A NEIGHBORHOOD-BASED PROBABILISTIC APPROACH FOR FIBER TRACKING IN HUMAN CARDIAC DTI

Hongying Li, Marc Robini, Feng Yang, and Yuemin Zhu

CREATIS (CNRS UMR 5220, Inserm U1044), INSA Lyon, 69621 Villeurbanne cedex, France

ABSTRACT

Fiber tracking techniques are essential for representing and visualizing the cardiac fiber architecture information encoded in diffusion-tensor imaging (DTI) data. We propose a neighborhood-based probabilistic fiber tracking method for cardiac DTI which accounts for spatial correlation and data uncertainty. The method consists in tracking fiber paths by sampling step directions from a normalized weighted sum of local fiber orientation probability distributions. The sum is over a predefined neighborhood of the current position in the fiber path, and the weights depend on fractional anisotropy, angular change and distance to the current position. Experimental results on synthetic data and on real human cardiac data show that, compared to the streamlining approach, the proposed method is meaningfully more robust to noise and produces smoother and more consistent fibers.

***Index Terms*—** Magnetic resonance imaging (MRI), diffusion-tensor imaging (DTI), fiber tracking, probabilistic tracking.

1. INTRODUCTION

Fiber tracking techniques are essential for representing and visualizing the fiber architecture information encoded in diffusion-tensor imaging (DTI) data [1]. Standard tracking methods such as streamlining generate fiber paths from tensor fields or from principal eigenvector fields [2] without taking data uncertainty into account. Consequently, standard tracking methods usually require some preprocessing to remove noise and artifacts from the raw diffusion-weighted (DW) data (or from the tensor field extracted from the DW data, or from the principal eigenvector field computed from the tensors). Fiber tracking can be improved by resorting to global optimization or probabilistic modeling. Global methods have a better ability to extract fiber bundles from human brain data [3] or from low SNR human cardiac data [4], but they are computationally very expensive. Probabilistic tracking methods explicitly account for data uncertainty by describing the local fiber orientation via a posterior probability distribution that incorporates data likelihood and prior information [5, 6] (fiber paths are generated by following the directions randomly sampled from this posterior), but they do not exploit the spatial correlation of the fiber orientations that is observed in the cardiac muscle tissue architecture.

In this paper, we propose an efficient probabilistic method that takes spatial correlation into account. The method consists in tracking fiber paths by sampling step directions from a normalized weighted sum of local fiber orientation probability distributions. The sum is over a predefined neighborhood of the current position in the fiber path, and the weights depend on fractional anisotropy, angular change and distance to the current position. Our method is described in the next section. We provide experimental results on synthetic data and on real human cardiac data in Section 3.

2. METHOD

The raw DTI data produced by MRI machines consists in N DW volumes associated with each gradient direction $g^{(i)}$, $i \in [\![1, N]\!]$ ($[\![1, N]\!]$ is a shorthand notation for $\{1, \dots N\}$). Each raw DW volume $y^{(i)}$ is defined on a 3-D voxel lattice V, and we let $y_v^{(i)}$ denote the intensity of $y^{(i)}$ at voxel $v \in V$. Assuming that the SNR is sufficiently high, the observed voxel intensities can be modeled by taking the logarithm of the true voxel intensities and adding the realizations of independent Gaussian random variables [7], that is, for any $i \in [\![1, N]\!]$ and any $v \in V$,

$$\ln y_v^{(i)} \;=\; \ln \mu_v^{(i)} + \varepsilon_v^{(i)} \tag{1}$$

with $\varepsilon_v^{(i)} \sim \mathcal{N}\big(0, \big(\sigma_v/\mu_v^{(i)}\big)^2\big)$. In this case, the conditional distribution of the log-transformed data $\mathbf{z}_v = \big(\ln(y_v^{(1)}), \dots, \ln(y_v^{(N)})\big)$ given $\boldsymbol{\mu}_v = (\mu_v^{(1)}, \dots, \mu_v^{(N)})$ and the residual noise standard deviation σ_v writes

$$\begin{aligned} &\mathrm{P}(\mathbf{z}_v |\, \boldsymbol{\mu}_v, \sigma_v) \\ &\qquad = \prod_{i=1}^{N} \frac{\mu_v^{(i)}}{\sigma_v \sqrt{2\pi}} \exp\!\left(- \frac{(\mu_v^{(i)})^2}{2\sigma_v^2} \left(\ln \frac{y_v^{(i)}}{\mu_v^{(i)}} \right)^2 \right). \end{aligned} \tag{2}$$

To express the data likelihood in terms of the principal diffusion direction field — denote it by $\{\mathbf{v}_v \in \mathbb{U};\, v \in V\}$ with $\mathbb{U}$ being the unit sphere — we model the true DW data $\boldsymbol{\mu}$ using the constrained model proposed in [8]. This model is obtained by assuming that the two smallest eigenvalues of each diffusion tensor $\mathbf{D}_v$ are equal; it approximates $\boldsymbol{\mu}_v$ by a function of

$\mathbf{v}_v$ that is parameterized by a set $\boldsymbol{\theta}_v$ of parameters computed from the eigen-decomposition of $\mathbf{D}_v$. The modified likelihood is given by

$$\mathrm{P}(\mathbf{z}_v|\,\mathbf{v}_v = \mathbf{u}, \boldsymbol{\theta}_v, \sigma_v) \;=\; \mathrm{P}(\mathbf{z}_v|\,\boldsymbol{\mu}_v = \widehat{\boldsymbol{\mu}}(\mathbf{u}; \boldsymbol{\theta}_v), \sigma_v), \quad (3)$$

where $\widehat{\boldsymbol{\mu}}(\,\cdot\,; \boldsymbol{\theta}_v)$ is the constrained-model approximation of the true voxel intensities: letting $\lambda_{1,v}$, $\lambda_{2,v}$ and $\lambda_{3,v}$ be the eigenvalues of $\mathbf{D}_v$ with $\lambda_{1,v} \geqslant \lambda_{2,v} \geqslant \lambda_{3,v}$, we have

$$\widehat{\mu}^{(i)}(\mathbf{u}; \boldsymbol{\theta}_v) \;=\; \mu_v^{(0)} \exp\bigl(-\, b^{(i)}\bigl(\theta_{v,1} + \theta_{v,2}\langle g^{(i)}, \mathbf{u}\rangle^2\bigr)\bigr) \quad (4)$$

$$\text{with} \quad \theta_{v,1} = \frac{\lambda_{2,v} + \lambda_{3,v}}{2} \quad \text{and} \quad \theta_{v,2} = \lambda_{1,v} - \theta_{v,1}, \quad (5)$$

where $\langle\cdot,\cdot\rangle$ denotes the standard Euclidean dot product, $\mu^{(0)}$ is the true non-diffusion-weighted volume (in practice, $\mu^{(0)}$ is replaced by its experimental measurement $y^{(0)}$), and the gradient directions $g^{(i)}$ and the b-values $b^{(i)}$ are known experimental parameters.

If we do not assume any prior knowledge on the diffusion direction or on the raw data, the function $\mathbf{u} \in \mathbb{U} \mapsto \mathrm{P}(\mathbf{z}_v|\,\mathbf{v}_v = \mathbf{u}, \boldsymbol{\theta}_v, \sigma_v)$ is proportional to the diffusion direction probability map at voxel v. The likelihood (3) is "voxel-wise": it only uses the information attached to the voxel under consideration, and thus it does not take into account the spatial correlation inherent to DTI data eligible for fiber tracking. To overcome this weakness, we propose to define the likelihood at a given voxel v as a weighted sum of the voxel-wise likelihoods within a spatial neighborhood of v. We give more weights to the voxels where the FA is high and to the voxels that define a small angular change in the direction of the fiber being tracked. More precisely, let $\mathcal{G}$ be a neighborhood system on V, that is, a collection $\mathcal{G} = \{\mathcal{G}(v)\,;\, v \in V\}$ of subsets of V such that (i) $v \notin \mathcal{G}(v)$ for all $v \in V$, and (ii) $w \in \mathcal{G}(v) \iff v \in \mathcal{G}(w)$ for all $(v, w) \in V^2$. Let $w \in \mathcal{G}(v)$ be a neighbor of the current voxel v, let $\mathbf{c}_w$ be the center of w, and let $\overrightarrow{\mathbf{s}_{k-1}\mathbf{s}_k}$, $\mathbf{s}_k \in v$, be the last computed fiber segment (see Fig. 1). Then, the weight assigned to w is given by

$$\xi_k(w) \;=\; (\mathrm{FA}(w))^{\alpha}\,|\cos(\varphi_k(\mathbf{c}_w))|^{\beta} \quad (6)$$

$$\text{with} \quad \cos(\varphi_k(\mathbf{a})) \;=\; \frac{\bigl\langle \overrightarrow{\mathbf{s}_{k-1}\mathbf{s}_k}, \overrightarrow{\mathbf{s}_k\mathbf{a}} \bigr\rangle}{\bigl\|\overrightarrow{\mathbf{s}_{k-1}\mathbf{s}_k}\bigr\|\,\bigl\|\overrightarrow{\mathbf{s}_k\mathbf{a}}\bigr\|}, \quad (7)$$

where α and β are positive parameters that control the influence of FA and angular change in the weighting process — note that the angular weight is symmetric with respect to $\pi/2$, so that forward and backward tracking of a given fiber are equally likely. We also weight the influence of the neighbors by a decreasing function of their distance to the current fiber position $\mathbf{s}_k$. Denoting this function by ρ, the weighted likelihood is defined by

$$\begin{aligned}&\mathrm{P}(\mathbf{z}_v|\,\mathbf{v}_v = \mathbf{u}, \mathbf{s}_{k-1}, \mathbf{s}_k, \boldsymbol{\theta}_v, \sigma_v)\\ &\quad\propto \sum_{w \in \mathcal{G}(v)\cup\{v\}} \rho\bigl(\bigl\|\overrightarrow{\mathbf{s}_k\mathbf{c}_w}\bigr\|\bigr)\xi_k(w)\mathrm{P}(\mathbf{z}_w|\,\mathbf{v}_w = \mathbf{u}, \boldsymbol{\theta}_w, \sigma_w),\end{aligned} \quad (8)$$

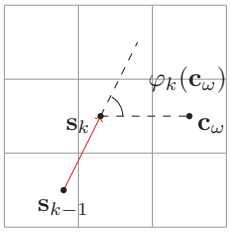


Fig. 1. Angular change used for weighting a neighbor w of the center voxel v: $\varphi_k(\mathbf{c}_w)$ is the angle between the last computed fiber segment going from $\mathbf{s}_{k-1}$ to $\mathbf{s}_k \in v$ and the vector going from the current position $\mathbf{s}_k$ to the center $\mathbf{c}_w$ of w

where it is understood that $\mathbf{s}_k \in v$.

Following [6], the prior knowledge about fiber regularity is modeled by a conditional probability $\mathrm{P}(\mathbf{s}_{k+1}|\mathbf{s}_{k-1}, \mathbf{s}_k)$ that encourages the new step direction $\overrightarrow{\mathbf{s}_k\mathbf{s}_{k+1}}$ to be close to the previous step direction $\overrightarrow{\mathbf{s}_{k-1}\mathbf{s}_k}$ and which decreases to zero as the angle between $\overrightarrow{\mathbf{s}_{k-1}\mathbf{s}_k}$ and $\overrightarrow{\mathbf{s}_k\mathbf{s}_{k+1}}$ goes to $\frac{\pi}{2}$:

$$\begin{aligned}&\mathrm{P}(\mathbf{s}_{k+1}|\,\mathbf{s}_{k-1}, \mathbf{s}_k)\\ &\qquad\propto \begin{cases}(\cos\varphi_k(\mathbf{s}_{k+1}))^{\gamma} & \text{if } \cos\varphi_k(\mathbf{s}_{k+1}) > 0,\\ 0 & \text{otherwise,}\end{cases}\end{aligned} \quad (9)$$

where $\cos\varphi_k(\cdot)$ is given by (7) and the parameter $\gamma \geqslant 0$ controls the fiber regularity degree (the larger γ, the more regular the reconstructed fibers). Note that the probability is zero for angular changes greater than $\frac{\pi}{2}$ in order to prevent fibers from going backwards.

Finally, the combination of the weighted likelihood (8) together with the prior (9) gives the posterior distribution used for sampling the local fiber directions:

$$\begin{aligned}&\mathrm{P}\bigl(\mathbf{s}_{k+1} = \mathbf{s}_k + \delta\mathbf{u}\,\bigl|\,\mathbf{z}_{v_k}, \mathbf{s}_{k-1}, \mathbf{s}_k, \boldsymbol{\theta}_{v_k}, \sigma_{v_k}\bigr)\\ &\qquad\propto \mathrm{P}\bigl(\mathbf{z}_{v_k}\bigl|\,\mathbf{v}_{v_k} = \mathbf{u}, \mathbf{s}_{k-1}, \mathbf{s}_k, \boldsymbol{\theta}_{v_k}, \sigma_{v_k}\bigr)\\ &\qquad\quad\times \mathrm{P}\bigl(\mathbf{s}_{k+1} = \mathbf{s}_k + \delta\mathbf{u}\,\bigl|\,\mathbf{s}_{k-1}, \mathbf{s}_k\bigr),\end{aligned} \quad (10)$$

where v_k denotes the voxel in which $\mathbf{s}_k$ lies and $\delta > 0$ is the step length of the tracking process. A fiber path is generated by iteratively sampling step directions from (10), starting from a given seed point $\mathbf{s}_0$ and until the curvature exceeds a user-defined threshold τ (the curvature κ at a point $\mathbf{s}$ of a fiber path $\mathfrak{F}$ is the reciprocal of the radius of the osculating circle of $\mathfrak{F}$ at $\mathbf{s}$).

On the practical side, the unit sphere $\mathbb{U}$ is restricted to a finite set $\mathbb{U}_0$ of possible step directions in order to facilitate the sampling process. The first step direction is sampled from

$$\begin{aligned}&\mathrm{P}\bigl(\mathbf{s}_1 = \mathbf{s}_0 + \delta\mathbf{u}\,\bigl|\,\mathbf{z}_{v_0}, \mathbf{s}_0, \boldsymbol{\theta}_{v_0}, \sigma_{v_0}\bigr)\\ &\qquad\propto \sum_{w \in \mathcal{G}(v_0)\cup\{v_0\}} \rho\bigl(\bigl\|\overrightarrow{\mathbf{s}_0\mathbf{c}_w}\bigr\|\bigr)(\mathrm{FA}(w))^{\alpha}\\ &\qquad\quad\times \mathrm{P}(\mathbf{z}_w|\,\mathbf{v}_w = \mathbf{u}, \boldsymbol{\theta}_w, \sigma_w),\end{aligned} \quad (11)$$

and the residual noise variance σ_v^2 is approximated by

$$\widehat{\sigma}_v^2 = \frac{1}{N-5}\sum_{i=1}^{N}(y_v^{(i)} - \widehat{\mu}^{(i)}(\mathbf{v}_v^\star;\boldsymbol{\theta}_v))^2, \qquad (12)$$

where $\mathbf{v}_v^\star$ is the unit principal eigenvector of $\mathbf{D}_v$ (we divide by $N-5$ because the constrained-model approximation has 5 parameters, namely, $\theta_{1,v}$, $\theta_{2,v}$ and the vector components of the principal diffusion direction). The procedure for generating a fiber path $\mathbf{s}_0, \mathbf{s}_1, \ldots, \mathbf{s}_K$ given the seed $\mathbf{s}_0$ is the following (given a PDF f on $\mathbb{U}$, we call the restriction of f to $\mathbb{U}_0$ the PDF $\mathbf{u} \in \mathbb{U}_0 \mapsto f(\mathbf{u})/\sum_{\mathbf{v}\in\mathbb{U}_0} f(\mathbf{v})$).

$\diamond$ set $k \leftarrow 0$ and $\kappa \leftarrow 0$;
while $\kappa \leqslant \tau$ **do**
  $\diamond$ add $\mathbf{s}_k$ to the fiber path;
  for each $v \in \mathcal{G}(v_k) \cup \{v_k\}$ **do**
    $\diamond$ estimate the tensor $\mathbf{D}_v$ from the raw DTI data;
    $\diamond$ compute the eigen-decomposition of $\mathbf{D}_v$;
    $\diamond$ compute $\boldsymbol{\theta}_v = \{\theta_{1,v}, \theta_{2,v}\}$ (see (5));
    $\diamond$ estimate the residual noise variance σ_v^2 (see (12));
    $\diamond$ compute the restriction of the modified likelihood $\mathrm{P}(\mathbf{z}_v|\,\mathbf{v}_v = \mathbf{u}, \boldsymbol{\theta}_v, \sigma_v)$ (see (3)) to $\mathbb{U}_0$;
  end(for)
  if $k = 0$ **then**
    $\diamond$ compute the restriction to $\mathbb{U}_0$ of the weighted likelihood (11) and draw $\mathbf{u}_0$ from it;
    $\diamond$ set $\mathbf{s}_1 \leftarrow \mathbf{s}_0 + \delta\mathbf{u}_0$;
  else
    $\diamond$ compute the restrictions to $\mathbb{U}_0$ of the weighted likelihood (8) and of the prior (9);
    $\diamond$ form the restriction to $\mathbb{U}_0$ of the posterior (10) and draw $\mathbf{u}_k$ from it;
    $\diamond$ set $\mathbf{s}_{k+1} \leftarrow \mathbf{s}_k + \delta\mathbf{u}_k$;
    $\diamond$ estimate the curvature κ at $\mathbf{s}_k$ of an interpolation of the fiber path $\mathbf{s}_0, \ldots, \mathbf{s}_{k+1}$;
  end(if)
  $\diamond$ set $k \leftarrow k+1$;
end(while)

3. EXPERIMENTS

3.1. The data

We assess our method with both synthetic and human cardiac DTI data.

The synthetic data (see Fig. 2) is generated from a spiral-shaped tensor field of size $50 \times 50 \times 20$. The diffusion tensor at voxel v is given by

$$\mathbf{D}_\mathbf{v} = \mathbf{e}_{1,v}\mathbf{e}_{1,v}^T + \frac{1}{2}\big(\mathbf{e}_{2,v}\mathbf{e}_{2,v}^T + \mathbf{e}_{3,v}\mathbf{e}_{3,v}^T\big), \qquad (13)$$

where the unit eigenvectors $\mathbf{e}_{1,v}$, $\mathbf{e}_{2,v}$ and $\mathbf{e}_{3,v}$ are defined as follows: the principal eigenvector $\mathbf{e}_{1,v}$ is perpendicular to the radial direction (the angle between $\mathbf{e}_{1,v}$ and the xy plane ranges from $-70°$ to $+70°$), $\mathbf{e}_{2,v}$ is parallel to the xy plane and points towards the z-axis, and $\mathbf{e}_{3,v} = \mathbf{e}_{1,v} \wedge \mathbf{e}_{2,v}$. The corresponding DW volumes are computed by using the Stejskal-Tanner diffusion equations with a constant non-diffusion-weighted signal.

The real data consists in ex-vivo human cardiac data acquired with a Siemens 1.5 T Magnetom Avanto in the Cardiology Hospital of Lyon with the following settings: TE = 98 ms, TR = 8600 ms, FOV = $256 \times 256\,\mathrm{mm}^2$, slice thickness = 2 mm, number of slices = 52, slice size = 128×128, diffusion sensitivity $b = 1000\ \mathrm{s{\cdot}mm^{-2}}$, gradient directions = 12. The sequence used was a 2D EPI diffusion-weighted sequence with a twice-refocused diffusion preparation.

3.2. Parameter setting and performance evaluation

We compare the performance of the proposed method to the standard streamlining approach. For our method, we take $\mathcal{G}$ to be the 26 nearest-neighbor system, the attenuation function ρ is defined by $\rho(t) = \exp(-t^2)$, the parameters α and β of the weighting function are set to 2.0, the parameter γ of the prior is set to 1.0, the step length δ is set to half the voxel size (that is, 1 mm for the real cardiac data), and the curvature threshold τ used for the termination of the tracking process is set to $1/\delta$. The streamlining approach uses a fourth-order Runge-Kutta integration scheme with integration step δ, and the criteria used for stopping the tracking are thresholds on FA (FA < 0.1) and on curvature ($\kappa \leqslant 1/\delta$).

To assess the relative performances of the fiber tracking methods in the case of the synthetic data, we measure the adequation of a reconstructed fiber path $\mathfrak{F}$ to the noise-free tensor distribution by computing the error term

$$\mathcal{E}(\phi) = \frac{1}{L(\phi)}\int_0^L \left(1 - \frac{\|\mathbf{D}_{\phi(l)}\,\phi'(l)\|}{\|\phi'(l)\|\,|||\mathbf{D}_{\phi(l)}|||}\right)\|\phi'(l)\|\,\mathrm{d}l\,, \qquad (14)$$

where $\phi : [0, L] \longrightarrow \mathbb{R}^3$ is a quadratic spline interpolation of $\mathfrak{F}$, $L(\phi) = \int_0^L \|\phi'(l)\|\,\mathrm{d}l$ denotes the length of ϕ, $\mathbf{D}_{\phi(l)}$ is the tensor at $\phi(l)$, and $|||\cdot|||$ is the spectral norm induced by the standard Euclidean norm $\|\cdot\|$ (the term between parenthesis is equal to one if $\mathrm{FA}(\phi(l)) = 1$ and if $\phi'(l)$ is parallel to the diffusion direction, and it is zero if $\mathrm{FA}(\phi(l)) = 0$ or if $\mathrm{FA}(\phi(l)) = 1$ and $\phi'(l)$ is perpendicular to the diffusion direction). This error term is averaged over all generated fiber paths to obtain a mean fiber error measure.

In the case of the real data, the noise-free tensor distribution is not available and thus (14) cannot be used. We therefore consider the following quantitative criteria: number of fibers with length greater than 6 mm, mean fiber-length, and standard deviation of the fiber length.

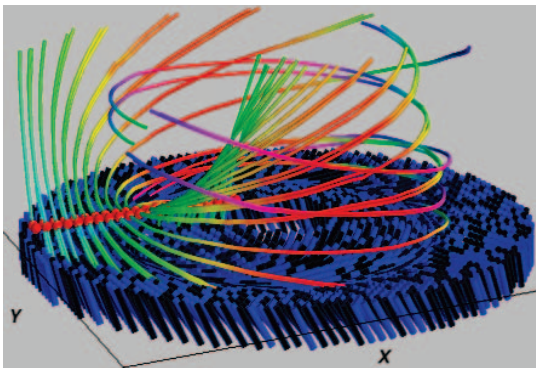


Fig. 2. Fibers paths generated by the proposed method from the 3-D noise-free synthetic tensor field, starting from the red point seeds (the cylinders represent the principal eigenvectors of the tensors on the xy plane).

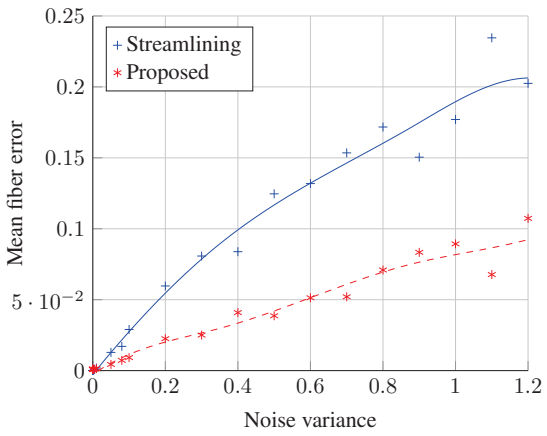


Fig. 3. Mean fiber error as a function of the noise level for the streamlining method and for the proposed method.

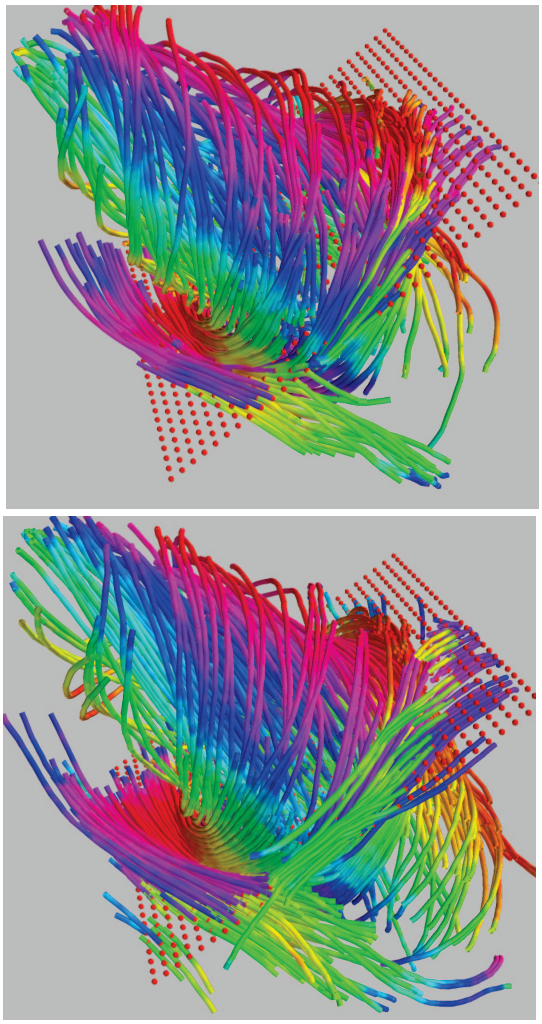


Fig. 4. Fibers tracked from the red point seeds using the streamlining method (top) and the proposed method (bottom).

3.3. Results

Fig. 3 shows the mean fiber error as a function of the noise variance for the streamlining method and for the proposed method (the values of the noise variance are normalized by the maximum DW intensity). For both methods, the fibers were initialized from the red point seeds shown in Fig. 2. We observe that the fibers generated by our method are always closer to the noise-free tensor data (in terms of the error measure (14)) than are the fibers produced by streamlining. Furthermore, the higher the noise level, the higher the improvement brought by our method over streamlining.

The fiber tracking results associated with our human cardiac data are shown in Fig. 4. The proposed method produces more regular and smoother fibers than does streamlining, as can be observed in the boxed regions. Quantitative results are given in Table 1 where both methods are compared in terms of the number of fibers, the mean fiber-length, and the fiber-length standard deviation. The proposed method outperforms the streamlining method in terms of all three criteria.

4. CONCLUSION

We have proposed a probabilistic fiber tracking method that exploits spatial correlation to better account for data uncertainty in human cardiac DTI data. The method is based on the estimation of a weighted mean likelihood over the spatial neighborhood of the current fiber position. Experimental results on both synthetic and real human cardiac DTI data showed that the proposed method is meaningfully more robust to noise than the streamlining method, and that it produces more regular and smoother fibers, which enables cardiac fiber configurations to be more clearly observed.

Table 1. Quantitative comparison between the streamlining method and the proposed method (real cardiac data)

	Number of fibers	Mean fiber-length	Fiber-length standard deviation
Streamlining	775	24.70 mm	18.01 mm
Proposed	1000	27.19 mm	17.88 mm

5. ACKNOWLEDGMENTS

This work was partly supported by French ANR 2009 (under ANR-09-BLAN-0372-01).

6. REFERENCES

[1] P J Basser, J Mattiello, and D LeBihan, "Estimation of the effective self-diffusion tensor from the NMR spin echo," *Journal of Magnetic Resonance. Series B*, vol. 103, no. 3, pp. 247–254, 1994.

[2] P Basser, S Pajevic, C Pierpaoli, J Duda, and A Aldroubi, "In vivo fiber tractography using DT-MRI data," *Magnetic Resonance in Medicine*, vol. 44, no. 4, pp. 625–632, 2000.

[3] P Fillard, M Descoteaux, A Goh, S Gouttard, B Jeurissen, J Malcolm, A Ramirez-Manzanares, M Reisert, K Sakaie, F Tensaouti, T Yo, J.-F. Mangin, and C Poupon, "Quantitative evaluation of 10 tractography algorithms on a realistic diffusion MR phantom," *NeuroImage*, vol. 56, no. 1, pp. 220–234, 2011.

[4] C Frindel, M Robini, J Schaerer, P Croisille, and Y.-M. Zhu, "A graph-based approach for automatic cardiac tractography," *Magnetic Resonance in Medicine*, vol. 64, no. 4, pp. 1215–1229, 2010.

[5] T Behrens, M Woolrich, M Jenkinson, H Johansen-Berg, R Nunes, S Clare, P Matthews, J Brady, and S Smith, "Characterization and propagation of uncertainty in diffusion-weighted MR imaging," *Magnetic Resonance in Medicine*, vol. 50, no. 5, pp. 1077–1088, 2003.

[6] O Friman, G Farnebck, and C.-F. Westin, "A bayesian approach for stochastic white matter tractography," *IEEE Transactions on Medical Imaging*, vol. 25, no. 8, pp. 965–978, 2006.

[7] R. Salvador, A. Pena, D. K Menon, T. A Carpenter, J. D Pickard, and E. T Bullmore, "Formal characterization and extension of the linearized diffusion tensor model," *Human brain mapping*, vol. 24, no. 2, pp. 144155, 2005.

[8] O. Friman and C.-F Westin, "Uncertainty in white matter fiber tractography," in *Proc. 8th Int. Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI'05)*, pp. 107–114. Berlin, Heidelberg, 2005.