1993. Novel approach to nonlinear non-Gaussian Bayesian state estimation. *Radar Signal Process. IEE Proc. F* 140 (2), 107–113.

Gossl, C., Fahrmeir, L., Putz, B., Auer, L., Auer, D., 2002. Fiber tracking from DTI using linear state space models: detectability of the pyramidal tract. *NeuroImage* 16, 378–388.

Hill, G., 1981. Algorithm 571: statistics for von Mises' and Fisher's distributions of directions. *ACM Trans. Math. Software* 7, 233–238.

Isard, M., Blake, A., 1998. Condensation-conditional density propagation for visual tracking. *Int. J. Comput. Vision* 29 (1), 5–28.

Jones, D., 2003. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. *Magn. Reson. Med.* 49, 7–12.

Lazar, M., Alexander, A., 2005. Bootstrap white matter tractography (BOOT-TRAC). *NeuroImage* 24, 524–532.

Lazar, M., Weinstein, D., Tsuruda, J., Hasan, K., Arfanakis, K., Meyerand, E., Badie, B., Rowley, H., Houghton, V., Field, A., Witwer, B., Alexander, A., 2003. White matter tractography using tensor deflection. *Hum. Brain Mapp.* 18, 306–321.

Liu, J., 1996. Metropolisized independent sampling with comparison to rejection sampling and importance sampling. *Stat. Comput.* 6, 113–119.

Liu, J., Chen, R., Logvinenko, T., 2001. *A Theoretical Framework for Sequential Importance Sampling and Resampling*. Sequential Monte Carlo Methods in Practice, New York.

Macovski, A., 1996. Noise in MRI. *Magn. Reson. Med.* 36, 494–497.

Mardia, K.V., Jupp, P.E., 2000. *Directional Statistics*. Wiley.

McGraw, T., Vemuri, B., Yeziarski, R., Mareci, T., 2006. Segmentation of high angular resolution diffusion MRI modeled as a field of von Mises–Fisher mixtures. In: *Proc. ECCV*, pp. 463–475.

Mori, S., van Zijl, P., 2002. Fiber tracking: principles and strategies – a technical review. *NMR Biomed.* 15, 468–480.

Mori, S., Crain, B., Chacko, V., van Zijl, P., 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann. Neurol.* 5, 265–269.

Parker, G., 2004. Analysis of MR diffusion weighted images. *Brit. J. Radiol.* 77, 176–185.

Parker, G., Alexander, D., 2003. Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. In: *Proc. Information Processing in Medical Imaging*, pp. 684–695.

Parker, G., Alexander, D., 2005. Probabilistic anatomical connectivity using persistent angular structure obtained from diffusion weighted imaging. *Phil. Trans. Roy. Soc. Ser. B* 360, 893–902.

Parker, G., Wheeler-Kingshott, C., Barker, G., 2002. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Trans. Med. Imag.* 21 (5), 505–512.

Parker, G., Haroon, H., Wheeler-Kingshott, C., 2003. A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *J. Magn. Reson. Imag.* 18, 242–254.

Perez, P., Blake, A., Gangnet, M., 2001. JetStream: probabilistic contour extraction with particles. In: *Proc. IEEE ICCV*, pp. 524–531.

Pichon, E., Westin, C., Tannenbaum, A., 2005. A Hamilton–Jacobi–Bellman approach to high angular resolution diffusion tractography. In: *Proc. MICCAI*, pp. 180–187.

- Prados, E., Lenglet, C., Pons, J., Wotawa, N., Deriche, R., Faugeras, O., Soatto, S., 2006. Control theory and fast marching techniques for brain connectivity mapping. In: Proc. CVPR, pp. 1076–1083.
- Salvador, R., Pena, A., Menon, D., Carpenter, T., Pickard, J., Bullmore, E., 2005. Formal characterization and extension of the linearized diffusion tensor model. *Hum. Brain Mapp.* 24, 144–155.
- Staempfli, P., Jaermann, T., Crelier, G., Kollias, S., Valavanis, A., Boesiger, P., 2006. Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. *NeuroImage* 30, 110–120.
- Tournier, J., Calamante, F., Gadian, D., Connelly, A., 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage* 23, 1176–1185.
- Tuch, D., 2004. Q-ball imaging. *Magn. Reson. Med.* 52, 1358–1372.
- Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., Wedeen, V., 2002. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn. Reson. Med.* 48 (4), 577–582.
- Ulrich, G., 1984. Computer generation of distributions on the m -sphere. *Appl. Stat.* 33, 158–163.
- Westin, C., Maier, S., Mamata, H., Nabavi, A., Jolesz, F., Kikinis, R., 2002. Processing and visualization for diffusion tensor MRI. *Med. Image Anal.* 6, 93–108.
- Wood, A.T.A., 1994. Simulation of the von Mises–Fisher distribution. *Commun. Stat. Simul. Comput.* 23, 157–164.
- Zhang, F., Goodlett, C., Hancock, E.R., Gerig, G., 2008. Probabilistic fiber tracking using particle filtering. In: Proc. MICCAI, pp. 144–152.
- Zhukov, L., Barr, A., 2002. Oriented tensor reconstruction: tracing neural pathways from diffusion tensor MRI. In: Proc. IEEE Visualization, pp. 387–394.