Adaptive Multi-modal Particle Filtering for Probabilistic White Matter Tractography

Aymeric Stamm[1,⋆], Olivier Commowick[1], Christian Barillot[1], and Patrick Pérez[2]

[1] VISAGES: INSERM U746 - CNRS UMR6074 - INRIA - Univ. of Rennes I, France
[2] Technicolor, Rennes, France
Aymeric.Stamm@irisa.fr

Abstract. Particle filtering has recently been introduced to perform probabilistic tractography in conjunction with DTI and Q-Ball models to estimate the diffusion information. Particle filters are particularly well adapted to the tractography problem as they offer a way to approximate a probability distribution over all paths originated from a specified voxel, given the diffusion information. In practice however, they often fail at consistently capturing the multi-modality of the target distribution. For brain white matter tractography, this means that multiple fiber pathways are unlikely to be tracked over extended volumes.

We propose to remedy this issue by formulating the filtering distribution as an adaptive M-component non-parametric mixture model. Such a formulation preserves all the properties of a classical particle filter while improving multi-modality capture. We apply this multi-modal particle filter to both DTI and Q-Ball models and propose to estimate dynamically the number of modes of the filtering distribution. We show on synthetic and real data how this algorithm outperforms the previous versions proposed in the literature.

1 Introduction

The advent of MRI technology has provided the medical community with a large amount of data that help clinicians in making decisions on a daily basis. Diffusion MRI is the sequence of choice for the study and analysis of the brain white matter (WM) neural network [16]. Yet, the **tractography** problem of inferring the WM neural system from noisy diffusion-weighted images is very challenging. It requires (i) an appropriate **diffusion model** that retrieves the diffusion information (e.g., diffusion orientations, diffusivities, anisotropies) and (ii) a **tracking algorithm** that generates pathways from the diffusion information.

Many diffusion models have been devised in the literature [3]. The very first proposed diffusion model in the literature is diffusion tensor imaging (DTI) [7], which can be viewed as the solution of the modified Bloch-Torrey equation for anisotropic media [24]. Despite its good performance in homogeneous regions of the brain, its robustness to noise and its low computational cost, DTI has

⋆ Corresponding author.

J.C. Gee et al. (Eds.): IPMI 2013, LNCS 7917, pp. 594–606, 2013.

shortcomings: since it summarizes the diffusion in a second-order tensor, it characterizes well uni-oriented anisotropic media but fails to describe multi-oriented ones, which occurs in approximately a third of the voxels in the brain [9]. Solutions to the intra-voxel fiber heterogeneity problem rely on higher-order tensors [19], mixture models [26,8,2,23] or non-parametric models such as the Q-Ball model [25]. The latter is especially appealing since its estimation is linear (and thus fast) and directly provides the orientation information.

The second ingredient for tractography is the tracking algorithm itself [15]. The tractography problem that we tackle here is the generation of possible fiber pathways from one seed voxel without constraining the other extremity of the paths. Existing algorithms that address this problem can be classified as *deterministic* or *probabilistic*. The first category generates a single fiber pathway either by following the estimated local orientations [6,18] or by fast marching front propagation [20,22]. Local deterministic algorithms suffer from a possible accumulation of errors during the tracking process. Fast marching methods partially address this issue by reconstructing the fiber pathway that minimizes these errors. On the other hand, probabilistic algorithms generate several weighted fiber pathways that quantify the relative connectivity of the arrival point to the seed voxel (e.g. [9,12]). Randomness is often introduced on the estimated local diffusion orientations to account for their uncertainty. The von Mises & Fisher distribution on the sphere [14] is particularly well adapted to perform this sampling and has been used within a particle filter in which fiber pathways are reconstructed as first order Markov chains [28,21].

The use of particle filters for WM tractography has been pioneered in [28], where the diffusion information is estimated through DTI. The particle filter is improved in [21] by using the Q-Ball model to account for multiple local fiber orientations. However, particle filters often fail at consistently capturing the multi-modality of the filtering distribution [27]. Multiple fiber pathways are thus unlikely to be tracked over extended volumes. We propose an adaptive multi-modal particle filter for WM tractography that improves the multi-modality capture. We apply it to both DTI and Q-Ball models (Section 2). We then design an experimental framework for validation of the proposed methods (Section 3) and show results on both synthetic data with an extensive comparison to their classic particle filter counterparts and real clinical data (Section 4). Another contribution is that the number of modes of the filtering distribution is dynamically estimated. We show that our proposed algorithm greatly improves WM tractography and, when the associated diffusion model accurately captures multiple fiber orientations, it is able to distinguish crossings from bifurcations.

2 Proposed Algorithm

2.1 Tractography as a Mixture Filtering Problem

Denote $\Omega \subseteq \mathbb{R}^3$ the image spatial domain. An image over this domain is a set $A = \{a_{\mathbf{x}}, \mathbf{x} \in \Omega\}$. Let then $\{S_i\}_{i=1}^n$ be a set of n raw gradient images, S_0 be one non-weighted diffusion image and Ψ one diffusion model image.

At step k, a fiber pathway is a sequence $X^k = [(\mathbf{x}_0, \mathbf{v}_{-1}), \ldots, (\mathbf{x}_k, \mathbf{v}_{k-1})] \in \left(\Omega, \mathbb{S}^2\right)^{k+1}$ of successive pairs of positions and arrival directions, related by $\mathbf{x}_{k+1} = \mathbf{x}_k + \rho \mathbf{v}_k$, where $\rho > 0$ is the step size, which is assumed to be constant. The state space at step k is the set of all possible fiber pathways X^k originated from a specified position $\mathbf{x}_0$ and a specified arrival direction $\mathbf{v}_{-1}$.

The filtering distribution at step k is the distribution of X^k, given the diffusion data $\mathcal{Y}_k(X^k) = \{y_{\mathbf{x}_0}, y_{\mathbf{x}_1}, \ldots, y_{\mathbf{x}_k}\}$ where $y_{\mathbf{x}_j} = \{s_{\mathbf{x}_j,i}, \psi_{\mathbf{x}_j,i}\}_{i=1}^n$, for all $j = 0, \ldots, k$. At each step k, this distribution is sequentially determined by successively computing the prediction distribution $p(X^k|\mathcal{Y}_{k-1}(X^{k-1}))$ out of the previous filtering distribution and then using Bayes' rule to obtain $p(X^k|\mathcal{Y}_k(X^k))$. These **prediction** and **update** stages require to specify respectively the evolution model $p(\mathbf{x}_{k+1}, \mathbf{v}_k|X^k)$ and the likelihood $p(\mathcal{Y}_{k+1}(X^{k+1})|X^{k+1})$.

We assume that fiber pathways are first order Markov chains. Consequently, the evolution model simplifies to $p(\mathbf{v}_k|\mathbf{x}_k, \mathbf{v}_{k-1})$. In the remainder of the article, according to [28,21], we use the following evolution model:

$$p(\mathbf{v}_k|\mathbf{x}_k, \mathbf{v}_{k-1}) = \mathrm{vMF}(\mathbf{v}_k; \mathbf{v}_{k-1}, \kappa) = \frac{\kappa}{4\pi \sinh \kappa} \exp\{\kappa \mathbf{v}_{k-1}^T \mathbf{v}_k\}\,, \tag{1}$$

where $\mathrm{vMF}(\cdot; \mathbf{v}_{k-1}, \kappa)$ is the von Mises & Fisher distribution [14] on the 2-dimensional sphere with mean direction $\mathbf{v}_{k-1} \in \mathbb{S}^2$ and concentration parameter $\kappa \geq 0$. The concentration parameter κ of the evolution model controls the smoothness of the reconstructed fiber pathways. Assuming conditional independence of the observations given a pathway, the observation model reads:

$$p(\mathcal{Y}_{k+1}(X^{k+1})|X^{k+1}) = \prod_{j=0}^{k} p(y_{\mathbf{x}_{j+1}}|\mathbf{v}_j)\,, \tag{2}$$

where $p(y_{\mathbf{x}_{j+1}}|\mathbf{v}_j)$ depends on the diffusion model and will be defined in Section 2.4. In order to better capture multi-modality, we follow the idea of [27] and formulate the filtering distribution as a mixture of M_k components:

$$p(X^k|\mathcal{Y}_k(X^k)) = \sum_{m=1}^{M_k} \pi_{m,k} p_m(X^k|\mathcal{Y}_k(X^k))\,, \tag{3}$$

where $\sum_{m=1}^{M_k} \pi_{m,k} = 1$. Such a formulation allows us to perform the filtering recursion for each component p_m individually, provided that each mixture weight is updated as the normalized weighted likelihood for the associated component.

2.2 Mixture Particle Filter

In general, there is no closed-form expressions for the filtering recursion equations. A popular strategy is to resort to particle filters. They approximate the filtering distribution by a set of samples that are properly weighted to represent the filtering distribution at each step. Using the notations in [27], let $\mathcal{P}_k = \{M_k, \Pi_k, \mathcal{X}_k, \mathcal{W}_k, \mathcal{C}_k\}$ be the particle representation of the filtering distribution where M_k is the number of components, $\Pi_k = \{\pi_{m,k}\}_{m=1}^{M_k}$ the set of

mixture weights, $\mathcal{X}_k = \{\mathbf{x}_k^{(\ell)}\}_{\ell=1}^N$ the set of N particles, $\mathcal{W}_k = \{w_k^{(\ell)}\}_{\ell=1}^N$ the set of particle weights and $\mathcal{C}_k = \{c_k^{(\ell)}\}_{\ell=1}^N$ the set of component indicators (i.e., $c_k^{(\ell)} = m$ if particle ℓ belongs to component m). Given $\mathcal{P}_k$, the particle approximation with mixture filtering distribution proceeds to step $k+1$ in five stages:

Proposition of New Samples: New samples are generated according to a proposal density $q(\cdot|\mathbf{v}_{k-1}^{(\ell)}, y_{\mathbf{x}_k^{(\ell)}})$ which depends on the previous direction and the diffusion information at step k:

$$\mathbf{v}_k^{(\ell)} \sim q(\mathbf{v}_k|\mathbf{v}_{k-1}^{(\ell)}, y_{\mathbf{x}_k^{(\ell)}}) \text{ and } \mathbf{x}_{k+1}^{(\ell)} = \mathbf{x}_k^{(\ell)} + \rho\mathbf{v}_k^{(\ell)}\,. \tag{4}$$

Update of Particle Weights: The weights of the new particles are updated in order to be representative of the filtering distribution according to [11]:

$$\tilde{w}_{k+1}^{(\ell)} = \frac{w_k^{(\ell)} p(y_{\mathbf{x}_{k+1}^{(\ell)}}|\mathbf{v}_k^{(\ell)}) p(\mathbf{v}_k^{(\ell)}|\mathbf{v}_{k-1}^{(\ell)})}{q(\mathbf{v}_k^{(\ell)}|\mathbf{v}_{k-1}^{(\ell)}, y_{\mathbf{x}_k^{(\ell)}})}\,. \tag{5}$$

The normalization of these weights is performed within each component:

$$w_{k+1}^{(\ell)} = \frac{\tilde{w}_{k+1}^{(\ell)}}{\sum_{j\in\mathcal{I}_{m,k}} \tilde{w}_{k+1}^{(j)}}\,, \tag{6}$$

where $\mathcal{I}_{m,k} = \{\ell \in [\![1, N]\!] : c_k^{(\ell)} = m\}$ is the set of indices of the particles that belong to the m-th mixture component at step k.

Update of Mixture Weights: The mixture weights need to be updated properly to ensure that the filter still acts on each component individually:

$$\pi_{m,k+1} = \frac{\pi_{m,k}\tilde{w}_{m,k+1}}{\sum_{i=1}^M \pi_{i,k}\tilde{w}_{i,k+1}} \text{ with } \tilde{w}_{m,k+1} = \sum_{\ell\in\mathcal{I}_{m,k}} \tilde{w}_{k+1}^{(\ell)}\,. \tag{7}$$

Resampling within Each Component: To avoid the degeneracy of the particle weights, occasional resampling is necessary [11]. The resampling stage can be performed within each subset of particles associated to a mixture component independently, according to the component particle weights [27]: we compute the effective number of particles in a mixture component as:

$$\mathrm{ESS}_m = \left(\sum_{\ell\in\mathcal{I}_{m,k}} \left(w_{k+1}^{(\ell)}\right)^2\right)^{-1}\,, \tag{8}$$

and perform resampling according to the weights in Eq.(6), if ESS_m is below a threshold $\alpha|\mathcal{I}_{m,k}|$, where $|\cdot|$ denotes the set size operator.

Reclustering of the Particles within New Components: The number of components M_k in the mixture is not known. At the end of each step, it is dynamically estimated by merging and/or splitting some of the components: M_k, $\mathcal{C}_k$ and $\mathcal{I}_{m,k}$ are updated to M_{k+1}, $\mathcal{C}_{k+1}$ and $\mathcal{I}_{m,k+1}$ accordingly.

In Section 2.3, we describe how the reclustering of the mixture filtering distribution is performed. In Section 2.4, we define the proposal density and the likelihood for the DTI and Q-Ball models following respectively [28] and [21].

2.3 Reclustering of the Filtering Distribution

After the resampling stage, we characterize each mixture component by a vMF distribution with mean direction $\boldsymbol{\mu}_{m,k}$ and concentration $\kappa_{m,k}$. The parameters of the distribution are estimated using the following equations:

$$\mathbf{r}_{m,k} := \frac{\sum_{\ell \in \mathcal{I}_{m,k}} \mathbf{v}_k^{(\ell)}}{|\mathcal{I}_{m,k}|}, \quad \boldsymbol{\mu}_{m,k} = \frac{\mathbf{r}_{m,k}}{\|\mathbf{r}_{m,k}\|}, \quad \kappa_{m,k} = \frac{\|\mathbf{r}_{m,k}\|(3 - \|\mathbf{r}_{m,k}\|^2)}{1 - \|\mathbf{r}_{m,k}\|^2}. \quad (9)$$

These estimators have been proposed in [4] and have been introduced for diffusion MRI in [10] for their unbiasedness and robustness.

We first test the components pairwise for merging. We merge two components if the two following conditions are met:

1. the Euclidean distance between the mean positions $\overline{\mathbf{x}}_{m,k} := \frac{\sum_{\ell \in \mathcal{I}_{m,k}} \mathbf{x}_k^{(\ell)}}{|\mathcal{I}_{m,k}|}$ is below a threshold ξ_1, and
2. the distance between the two vMF distributions computed with Eq.(9) is below a threshold ξ_2.

We compute the distance between two vMF distributions as proposed in [17]:

$$d\left((\kappa_i, \boldsymbol{\mu}_i), (\kappa_j, \boldsymbol{\mu}_j)\right) = \sqrt{\log^2\left(\frac{\kappa_j}{\kappa_i}\right) + \arccos^2\left(\boldsymbol{\mu}_i^T \boldsymbol{\mu}_j\right)}. \quad (10)$$

We then test each component for splitting. We split a component if its concentration parameter $\kappa_{m,k}$ drops below a threshold ξ_3.

Finally, the number of mixture components is updated to M_{k+1}, the component indicators to $\mathcal{C}_{k+1}$ and the set of indices to $\mathcal{I}_{m,k+1}$. In order to maintain a properly weighted sample from the filtering distribution and thus to preserve the convergence properties of the particle filter, we perform the following update of mixture and particle weights [27]:

$$\pi^\star_{m,k+1} = \sum_{\ell \in \mathcal{I}_{m,k+1}} \pi_{c_k^{(\ell)},k+1} w_{k+1}^{(\ell)}, \qquad w_{k+1}^{(\ell)\star} = \frac{\pi_{c_k^{(\ell)},k+1} w_{k+1}^{(\ell)}}{\pi^\star_{c_{k+1}^{(\ell)},k+1}}. \quad (11)$$

2.4 Diffusion Models: Associated Proposal Densities and Likelihoods

The DTI Model. It provides a 2nd order diffusion tensor represented by its eigensystem $\{\lambda_1, \lambda_2, \lambda_3, \mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$, of which we extract the fractional anisotropy FA [5], the linear coefficient c_l defined as in [28], the mean diffusivity

$\overline{\lambda} = (\lambda_1 + \lambda_2 + \lambda_3)/3$, the perpendicular diffusivity $\lambda_\perp = (\lambda_2 + \lambda_3)/2$, the principal eigenvector $\mathbf{e}_1$ and the minor eigenvector $\mathbf{e}_3$. The likelihood is given by:

$$p(y_{\mathbf{x}_{k+1}}|\mathbf{v}_k) = \begin{cases} \left[\prod_{i=1}^{n} \frac{s^\star_{\mathbf{x}_{k+1},i}}{\sigma_i\sqrt{2\pi}} \exp\left\{-\frac{(s^\star_{\mathbf{x}_{k+1},i})^2(\log s_{\mathbf{x}_{k+1},i} - \log s^\star_{\mathbf{x}_{k+1},i})^2}{2\sigma_i^2}\right\}\right]^{\frac{1}{n}}, & c_l > \tau\,, \\ \frac{1}{\sigma\sqrt{(2\pi)^3}} \exp\left\{-\frac{(\arccos(\mathbf{v}_k^T\mathbf{e}_3) - \pi/2)^2}{2\sigma^2}\right\}, & c_l \leq \tau\,, \end{cases} \quad (12)$$

where the diffusion tensor is estimated at position $\mathbf{x}_{k+1}$, σ_i and $\mathbf{g}_i$ are the standard deviation and the gradient direction of the i-th gradient image respectively, estimated by least square estimation and pseudo-residuals [13], σ is a user-defined standard deviation and $s^\star_{\mathbf{x}_{k+1},i} = s^\star_{\mathbf{x}_{k+1},0} \exp\{-b(\lambda_\perp + 3(\mathbf{v}_k^T\mathbf{g}_i)^2(\overline{\lambda} - \lambda_\perp))\}$ is the diffusion signal simulated from the diffusion tensor cylindrically constrained along the sampled direction $\mathbf{v}_k$. The proposal density is given by:

$$q(\mathbf{v}_k|\mathbf{v}_{k-1}, y_{\mathbf{x}_k}) = \begin{cases} \text{vMF}(\mathbf{v}_k; \mathbf{e}_{1,k}, \nu_k)\,, & c_l > \tau\,, \\ p(\mathbf{v}_k|\mathbf{v}_{k-1})\,, & c_l \leq \tau\,, \end{cases} \quad (13)$$

where the diffusion tensor is estimated at $\mathbf{x}_k$ and ν_k is a function of FA [28].

The Q-Ball model. It provides an orientation distribution function (ODF) of which we extract the set Λ of maxima $\boldsymbol{\mu}$, the value of the ODF at its maxima $\psi(\boldsymbol{\mu})$ and the mean curvature of the ODF at its maxima $H(\boldsymbol{\mu})$. Borrowing ideas from [21], we define the likelihood as follows:

$$p(y_{\mathbf{x}_{k+1}}|\mathbf{v}_k) = \left[\prod_{i=1}^{n} \frac{1}{\sigma_i\sqrt{2\pi}} \exp\left\{-\frac{(s_{\mathbf{x}_{k+1},i} - s^\star_{\mathbf{x}_{k+1},i})^2}{2\sigma_i^2}\right\}\right]^{\frac{1}{n}}, \quad (14)$$

where the ODF is estimated at position $\mathbf{x}_{k+1}$ and $s^\star_{\mathbf{x}_{k+1},i}$ is the diffusion signal simulated according to [1] from the ODF that has been rotated to align the sampling direction to the sampled one. The proposal density is given by:

$$q(\mathbf{v}_k|\mathbf{v}_{k-1}, y_{\mathbf{x}_k}) = \begin{cases} \sum_{\boldsymbol{\mu}\in\Lambda} \omega_{\boldsymbol{\mu}} \text{vMF}(\mathbf{v}_k; \boldsymbol{\mu}, \kappa_{\boldsymbol{\mu}})\,, & \Lambda \neq \emptyset\,, \\ p(\mathbf{v}_k|\mathbf{v}_{k-1})\,, & \Lambda = \emptyset\,, \end{cases} \quad (15)$$

where the ODF is estimated at $\mathbf{x}_k$, $\omega_{\boldsymbol{\mu}} \propto \psi(\boldsymbol{\mu})$ (normalized) and $\kappa_{\boldsymbol{\mu}} \propto H(\boldsymbol{\mu})$.

3 Experimental Setup and Evaluation Metrics

3.1 Phantom Diffusion Weighted Data

Two synthetic diffusion weighted phantoms were created for validation and are illustrated in Fig. 1: a case of two crossing fibers at a 90° angle and a case of

one fiber splitting into two fibers at a 60° angle. For both phantoms, one non-weighted diffusion image and 81 raw gradient images with a single b-value of 3000 s.mm^{-2} were simulated using an equally weighted multi-tensor model at each voxel. Rician noise was then added on the noise-free images with a relative standard deviation of 5%, to generate 50 samples of each phantom.

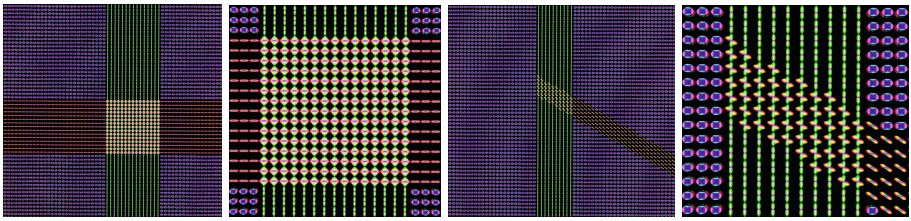

Fig. 1. Noise Free Diffusion Weighted Phantoms. ODF visualization of crossing fibers (a) with close-up view (b), and the bifurcating fibers (c) with close-up view (d).

3.2 Evaluation Measures for Phantom Data

Four methods were utilized: DTI-based "mono-modal" (i.e., without particle clustering mechanism) tractography (DTI mono), DTI-based multi-modal tractography (DTI multi), ODF-based mono-modal tractography (ODF mono) and ODF-based multi-modal tractography (ODF multi). For each phantom, a single seed voxel was placed in the upper branch. The initial direction of propagation was set towards the bottom. The following common parameters were used in all algorithms: resampling threshold $\alpha = 0.4$, number of particles $N = 1000$, step length $\rho = 1$ mm, prior concentration $\kappa = 30$, and merge and split thresholds $\xi_1 = 1$ mm, $\xi_2 = 1$ and $\xi_3 = 40$. Tensors were considered as oblate for $\tau = 0.25$. The output fiber pathways are the averaged of each cluster of particles (a single one in mono-modal versions).

The tractography results were evaluated visually and with 3 different quantitative measures: (i) the proportion of fibers branching from the main direction, which is an indicator of branching capacity (ii) the root-mean-square error between the end point (after $L_{\max}$ iterations of the particle filter) of each fiber following the main path and the expected arrival position (known in the phantoms), which gives an idea of how spread the fibers are around the true one and (iii) the local curvature along each branching fiber for the bifurcation phantom, which translates how each branch was created from the main direction (either by an uncertain turn or by a sharper local turn).

The expected arrival position in the bifurcation phantom is not obvious. In this phantom, fibers going straight follow the inaccurate diffusion orientations given by the diffusion model in the heterogeneous region and are thus expected to deviate exclusively towards the right border of the vertical band. Therefore, we have chosen the end position for the bifurcation phantom at the center of the segment joining the center of the vertical band and its right border.

3.3 Clinical Diffusion Acquisition

The 4 algorithms were also applied on real clinical scans, acquired on a Siemens 3T scanner with a matrix size of 128x128, 60 slices (voxel size 2x2x2 mm^3). The diffusion acquisition consisted of one non-weighted diffusion image and 30 gradient images with a b-value of 1000 $s.mm^{-2}$. Seed regions were placed by a radiologist at the basis of the left and right cortico-spinal tracts (CST) in the mesencephalon, with filtering regions in the posterior limb of the internal capsule to keep only the CST. The same parameters as for synthetic data were utilized for real data. A particle filter was initiated at each voxel of the seed regions and the fiber pathways are the averaged fibers of each cluster of particles.

4 Results

4.1 Experiments on Synthetic Data

We present a representative example of the results achieved by each method in Fig. 2. We clearly notice that the 2 mono-modal methods fail to capture the multi-modality of the bifurcation phantom and therefore follow only one of the two directions. On the contrary, the 2 multi-modal methods are visually well able to capture the two branching fiber tracts, thanks to the adaptive clustering based on the proposed directions. It may be noted that DTI multi tends to obtain more fanning fibers, because the observation model is wide for oblate tensors. These visual results are valid for both crossing and bifurcation phantoms. However, in the crossing, fibers are only expected to go straight since the crossing ones are

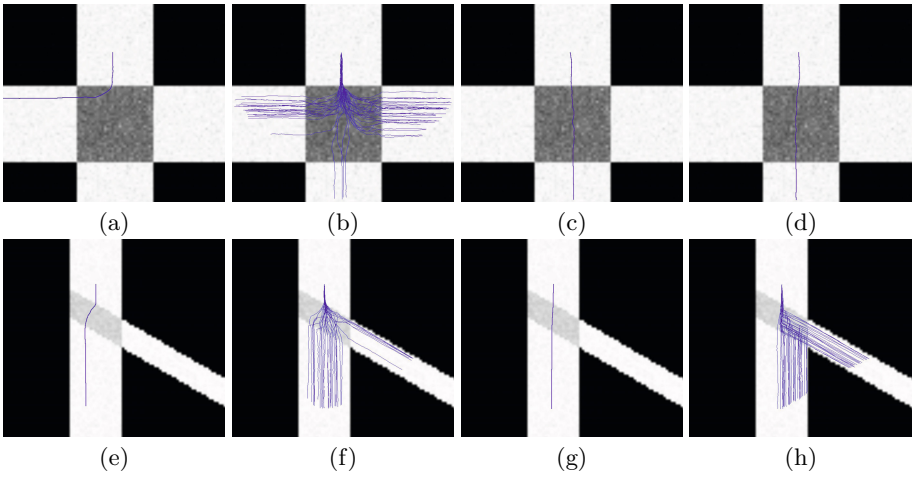

Fig. 2. Tractography Results on Phantom Data. First row: crossing phantom, Second row: bifurcation phantom. Methods used were DTI mono (a,e), DTI multi (b,f), ODF mono (c,g), ODF multi (d,h).

not part of the same pathway. Therefore, ODF mono and multi are performing well while DTI multi tends to capture too many branches.

We report in Table 1 the proportion of branching fibers for each phantom and each method. These quantitative results confirm the visual ones. When utilizing the mono-modal methods, only one of the two branches of each phantom is explored. On the contrary, multi-modal methods capture much better the 2 modes in the bifurcation phantom, with ODF multi being the closest to the half/half ground truth in each branch.

Table 1. Evaluation of Fibers on Phantom Data. Proportion of branching and straight fibers and Root Mean Squared error of fibers going straight with respect to true expected position, for each phantom.

	DTI mono	DTI multi	ODF mono	ODF multi
Crossing Phantom				
Fiber proportion straight (%)	0 ± 0	8.75 ± 3.06	100 ± 0	98.76 ± 6.14
Fiber proportion branch (%)	100 ± 0	91.25 ± 3.06	0 ± 0	1.24 ± 6.14
RMS (mm)	N/A	9.71 ± 3.71	2.09 ± 1.35	2.05 ± 1.15
Bifurcation Phantom				
Fiber proportion straight (%)	93.88 ± 24.22	66.53 ± 4.06	100 ± 0	56.63 ± 4.59
Fiber proportion branch (%)	6.12 ± 24.22	33.47 ± 4.06	0 ± 0	43.37 ± 4.59
RMS (mm)	20.16 ± 6.07	15.36 ± 1.28	12.94 ± 0.29	9.08 ± 0.41

In addition, Table 1 displays the RMS error towards the expected arrival point of the straight fibers only for those fibers which go in the straight branch of each phantom. For both phantoms, DTI mono and multi perform worse as DTI does not handle multiple directions. ODF multi outperforms the other methods, being able to better recover the final positions of the fibers.

The last metric, only for the bifurcation phantom, is the local curvature of each mean fiber that deviates from the main vertical path. We report one representative example of the obtained curves as well as a box-plot of inter-quantile

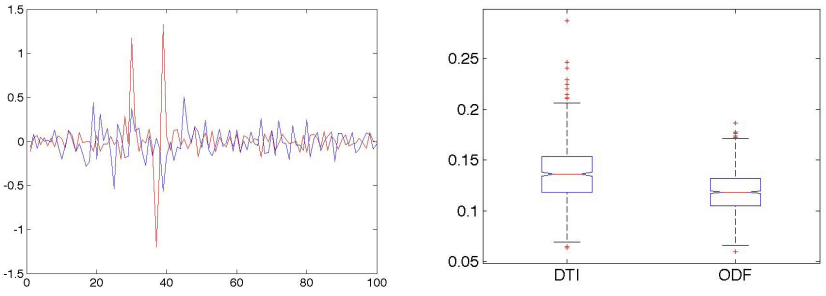

(a) (b)

Fig. 3. Local Curvature of Branching Fibers. Evaluation of local curvature for the branching fibers of the splitting fibers phantom. (a): Local curvature on one single fiber (blue: DTI-based, multimodal, red: ODF-based, multimodal), (b): box-plot representation of inter-quantile range for DTI and ODF over all fibers and repetitions.

ranges of the curvatures along each mean fiber for all repetitions (Fig. 3). The curves of local curvatures clearly show a more peaked behavior for ODF multi (red curve), indicating that it branches more sharply. This was expected as the ODF model captures the 2 fiber orientations in the splitting region, whereas the DTI model inaccurately estimates the fiber orientations and thus the particle filter is mainly driven by the previous direction with a wide observation model. Also, a one-way ANOVA quantitatively shows a significant difference ($p \ll 10^{-3}$) between the standard deviations of the curvature of the 2 methods: away from the peaks, the curvature varies less with ODF multi than with DTI multi.

4.2 Experiments on Real Data

To illustrate the capacities of the proposed algorithms on real clinical datasets, we report in Fig. 4 the left and right cortico-spinal tracts (CST) obtained on

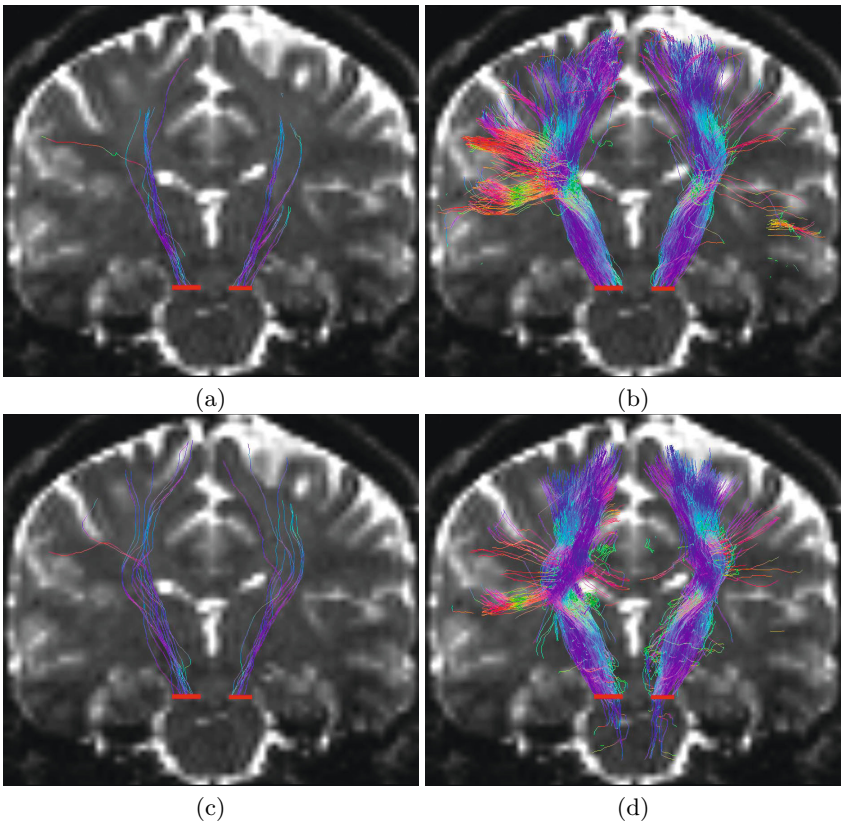

Fig. 4. Tractography of the Left and Right Corticospinal Tracts. Coronal view of both corticospinal tracts using the four proposed methods: (a): DTI mono, (b): DTI multi, (c): ODF mono, (d): ODF multi. Fiber tracts are overlaid on the T2-weighted volume from the diffusion acquisition, red bars indicate the seeding regions.

a normal control subject. The results obtained here are consistent with those obtained on synthetic data. While the mono-modal methods do not capture branchings to lateral parts of the CST, both DTI multi and ODF multi are able to capture branches to the hand area or even sometimes the face area. Interestingly, DTI multi seems a bit more able to capture branches than ODF multi especially on the right side of the brain (left in the images). Since the data was acquired with only 30 gradient directions and a single low b-value of 1000 s.mm^{-2}, it might indeed not be enough for the ODF model to identify accurately multiple orientations. However, ODF mono and multi seem overall able to capture more accurate directions of the main tract.

5 Discussion

We have presented a new adaptive multi-modal particle filter algorithm for diffusion MRI probabilistic tractography. It relies on the adaptive clustering of the filtering distribution through a new scheme for splitting and merging clusters based on the distribution of directions and positions among clusters. This strategy is applied after each step of the filtering recursion. We have implemented this algorithm with two different diffusion models: DTI and ODF on a spherical harmonics basis.

We have demonstrated through experiments on synthetic and real data that our proposed algorithms outperform more classical particle filtering approaches available in the literature, being more able to capture branching and crossing fibers. In addition, the ODF multi-modal algorithm produces more accurate branchings and differentiate crossing fibers from splitting fibers, thanks to the diffusion model that already accurately captures multiple orientations.

As noted in Section 4.2, it will be interesting in the near future to apply this algorithm using other diffusion models such as multi-compartment models [26,8,2,23], more adapted to data with a small number of gradient directions and a single low b-value. The proposed algorithm is very generic and can thus easily be extended to new diffusion models: only the proposal and the observation densities require to be modified to reflect the underlying diffusion model. Future work will also include evaluation on more real data cases, also with HARDI acquisitions, as well as an in-depth study of the usual parameters of a particle filter and their influence on the results.

References

1. Aganj, I., Lenglet, C., Sapiro, G., Yacoub, E., Ugurbil, K., Harel, N.: Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. MRM 64(2), 554–566 (2010)
2. Assaf, Y., Basser, P.: Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. NeuroImage 27(1), 48–58 (2005)
3. Assemlal, H., Tschumperlé, D., Brun, L., Siddiqi, K.: Recent advances in diffusion MRI modeling: Angular and radial reconstruction. MedIA (2011)

4. Banerjee, A., Dhillon, I., Ghosh, J., Sra, S.: Clustering on the unit hypersphere using von Mises-Fisher distributions. J. of Machine Learning 6, 1345–1382 (2006)
5. Basser, P., Pierpaoli, C.: Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. J. Magn. Reson. 111(3), 209–219 (1996)
6. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A.: In vivo fiber tractography using DT-MRI data. MRM 44(4), 625–632 (2000)
7. Basser, P., Mattiello, J., Le Bihan, D.: MR diffusion tensor spectroscopy and imaging. Biophysical journal 66(1), 259–267 (1994)
8. 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. MRM 50(5), 1077–1088 (2003)
9. Behrens, T.E.J., Berg, H.J., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34(1), 144–155 (2007)
10. Bhalerao, A., Westin, C.-F.: Hyperspherical von Mises-Fisher Mixture (HvMF) Modelling of High Angular Resolution Diffusion MRI. In: Ayache, N., Ourselin, S., Maeder, A. (eds.) MICCAI 2007, Part I. LNCS, vol. 4791, pp. 236–243. Springer, Heidelberg (2007)
11. Doucet, A., Godsill, S., Andrieu, C.: On sequential Monte Carlo sampling methods for Bayesian filtering. Statistics and Computing 10(3), 197–208 (2000)
12. Friman, O., Farnebäck, G., Westin, C.: A Bayesian approach for stochastic white matter tractography. IEEE TMI 25(8), 965–978 (2006)
13. Gasser, T., Sroka, L., Jennen-Steinmetz, C.: Residual variance and residual pattern in nonlinear regression. Biometrika 73(3), 625–633 (1986)
14. Jupp, P., Mardia, K.: A unified view of the theory of directional statistics. International Statistical Review 57(3), 261–294 (1989)
15. Lazar, M.: Mapping brain anatomical connectivity using white matter tractography. NMR in Biomedicine 23(7), 821–835 (2010)
16. Le Bihan, D.: Looking into the functional architecture of the brain with diffusion MRI. Nature reviews. Neuroscience 4(6), 469–480 (2003)
17. McGraw, T., Vemuri, B.: Von Mises-Fisher mixture model of the diffusion ODF. In: IEEE ISBI, pp. 65–68 (2006)
18. Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.: Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Annals of Neurology 45(2), 265–269 (1999)
19. Ozarslan, E., Mareci, T.: Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. MRM 50(5), 955–965 (2003)
20. Parker, G., Wheeler-Kingshott, C., Barker, G.: Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. IEEE TMI 21(5), 505–512 (2002)
21. Pontabry, J., Rousseau, F.: Probabilistic tractography using Q-ball modeling and particle filtering. In: Fichtinger, G., Martel, A., Peters, T. (eds.) MICCAI 2011, Part II. LNCS, vol. 6892, pp. 209–216. Springer, Heidelberg (2011)
22. Staempfli, P., Jaermann, T., Crelier, G.R., Kollias, S., Valavanis, A., Boesiger, P.: Resolving fiber crossing using advanced fast marching tractography based on diffusion tensor imaging. NeuroImage 30(1), 110–120 (2006)
23. Stamm, A., Pérez, P., Barillot, C.: A new multi-fiber model for low angular resolution diffusion MRI. In: IEEE ISBI, pp. 936–939 (2012)

24. Stejskal, E.O.: Use of spin echoes in a pulsed magneticfield gradient to study anisotropic, restricted diffusion and flow. J. Chem. Phys. 43, 3597 (1965)
25. Tuch, D.S.: Q-ball imaging. MRM 52(6), 1358–1372 (2004)
26. Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., Wedeen, V.: High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. MRM 48(4), 577–582 (2002)
27. Vermaak, J., Doucet, A., Pérez, P.: Maintaining multimodality through mixture tracking. In: IEEE ICCV. vol. 2, pp. 1110–6 (2003)
28. Zhang, F., Hancock, E., Goodlett, C., Gerig, G.: Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. MedIA 13(1), 5–18 (2009)