

Probabilistic Fiber Tracking Using Particle Filtering

Fan Zhang¹, Casey Goodlett², Edwin Hancock¹, and Guido Gerig²

¹ Dept. of Computer Science, University of York, York, YO10 5DD, UK

² Dept. of Computer Science, University of North Carolina at Chapel Hill, USA

Abstract. This paper presents a novel and fast probabilistic method for white matter fiber tracking from diffusion weighted MRI (DWI). We formulate fiber tracking on a nonlinear state space model which is able to capture both smoothness regularity of fibers and uncertainties of the local fiber orientations due to noise and partial volume effects. The global tracking model is implemented using particle filtering, which allows us to recursively compute the posterior distribution of the potential fibers. The fiber orientation distribution is theoretically formulated for prolate and oblate tensors separately. Fast and efficient sampling is realised using the von Mises-Fisher distribution on unit spheres. Given a seed point, the method is able to rapidly locate the global optimal fiber and also provide a connectivity map. The proposed method is demonstrated on a brain dataset.

1 Introduction

White matter fiber tracking or "tractography" estimates possible fiber paths by tracing the local fiber orientations. However, the local fiber orientations measured by DTI are not completely reliable due to both noise, partial volume effects and crossing fibers. To deal with this uncertainty, probabilistic fiber tracking methods have received considerable interest recently [1,2,3]. Instead of reconstructing the fiber pathways, they aim to measure the probability of connectivity between brain regions. These methods can be described in terms of two steps. Firstly, they model the uncertainty in fiber orientation measurements at each voxel using a probability density function (PDF) [1,3]. Secondly, the probabilistic tracking algorithms simply repeat a streamline propagation process 1000 ~ 10000 times with propagation directions randomly sampled from the PDF. The fraction of the streamlines that pass through a voxel provides an index of the strength of connectivity between that voxel and the starting point. Most previous methods estimate the connectivity map by sampling directly from the PDF for fiber orientations. The sampling process is challenging, thus it is necessary to resort to MCMC methods [1] or to evaluate the PDF discretely with low angular resolution [3]. The intrinsic drawback of the previous methods is their computational complexity (often more than several hours on a modern PC [1,3]), and this is unacceptable in practice.

In this paper, we propose a new probabilistic method for white matter fiber tracking. Our contributions are threefold: First, we formulate fiber tracking using a nonlinear state space model and recursively compute the posterior distribution using particle filtering [4]. The proposed model can capture both smoothness regularity of the fibers and the uncertainties of the local fiber orientations. Our second contribution concerns the PDF modeling of local fiber orientations from DTI. To do so, we build PDFs for prolate and

oblate tensors separately. For prolate tensors, we characterise the uncertainty using the axially symmetric model [5]. For oblate tensors, we model the PDF using the normal distribution of the angle with the smallest eigenvectors of tensors. Finally, we use the von Mises-Fisher distribution to model prior and the importance density for sampling particles. This spherical distribution provides a natural way to model noisy directional data, and it can be efficiently sampled using the Wood's simulation algorithm [6].

2 Tracking Algorithm

The problem of fiber tracking is to extract the best possible fiber pathway from a predefined seed point. We formulate the problem using a nonlinear state space model.

2.1 Global Tracking Model

A white matter fiber path can be modeled 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 image data, and the data observed at \hat{v}_i is $y_i = \mathcal{Y}(\hat{v}_i) = \mathcal{Y}(x_i)$. Our goal is to propagate a sequence of unit vectors that best estimates the true fiber based on prior density $p(\hat{v}_{i+1}|\hat{v}_{0:i})$ and the observation model $p(\mathcal{Y}|\hat{v}_{0:i})$.

We assume that the tracking dynamics forms a Markov chain, so that $p(\hat{v}_{i+1}|\hat{v}_{0:i}) = p(\hat{v}_{i+1}|\hat{v}_i)$. Thus, the prior of the fiber path is $p(\hat{v}_{0:n}) = p(\hat{v}_0) \prod_{i=1}^n p(\hat{v}_i|\hat{v}_{i-1})$. Another assumption is that the diffusion measurements are conditionally independent given $\hat{v}_{0:n}$, i.e. $p(\mathcal{Y}|\hat{v}_{0:n}) = \prod_{r \in \Omega} p(\mathcal{Y}(r)|\hat{v}_{0:n})$. We also assume that the measurement at a point does not depend on any points in the history of the path, i.e. $p(y_i|\hat{v}_{0:i}) = p(y_i|\hat{v}_i)$. Using the prior $p(\hat{v}_{0:n})$, the posterior distribution $p(\hat{v}_{0:n}|\mathcal{Y})$ can be expanded to

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

Applying Bayes theorem, we have

$$p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y}) = \frac{p(y_i|\hat{v}_i)p(\hat{v}_i|\hat{v}_{i-1})}{p(y_i)}, \quad (2)$$

where $p(y_i)$ is a constant regularity factor, i.e. $p(y_i) = \int_{\hat{v}_i} p(y_i|\hat{v}_i)p(\hat{v}_i|\hat{v}_{i-1})$. Most previous probabilistic methods [2,1] estimate the posterior $p(\hat{v}_{0:n}|\mathcal{Y})$ by sampling streamline paths from $p(y_i|\hat{v}_i)$. The sampling is difficult and time consuming. Moreover, it does often not take into account the smoothness constraint for fibers. In contrary, Friman et al. [3] estimate the posterior by sampling from $p(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y})$. To avoid sampling difficulties, they discretise the problem using a finite set of directions. In addition to introducing errors, this discretised sampling is still very time consuming. Moreover, some simple sampling methods may degenerate as path becomes longer [4].

2.2 Recursive Posterior Using Particle Filtering

We wish to estimate the posterior distribution iteratively in time. By inserting Equation (2) into Equation (1), we have

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

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. We recast the problem of tracking the expected fiber path as that of approximating the MAP path from the posterior distribution. It is straightforward to obtain the following recursive formula for the posterior from Equation (3)

$$p(\hat{v}_{0:i+1}|\mathcal{Y}) = p(\hat{v}_{0:i}|\mathcal{Y}) \frac{p(\hat{v}_{i+1}|\hat{v}_i)p(y_{i+1}|\hat{v}_{i+1})}{p(y_{i+1})}. \quad (4)$$

Since the denominator contains a complex high-dimensional integral, it is not feasible to locate the maximum likelihood path analytically. Like methods discussed above, we evaluate the posterior using a large number of samples which efficiently characterise the posterior. Thus, the statistical quantities, such as the mean, variance and maximum likelihood, can be approximated based on the sample set. Since it is seldom possible to obtain samples from the posterior directly, we use the particle filtering to recursively compute a finite set of sample paths from the posterior based on the Equation (4).

To sample a set of K paths, we set K particles at the starting location 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*. 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 [4] $\pi(\hat{v}_{0:n}|\mathcal{Y}) = \pi(\hat{v}_0|\mathcal{Y}) \prod_{i=1}^n \pi(\hat{v}_i|\hat{v}_{i-1}, \mathcal{Y})$. In the first *prediction* stage, each 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 generally is not an efficient approximation of the posterior distribution at time $i + 1$. Thus, in the second *weighting* stage, we measure the reliability of the approximation using a ratio, referred to as the importance weight, between the truth and the approximation $w_{i+1}^{(k)} = p(\hat{v}_{0:i+1}^{(k)}|\mathcal{Y}) / (\pi(\hat{v}_{0:i}|\mathcal{Y})\pi(\hat{v}_{i+1}^{(k)}|\hat{v}_i^{(k)}, \mathcal{Y}))$. We are more interested in the normalised weights $\tilde{w}_{i+1}^{(k)} = w_{i+1}^{(k)} / \sum_{l=1}^K w_{i+1}^{(l)}$. Inserting Equation (4) and expression of $w_{i+1}^{(k)}$ into the expression of $\tilde{w}_{i+1}^{(k)}$, it goes as

$$\tilde{w}_{i+1}^{(k)} \propto \tilde{w}_i^{(k)} \frac{p(\hat{v}_{i+1}^{(k)}|\hat{v}_i^{(k)})p(y_{i+1}|\hat{v}_{i+1}^{(k)})}{\pi(\hat{v}_{i+1}^{(k)}|\hat{v}_i^{(k)}, \mathcal{Y})}. \quad (5)$$

The choice of importance function will be detailed in the next section. 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 more and more skewed as time increases. The purpose of the last *selection* stage is to avoid this degeneracy. We measure the degeneracy of the algorithm using the effective sample size N_{eff} [4], i.e. $N_{eff} = 1 / \sum_{k=1}^K (\tilde{w}_{i+1}^{(k)})^2$. When N_{eff} is below a fixed threshold N_s , then a resampling procedure is used. 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 can be easily solved based on the discrete distribution of the posterior. The MAP estimate of the true fiber is the path with the maximal importance weight. The probability of connectivity between x_0 and a specific voxel is computed as the fraction of particles that pass through that voxel.

3 Algorithm Ingredients

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

3.1 Observation Density

Let $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$ be the decreasing eigenvalues of D and $\hat{e}_1, \hat{e}_2, \hat{e}_3$ be the corresponding eigenvectors. We can classify the diffusion tensors into prolate tensors ($\lambda_1 > (\lambda_2 \approx \lambda_3)$) and oblate tensors ($(\lambda_1 \approx \lambda_2) < \lambda_3$) by using fractional anisotropy (FA) [7] and the metric proposed by Westin et al. [8], i.e. $c_l = (\lambda_1 - \lambda_2) / \sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}$.

In the case of prolate tensors, we assume that a single dominant diffusion direction, \hat{e}_1 , is collinear with the true underlying fiber orientation \hat{v} . Borrowing ideas from [5], we suppose the prolate tensor is axially-symmetric. Let λ_\perp be the diffusivity of directions perpendicular to \hat{v} . Then, the diffusion along a gradient direction \hat{g}_j can be written as [5] $\hat{g}_j^T D \hat{g}_j = \lambda_\perp + 3(\hat{v} \cdot \hat{g}_j)^2 (\bar{\lambda} - \lambda_\perp)$, where $\bar{\lambda} = \text{tr}(D)/3$. By inserting this expression into the Stejskal-Tanner equation [7], we have $s_j = s_0 e^{-b_j(\lambda_\perp + 3(\hat{v} \cdot \hat{g}_j)^2 (\bar{\lambda} - \lambda_\perp))}$, where b_j is the scanner parameter and s_0 is the intensity of the baseline image. Due to noise, the intensity u_j measured by the scanner is a noisy observation of the true intensity s_j . In [9], Salvador et al. showed $\epsilon_j = \log(u_j) - \log(s_j) \sim N(0, \varrho_j^{-1})$, where $\varrho_j = s_j/\sigma_j$ is the signal-to-noise ratio (SNR). Let the intensities observed at i be $y_i = \{u_0, u_1, \dots, u_M\}$. Then the observation density for prolate tensors is given as

$$p(y_i | \hat{v}_i) = \prod_{j=1}^M \frac{\varrho_j}{\sqrt{2\pi}} e^{-\frac{\varrho_j^2 (\log u_j - \log s_j)^2}{2}}. \quad (6)$$

Panels (a) and (c) of Fig. 1 show two examples. The figure tells that the orientation distribution is very concentrated when its FA and c_l are relatively large.

In the case of oblate tensors, the dominant direction of diffusion is ambiguous and Equation (6) is inappropriate. It is possible that the plane defined by \hat{e}_1 and \hat{e}_2 contains




Fig. 1. Examples of observation density. (a) a prolate tensor $FA = 0.9299$, $c_l = 0.9193$. (b) a prolate tensor $FA = 0.3737$, $c_l = 0.3297$. (c) an oblate tensor $FA = 0.7115$, $c_l = 0.2157$.

several crossing fiber tracts. In this case, we 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. The vector \hat{v} is mainly distributed on the plane spanned by \hat{e}_1 and \hat{e}_2 . Hence, we choose the distribution of θ' to be normal with mean $\pi/2$ and standard deviation σ . The azimuth ψ' is assumed to have a uniform distribution over $[0, 2\pi]$. Thus, our fiber orientation distribution for oblate tensors is given by

$$p(y_i|\hat{v}) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\arccos(\hat{v} \cdot \hat{e}_3) - \pi/2)^2}{2\sigma^2}\right) \cdot \frac{1}{2\pi}. \quad (7)$$

Panel (c) of Fig. 1 shows an example of the density of an oblate tensor in white matter.

3.2 Prior Density

The transition density $p(\hat{v}_{i+1}|\hat{v}_i)$ 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 [10] over a unit sphere. For a d -dimensional unit random vector \mathbf{x} , the vMF distribution is given by

$$f_d(\mathbf{x}; \mu, \kappa) = \frac{\kappa^{d/2-1}}{(2\pi)^{d/2} I_{d/2-1}(\kappa)} \exp(\kappa \mu^T \mathbf{x}), \quad (8)$$

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 . 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 dimensional unit sphere in R^3 , i.e. $d = 3$. Thus, we choose our prior density as the vMF distribution with mean \hat{v}_i and concentration parameter κ , i.e.

$$p(\hat{v}_{i+1}|\hat{v}_i) = f_3(\hat{v}_{i+1}; \hat{v}_i, \kappa). \quad (9)$$

The value of the concentration parameter κ here controls the smoothness regularity of the tracked paths. It is set manually to optimally balance the prior constraints on smoothness against the evidence of v_{i+1} observed from the image data.

3.3 Importance Density Function

As discussed in Doucet et al. [4], the optimal importance density is $p(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y})$. However, it is difficult for sampling. Thus, our aim is to devise a suboptimal importance function that best represents $p(y_{i+1}|\hat{v}_{i+1})p(\hat{v}_{i+1}|\hat{v}_i)$ subject to the constraint that it can be sampled from. A most popular choice is to use the prior distribution as the importance function, i.e

$$\pi(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y}) = f_3(\hat{v}_{i+1}; \hat{v}_i, \kappa). \quad (10)$$

The vMF distribution in Equation (9) can be efficiently sampled from using the simulation algorithm developed by Wood [6]. However, the prior importance function is

not very efficient. Since no observation information is used, the generated particles are often outliers of the posterior distributions. 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, which is difficult to sample from. To overcome this problem, we model the observation density in Equation (6) using the vMF distribution. Since we use an axially symmetric tensor model, the distribution in Equation (6) is also rotationally symmetric around the direction of largest probability (see Fig. 1). We thus use the leading eigenvector, \hat{e}_1^i , of tensor D_i as the mean direction. We have found experimentally that \hat{e}_1^i is almost identical to the direction of maximum probability in Equation (6). The average difference between them is less than 2° due to a test on 1000 prolate tensors from brain MRI data. The concentration parameter ν_i at each state \hat{v}_i is set to $\nu_i = 90 \times c_l(D_i)$. This choice is based on empirical trial and error. A better way is to fit the vMF distribution to Equation (6) using the algorithm presented in [11]. However, it would need more computation time. Moreover, particle filtering requires an importance density that is close but not necessarily identical to the observation density. Therefore, for prolate tensors we set the importance density as

$$\pi(\hat{v}_{i+1}|\hat{v}_i, \mathcal{Y}) = f_3(\hat{v}_{i+1}; \hat{e}_1^i, \nu_i). \quad (11)$$

For oblate tensors, since the observation density in Equation (7) is wide, we can still use the prior as the importance density given in Equation (10).

3.4 Algorithm Outline

Given K particles at step i : $\hat{v}_{0:i}^{(k)}$, $k = 1, \dots, K$, the iteration steps is summarised as

-
- compute diffusion tensor $D_i^{(k)}$ for each particle k
 - **Prediction:** for $k = 1, \dots, K$
 - if $D_i^{(k)}$ is a prolate tensor, sample $\hat{v}_{i+1}^{*(k)}$ using Equation (11)
 - if $D_i^{(k)}$ is an oblate tensor, sample $\hat{v}_{i+1}^{*(k)}$ using Equation (10)
 - **Weighting:** for $k = 1, \dots, K$
 - if prolate, compute $\tilde{w}_{i+1}^{(k)}$ from Equation (5) using Equation (6), (9) and (11)
 - if oblate, compute $\tilde{w}_{i+1}^{(k)}$ from Equation (5) using Equation (7), (9) and (10)
 - **Selection:** normalise all the weights and evaluate N_{eff}
 - If $N_{eff} \geq N_s$, then for $k = 1, \dots, K$, $\hat{v}_{i+1}^{(k)} = \hat{v}_{i+1}^{*(k)}$
 - If $N_{eff} < N_s$, 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 $\hat{v}_{i+1}^{(k)} = \hat{v}_{i+1}^{*(z(k))}$, $\tilde{w}_{i+1}^{(k)} = \frac{1}{N}$
-

4 Experimental Results

We have tested the algorithm on a brain dataset with size $128 \times 128 \times 58$ and $2 \times 2 \times 2mm$ resolution. A six-direction gradient scheme was used with $b = 1000s/mm^2$. Since our particles propagate in a continuous domain, we choose the trilinear method in [12] for interpolation. A step length of 1mm and 5000 particles were used for all examples. The propagation of a particle is stopped when it exits white matter ($FA < 0.2$). We distinguish prolate tensors ($c_l > \tau$) and oblate tensors ($c_l \leq \tau$) by using a threshold $\tau = 0.27$.

Fig. 2(a) shows the trajectories of 1000 particles seeded from a point in the Corpus callosum. The figure shows that the sampled paths provide a robust delineation of the expected fiber bundle. Fig. 2(b) gives another example with two seed points in the Cingulum bundles. This example reveals how our probabilistic algorithm is able to handle splitting fibers and ambiguous neighborhoods. Fig. 2(c) shows the global optimal MAP paths of the examples in Fig. 2(a) and Fig. 2(b). We also compared our result with that of Friman’s method using the same seed points in Cingulum bundles, as shown in Fig. 2(c). 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. In contrast, the sampled paths of the Friman’s method are more dispersed, with a number of paths which have low probabilities. Thus, our method samples more representative paths surrounding the optimal candidate. Moreover, our algorithm runs

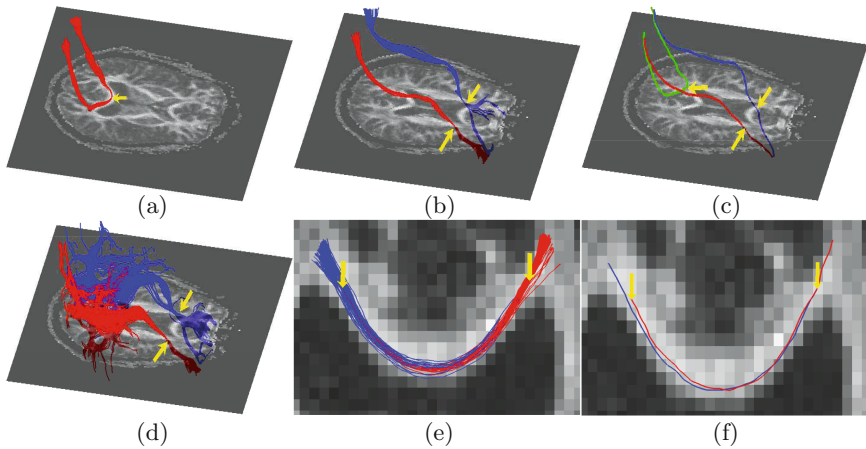


Fig. 2. (a): 1000 particle traces from a seed point in Corpus callosum. (b): from two seed points in the left and right Cingulum bundles. (c): Optimal MAP paths of (a) and (b). (d): 1000 path samples using Friman’s method [3] from the same seed points as in (b). (e): Zoomed particle traces of two seed points from the MAP path of example (a). (f): Optimal MAP paths of (e).

much faster than Friman’s algorithm (more than 30 times faster due to our MATLAB implementation). To further evaluate the algorithm, we set two seed points from the MAP path of the example in Fig. 2(a) and let the algorithm track from them toward each other. Fig. 2(e) shows 1000 sample paths from each seed point. The figure tells that the sampled paths from two seed points are almost overlapped with each other. Fig. 2(f) gives their two optimal MAP paths, which are very close to each other. Thus, the second seed point can successfully go back to the first one along with 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 thus can produce a probability map of connections between the seed and all other voxels. In Fig. 3(a), we show the probability map depicted from




Fig. 3. Probability map of our algorithm from (a): a seed point in the Corpus callosum, and, (b): from two seed points in the Cingulum bundles. (c): Probability map of Friman's method from the same seed points as in (b).

a seed point in Corpus callosum. The coloring shows the belief of our algorithm that a fiber initiated at the seed voxel reaches respective voxel. Fig. 3(b) gives a probability map of longer fiber tracts seeded from Cingulum bundles. The result here is compared to that of Friman's method, as shown in Fig. 3(c), 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. Fiber orientation distribution is formulated in a theoretical way for both prolate and oblate tensors. Fast and efficient sampling is realised using the vMF distribution. As a consequence, there is no need to apply MCMC sampling [1] or to discretise the fiber orientation distribution [3] for sampling paths. Unlike previous methods [1,3] which are computationally expensive, our method is able to rapidly locate the global optimal fiber and compute the connectivity map for a given seed point.

References

1. Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., Clare, S., Matthews, P., Brady, J., Smith, S.: Characterization and Propagation of Uncertainty in Diffusion-Weighted MR Imaging. *Magn. Reson. Med.* 50, 1077–1088 (2003)
2. Bjornemo, M., Brun, A., Kikinis, R., Westin, C.: Regularized Stochastic White Matter Tractography Using Diffusion Tensor MRI. In: Dohi, T., Kikinis, R. (eds.) *MICCAI 2002*. LNCS, vol. 2488, pp. 435–442. Springer, Heidelberg (2002)
3. Friman, O., Farneback, G., Westin, C.: A Bayesian Approach for Stochastic White Matter Tractography. *IEEE Trans. on Med. Imag.* 25(8), 965–978 (2006)
4. Doucet, A., de Freitas, N., Gordon, N. (eds.): *Sequential Monte Carlo Methods in Practice*. Springer, Heidelberg (2001)
5. Anderson, A.: Measurement of Fiber Orientation Distributions Using High Angular Resolution Diffusion Imaging. *Magn. Reson. Med.* 54, 1194–1206 (2005)

6. Wood, A.T.A.: Simulation of the von Mises-Fisher distribution. *Communications in Statistics. Simulation and Computation* 23, 157–164 (1994)
7. Basser, P., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 66, 259–267 (1994)
8. Westin, C., Maier, S., Mamata, H., Nabavi, A., Jolesz, F., Kikinis, R.: Processing and visualization for diffusion tensor MRI. *Medical Image Analysis* 6, 93–108 (2002)
9. Salvador, R., Pena, A., Menon, D., Carpenter, T., Pickard, J., Bullmore, E.: Formal Characterization and Extension of the Linearized Diffusion Tensor Model. *HBM* 24, 144–155 (2005)
10. McGraw, T., Vemuri, B., Yeziarski, R., Mareci, T.: Segmentation of High Angular Resolution Diffusion MRI Modeled as a Field of von Mises-Fisher Mixtures. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *ECCV 2006*. LNCS, vol. 3953, pp. 463–475. Springer, Heidelberg (2006)
11. Hill, G.: Algorithm 571: Statistics for von Mises' and Fisher's Distributions of Directions. *ACM Trans. on Math. Software* 7, 233–238 (1981)
12. Zhukov, L., Barr, A.: Oriented Tensor Reconstruction: Tracing Neural Pathways from Diffusion Tensor MRI. In: *Proc. IEEE Visualization*, pp. 387–394. IEEE Computer Society Press, Los Alamitos (2002)