Stochastic Tractography in 3-D Images via Nonlinear Filtering and Spherical Clustering

H. Ertan Çetingül[1], Gernot Plank[2], Natalia Trayanova[3], and René Vidal[1]

[1] Center for Imaging Science, Johns Hopkins University, USA
[2] Institute of Physiology, Medical University of Graz, Austria
[3] Institute for Computational Medicine, Johns Hopkins University, USA
{ertan,ntrayanova,rvidal}@jhu.edu, gernot.plank@meduni-graz.at

Abstract. This paper presents a stochastic tractography algorithm to identify branching fibrous structures in 3-D images. Specifically, we employ a Bayesian formulation that involves sequential importance resampling of multiple fiber trajectories (particles). At each step, we sample local orientation candidates from a discrete importance density that is estimated by utilizing a nonlinear oriented filter. The weights of the particles are subsequently obtained by using an observation density, which models the intensity coherence along the particle segments of inferred widths, and a prior density that enforces local smoothness of the trajectories. At the same time, we successfully build multiple fiber trajectories along all branches of a fiber by detecting dominant local fiber orientations via the mean shift algorithm. The performance of the proposed method is evaluated by tracking synthetic fibers at different noise levels as well as by extracting selected free-running cardiac Purkinje fibers in structural magnetic resonance images. Our experiments show that the method achieves a mean tracking error of about 4 voxels in the case of Purkinje fibers and remains robust to moderate amount of noise.

1 Introduction

The development of robust processing methods to quantitatively characterize fibrous structures constitutes an important yet challenging problem in medical image analysis. Specifically, the extraction of complex fiber networks in intensity data finds a wide range of applications in biological studies (see [1, 2] and references therein). For instance, by extracting the anatomical topology of different structures such as the human vasculature or pulmonary airways, one could perform a quantitative assessment of various pathologies. Furthermore, advanced simulations of arrhythmias could benefit from an electrophysiological model of the heart with a realistic conduction system [3]. This would involve extracting the Purkinje system, which is responsible for the propagation of the electrical impulse initiating the contraction of the ventricular myocardium. Modern ex vivo magnetic resonance imaging (MRI) techniques provide sufficient resolution to identify the *free-running* Purkinje fibers, which activate endocardial structures such as the papillary muscle. However, tracking these fibers in an automated manner is difficult due to the presence of numerous bifurcations and image noise.

From an image processing standpoint, a fibrous structure comprises a spatially coherent appearance pattern that can be quantified via different feature-based or model-based approaches. Typical examples reviewed in [4] include skeletons, matched filters, region growing, active contours, and image Hessian-guided streamline tracking. However, most of these methods are *deterministic* in the sense that the uncertainties associated with the resulting fibers are not inferred. In addition, they do not usually incorporate a priori geometric knowledge. Similarly, the majority of the existing tracking techniques on diffusion weighted imaging (DWI) estimate fiber tracts to be proportional to principal diffusion directions [5]. However, despite their practicality for visualization, they also do not estimate the aforementioned uncertainties, which are critical for eliminating anatomically incorrect trajectories caused by local tracking errors and their accumulation.

These issues have inspired recent works [6–10], which aim at *stochastically* tracking fibrous structures at the expense of increased computational complexity. These methods choose the most probable fiber trajectory among multiple possible trajectories that are generated by statistical sampling. For instance, [7] proposes a bootstrap filtering scheme to track the cross-sections of non-branching cerebral arteries. In order to extract non-branching vessels, [8] employs a particle filtering formulation with a sophisticated observation density that requires several parameters to be tuned, whereas [9] proposes a modified Bayesian formulation that requires fewer particles. For the analysis of branching fibers, [6] and [10] employ the k-means algorithm to cluster the trajectories of a particle filter to segment coronary arteries and brain vasculature, respectively. However, the extracted vessel topology is limited since k-means requires the number of clusters (branches) to be known beforehand. Moreover, there exist several works providing different stochastic tracking approaches on DW images [11–14]. In particular, [12] proposes a Bayesian formulation that draws samples from the posterior density, whereas [14] employs sequential importance resampling to track white matter fibers. [15] estimates complex fiber orientation distributions from high angular resolution diffusion images and performs random walks to infer brain connectivity.

We believe that the efficiency of stochastic tractography methods can be substantially improved by developing 1) practical probabilistic models for sampling, and 2) accurate clustering techniques for detecting bifurcations. We thereby present such a stochastic approach to track branching fibrous structures in 3-D images. Our contribution is to fuse the notions of filter-based local orientation extraction and multimodal fiber orientation distributions into a particle filtering formulation to infer the uncertainties of the fiber trajectories. Specifically, we utilize a nonlinear *pivoting filter* that estimates a reliable discrete importance density for sampling trajectories (particles). Furthermore, we use this density for detecting bifurcations by clustering the local fiber orientations via *spherical mean shift*. We recursively compute the weights of the particles using an observation density that quantifies intensity coherence along the particle segments of inferred widths and a prior density that enforces locally smooth trajectories. We evaluate our tracking method on synthetic fibers at different signal-to-noise ratios (SNRs) as well as on selected free-running cardiac Purkinje fibers in structural MRI.

2 Preliminaries on Stochastic Tracking

Our tracking scheme follows a well-known particle filtering technique called *sequential importance resampling* (SIR), which implements a recursive Bayesian filter by Monte Carlo simulations [16, 17]. In the SIR formulation, we represent a fiber trajectory as a sequence of 3-D points (states) $\boldsymbol{x}_{0:T} \doteq \{\boldsymbol{x}_t\}_{t=0}^{T}$ in the image domain $\Upsilon \subset \mathbb{R}^3$. Without loss of generality, we assume that the root of the fiber of interest is determined, either by the user or through some prior automatic detection. We define the observations as a sequence of image intensities $y_{1:T} \doteq \{y_t\}_{t=1}^{T} = \{I(\boldsymbol{x}_t)\}_{t=1}^{T}$, where $I(\boldsymbol{p})$[4] denotes the intensity value at voxel $\boldsymbol{p} \in \Upsilon$. We further assume that 1) the states are modeled as a first-order Markov process, 2) the observations are mutually independent, and 3) the observation at time t only depends on the current state, i.e., $p(y_t|\boldsymbol{x}_{0:t}) = p(y_t|\boldsymbol{x}_t)$.

The key idea in SIR is to represent the required posterior density $p(\boldsymbol{x}_{0:t}|y_{1:t})$ by a set of N weighted random samples $\{\boldsymbol{x}_{0:t}^{(n)}\}_{n=1}^{N}$ and the associated weights $\{w_t^{(n)}\}_{n=1}^{N}$, which are chosen using the principle of *importance sampling*. This specifically involves defining an *importance density* $q(\cdot)$ from which the random samples can be easily drawn. Accordingly, the weights measure the reliability of the samples as $w_t^{(n)} \propto p(\boldsymbol{x}_{0:t}^{(n)}|y_{1:t})/q(\boldsymbol{x}_{0:t}^{(n)}|y_{1:t})$.

In a sequential setting, we consider N samples that represent $p(\boldsymbol{x}_{0:t-1}|y_{1:t-1})$ and then approximate $p(\boldsymbol{x}_{0:t}|y_{1:t})$ with a new set of samples. Using Bayes' rule, a recursive formula for the posterior density can be written in terms of the *observation density* $p(y_t|\boldsymbol{x}_t)$ and the *prior density* $p(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$ as

$$p(\boldsymbol{x}_{0:t}|y_{1:t}) \propto p(\boldsymbol{x}_{0:t-1}|y_{1:t-1})p(y_t|\boldsymbol{x}_t)p(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}). \tag{1}$$

If the importance density satisfies $q(\boldsymbol{x}_{0:t}|y_{1:t}) = q(\boldsymbol{x}_{0:t-1}|y_{1:t-1})q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}, y_{1:t})$, the formulation only requires a filtered estimate $p(\boldsymbol{x}_t|y_{1:t})$ at each time step. Consequently, we can recursively compute the weights as

$$w_t^{(n)} \propto w_{t-1}^{(n)} \frac{p(y_t|\boldsymbol{x}_t^{(n)})p(\boldsymbol{x}_t^{(n)}|\boldsymbol{x}_{t-1}^{(n)})}{q(\boldsymbol{x}_t^{(n)}|\boldsymbol{x}_{t-1}^{(n)}, y_{1:t})}. \tag{2}$$

Then the posterior filtered density $p(\boldsymbol{x}_t|y_{1:t})$ can be approximated as

$$p(\boldsymbol{x}_t|y_{1:t}) \approx \hat{p}(\boldsymbol{x}_t|y_{1:t}) = \sum_{n=1}^{N} w_t^{(n)} \delta(\boldsymbol{x}_t - \boldsymbol{x}_t^{(n)}), \tag{3}$$

where $\delta(\cdot)$ denotes the Dirac delta function and the weights are normalized such that $\sum_{n=1}^{N} w_t^{(n)} = 1$. As the number of samples $N \to \infty$, the approximate posterior $\hat{p}(\boldsymbol{x}_t|y_{1:t})$ approaches the true posterior $p(\boldsymbol{x}_t|y_{1:t})$.

It is also worth noting that during SIR, the distribution of the weights becomes more skewed as t increases [16]. To avoid this degeneracy, a resampling procedure

[4] When the point $\boldsymbol{p}$ lies outside the discrete grid, we compute the corresponding intensity value by trilinear interpolation.

is usually included at each step. More specifically, if the *effective sample size* $N_{\text{eff}} \doteq \lceil 1/\sum_{n=1}^{N}(w_t^{(n)})^2 \rceil$ is less than a fixed threshold τ, we resample N particles from the discrete density $\hat{p}(\boldsymbol{x}_t|y_{1:t})$ and set $w_t^{(n)} = 1/N, \forall n$. The reader is referred to [16, 17] for further details on sequential Monte Carlo methods.

3 Estimation of the Density Functions in SIR

This section explains the estimation of the importance, observation and prior densities in the SIR formulation. These densities are primarily used to recursively compute the importance weights of the particles. For the sake of clarity, we drop the particle superscripts in the following discussions.

3.1 Importance Density

In the SIR formulation, modeling the importance density with the prior and solely using the observation density to compute the weights from (2) constitute a popular yet trivial choice. However, we believe that sampling from a density that models local fiber structure is more appropriate for tracking purposes. We thereby estimate the importance density $q(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}, y_{1:t})$ using a nonlinear *pivoting filter*, which is a 3-D variant of the matched filter introduced in [18]. Its efficacy in extracting local fiber orientations has been recently demonstrated in [19].

As depicted in Fig. 1(a), the pivoting filter is centered at a point of interest $\boldsymbol{x}$ with a fixed backward point $\boldsymbol{b}$ and a pivoting forward point $\boldsymbol{f}$ located at a distance l from $\boldsymbol{x}$. The segments $\overline{\boldsymbol{b}\boldsymbol{x}}$ and $\overline{\boldsymbol{x}\boldsymbol{f}}$ define the main segments of the filter and are designed to align with the fiber of interest. The purpose of the remaining points $\{\boldsymbol{b}_k, \boldsymbol{f}_k\}_{k=1}^{2K}$ is to fully encapsulate the fiber. Specifically, the points $\{\boldsymbol{b}_k\}$ (and $\{\boldsymbol{f}_k\}$) are placed by rotating the $2w$-apart antipodal pair $(\boldsymbol{b}_k, \boldsymbol{b}_{k+K})$ (and $(\boldsymbol{f}_k, \boldsymbol{f}_{k+K})$) by an angular step α until one circular round is completed.

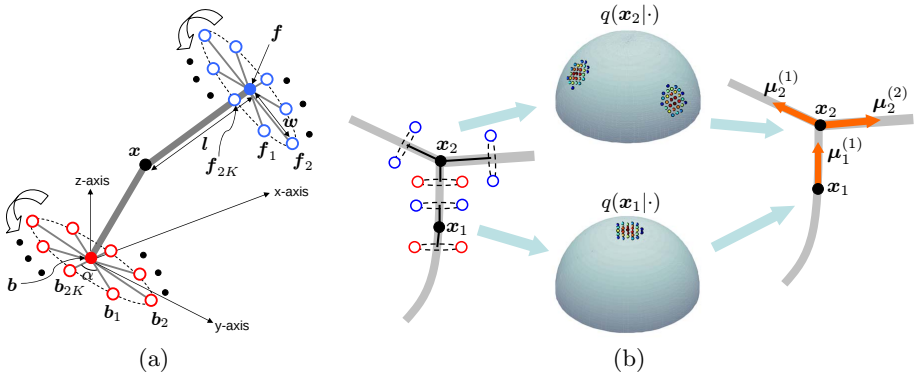

Fig. 1. (a) 3-D pivoting filter with key parameters, (b) Discrete importance densities estimated at two analysis points and the modes detected via spherical mean shift.

The operation of the filter can be summarized as follows: After fixing the backward segment $\overline{\boldsymbol{b}\boldsymbol{x}}$ along the known portion of the fiber, the forward segment

$\overline{\boldsymbol{xf}}$ is allowed to rotate for scanning multiple local orientations. Specifically, consider a particular orientation $\boldsymbol{s}$ in $\mathcal{S} \doteq \{\boldsymbol{s} : \langle \boldsymbol{s}, \boldsymbol{x} - \boldsymbol{b} \rangle > 0, \boldsymbol{s} \in \mathbb{S}^2\}$ such that $\boldsymbol{f} = \boldsymbol{x} + l\boldsymbol{s}$. For each pair of antipodes $(\boldsymbol{b}_k, \boldsymbol{b}_{k+K})$ and $(\boldsymbol{f}_k, \boldsymbol{f}_{k+K})$, we compute

$$h_k(\boldsymbol{s}; \boldsymbol{x}, \boldsymbol{x} - \boldsymbol{b}, I) = \begin{cases} 1 & \text{if } |I(\boldsymbol{b}) - I(\boldsymbol{f})| \leq \min\limits_{j=k,k+K} \{|I(\boldsymbol{b}) - I(\boldsymbol{b}_j)|, |I(\boldsymbol{f}) - I(\boldsymbol{f}_j)|\} \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

In other words, the filter detects an oriented structure if the absolute value of the intensity variation along the structure is less than the minimum absolute intensity variation orthogonal to the structure. The expression in (4) can be considered as a *partial* filter response for a fixed $k \in \{1, 2, \ldots, K\}$. The overall filter response is subsequently computed by summing (4) over all pairs of antipodes as $h(\boldsymbol{s}; \boldsymbol{x}, \boldsymbol{x} - \boldsymbol{b}, I) = \sum_{k=1}^{K} h_k(\boldsymbol{s}; \boldsymbol{x}, \boldsymbol{x} - \boldsymbol{b}, I)$.

In the sequential setting, having fixed the backward segment along the preceding orientation, i.e., $\boldsymbol{s}_{t-1} \doteq (\boldsymbol{x}_{t-1} - \boldsymbol{x}_{t-2})/l$, the filter response at $\boldsymbol{x}_{t-1}$ gives a coarse estimate of the probability of having a structure oriented along any $\boldsymbol{s} \in \mathcal{S}$. This is further refined by using $g(\boldsymbol{s}; \boldsymbol{x}, \boldsymbol{x} - \boldsymbol{b}, I) = \int_0^1 |I(\boldsymbol{x} + \lambda l \boldsymbol{s}) - I(\boldsymbol{x})|^2 \mathrm{d}\lambda$, which quantifies the intensity coherence along the forward segment [4]. The importance density is then estimated as

$$q(\boldsymbol{x}_t | \boldsymbol{x}_{t-1}, y_{1:t}) \propto \Big(h(\boldsymbol{s}_t; \boldsymbol{x}_{t-1}, \boldsymbol{s}_{t-1}, I) \times \exp(-g(\boldsymbol{s}_t; \boldsymbol{x}_{t-1}, \boldsymbol{s}_{t-1}, I)) \Big). \tag{5}$$

Given $q(\cdot)$, we can sequentially sample N local orientations $\{\boldsymbol{s}_t^{(n)}\}_{n=1}^N$ that also define the points $\{\boldsymbol{x}_t^{(n)}\}$ given $\{\boldsymbol{x}_{t-1}^{(n)}\}$. The resulting discrete density can be further used to detect branches via spherical clustering, as described in §4.

3.2 Observation Density

Given a current sample $\boldsymbol{s}_t$ (and $\boldsymbol{x}_t$), we obtain an estimate of the observation density $p(y_t | \boldsymbol{x}_t)$. For this purpose, we assume that 1) the fiber portion along $\boldsymbol{s}_t$ has the same width (diameter) as the preceding portion, and 2) the intensities of the voxels in a neighborhood of $\boldsymbol{x}_t$, denoted by $\Omega_{\boldsymbol{x}_t}$, are also observed.

We first find the cylinder that optimally encapsulates the preceding fiber portion and then use it to find the voxels around the segment $l\boldsymbol{s}_t = \boldsymbol{x}_t - \boldsymbol{x}_{t-1}$. Specifically, we encapsulate the fiber portion along $\boldsymbol{s}_{t-1}$ with (oriented) cylinders of height l and of different radii, and estimate the optimal radius r^e as

$$r^e = \underset{r \in \mathcal{R}}{\operatorname{argmax}} \ \frac{1}{2\pi l r} \int_{\mathcal{F}(r)} \Big(\langle \nabla I(\boldsymbol{p}), \overrightarrow{N}(\boldsymbol{p}) \rangle \Big)^2 \mathrm{d}\boldsymbol{p}. \tag{6}$$

Notice that the expression in (6) measures the alignment of the image gradients with the normals at the cylindrical surface, where $\mathcal{R}$ denotes the set of radii, $\boldsymbol{p}$ is a point on the lateral surface $\mathcal{F}(r)$ of the cylinder of radius r, $\nabla I(\boldsymbol{p})$ is the image gradient and $\overrightarrow{N}(\boldsymbol{p})$ is the outward normal to the surface $\mathcal{F}(r)$ at point $\boldsymbol{p}$.

We subsequently use the inferred radius to form a cylinder (of radius r^e and height l) oriented along $\boldsymbol{s}_t$, define $\Omega_{\boldsymbol{x}_t}$ as the voxels inside this cylinder, and compute the observation density as

$$p(y_t|\boldsymbol{x}_t) \propto \exp\left(-\frac{(I(\boldsymbol{x}_t) - \bar{I}_{t;r^e})^2}{2\sigma^2_{t;r^e}}\right). \tag{7}$$

Here $\bar{I}_{t;r^e}$ and $\sigma_{t;r^e}$ denote the mean and standard deviation of the intensities of the voxels inside the cylinder, respectively.

3.3 Prior Density

In order to enforce a certain level of smoothness in the fiber trajectories, we choose to model the prior density based on the von Mises-Fisher (vMF) distribution over the 2-sphere [20]. It constitutes a unimodal parametric distribution for directional data with the probability density function (pdf) of the form

$$p_{\text{vMF}}(\boldsymbol{s}; \boldsymbol{\mu}, \kappa) = \frac{\kappa}{4\pi \sinh \kappa} \exp(\kappa \langle \boldsymbol{\mu}, \boldsymbol{s} \rangle), \tag{8}$$

where $\boldsymbol{\mu} \in \mathbb{S}^2$ is the mean direction and $\kappa > 0$ is a parameter regulating the concentration around the mean direction. In our experiments, we manually set the value of κ and estimate the prior density as

$$p(\boldsymbol{x}_t|\boldsymbol{x}_{t-1}) \propto p_{\text{vMF}}(\boldsymbol{s}_t; \boldsymbol{s}_{t-1}, \kappa). \tag{9}$$

4 Algorithm Overview and Implementation Details

Our method is initiated at two user-specified seed points $\{\boldsymbol{x}_0, \boldsymbol{x}_1\}$, which place the backward segment $\overline{\boldsymbol{bx}}$ of the filter along $\boldsymbol{s}_1 = (\boldsymbol{x}_1 - \boldsymbol{x}_0)/l$. We subsequently estimate the width of this fiber segment from (6). In order to have a finite number of search orientations, we discretize the unit sphere at 642 predefined vectors obtained by a threefold tessellation of an icosahedron. Accordingly, at each step, we form the set of candidate orientations $\mathcal{S}_t \doteq \{\boldsymbol{s} : \langle \boldsymbol{s}, \boldsymbol{s}_{t-1} \rangle > 0, \boldsymbol{s} \in \mathbb{S}^2\}$, estimate the discrete importance density as described in §3.1, and use the resulting pdf to analyze local fiber structure, i.e., to detect branches at the point of interest.

The rationale behind analyzing local fiber structure is that in the case of branching fibers, one of the branches may cause the loss of the remaining ones by attracting most of the particles. To avoid this problem, before sampling the particles, we employ the mean shift (MS) algorithm [21], which automatically detects the number and directions of the branches. Specifically, we utilize a weighted spherical MS formulation with the Fisher kernel for clustering local orientations using their importance density values as weights [19]. This scheme converges to the modes $\{\boldsymbol{\mu}_t^{(c)}\}_{c=1}^{C}$ of the importance density $q(\boldsymbol{x}_t|\cdot)$ and the resulting modes are identified as the directions of the branches rooted at $\boldsymbol{x}_t$. Fig. 1(b) illustrates the operation of the filter and our branch detection strategy.

Algorithm 1 SIR-based tractography in 3-D images

1. At the t-th step, given $N = 1000$ weighted particles $\{(\boldsymbol{x}_{0:t-1}^{(n)}, w_{t-1}^{(n)})\}_{n=1}^{N}$ and the mean trajectory $\{\bar{\boldsymbol{x}}_{0:t-1}\}$, consider $\boldsymbol{s}_{t-1} = \bar{\boldsymbol{x}}_{t-1} - \bar{\boldsymbol{x}}_{t-2}$.
2. Place the pivoting filter such that $\boldsymbol{b} = \bar{\boldsymbol{x}}_{t-2}$ and $\boldsymbol{x} = \bar{\boldsymbol{x}}_{t-1}$.
3. Obtain the set $\mathcal{S}_t = \{\boldsymbol{s} : \langle \boldsymbol{s}, \boldsymbol{s}_{t-1} \rangle > 0, \boldsymbol{s} \in \mathbb{S}^2\}$.
4. Estimate the discrete importance density $\hat{q}(\boldsymbol{x}_t | \bar{\boldsymbol{x}}_{t-1}, y_{1:t})$ from (10) using (5).
5. Perform branch analysis by detecting the modes $\{\boldsymbol{\mu}_t^{(c)}\}_{c=1}^{C}$ of the importance density $\hat{q}(\boldsymbol{x}_t|\cdot)$ via spherical MS and check the number of branches C:
   - If $C > 1$, stop tracking the "parent" fiber, go to step 9 and generate separate SIR schemes for the branches.
   - If $C = 1$, sample N points $\{\tilde{\boldsymbol{x}}_t^{(n)}\}$ from the importance density.
6. For each $\tilde{\boldsymbol{x}}_t^{(n)}$, compute the weight $w_t^{(n)}$ from (2) using (5), (7), and (9) with $\kappa = 3$.
7. Calculate N_{eff} as described in §2, set $\tau = 2N/3$, and check N_{eff}:
   - If $N_{\text{eff}} \geq \tau$, then for $n = 1, 2, \ldots, N$, set $\boldsymbol{x}_t^{(n)} = \tilde{\boldsymbol{x}}_t^{(n)}$ and update the n-th trajectory as $\boldsymbol{x}_{0:t}^{(n)}$.
   - If $N_{\text{eff}} < \tau$, then for $n = 1, 2, \ldots, N$, sample an index $z(n)$ from the discrete distribution of $\{w_t^{(n)}\}_{n=1}^{N}$, set $\boldsymbol{x}_t^{(n)} = \tilde{\boldsymbol{x}}_t^{z(n)}$, $w_t^{(n)} = 1/N$ and update the n-th trajectory as $\boldsymbol{x}_{0:t}^{(n)}$.
8. Iterate between 1-8 by setting $t = t + 1$ until $\{w_t^{(n)}\}$ are small, a user-defined stopping criterion is met and/or a branching point is detected.
9. Identify the MAP fiber (or branch) as $\boldsymbol{x}_{0:t}^{(\zeta)}$, where $\zeta = \text{argmax}_{n=1,2,\ldots,N}\, w_t^{(n)}$, i.e., the trajectory with the maximum importance weight.

It depicts the importance densities, where the probabilities of the vectors are color-coded (blue∼low, red∼high), at two points as well as the resulting modes.

At this point, it is worth noting that the pivoting filter needs to be used N times per iteration to estimate the importance density given N preceding particles. This computational load can be reduced with the approximation

$$q(\boldsymbol{x}_t^{(n)} | \boldsymbol{x}_{t-1}^{(n)}, y_{1:t}) \approx \hat{q}(\boldsymbol{x}_t^{(n)} | \bar{\boldsymbol{x}}_{t-1}, y_{1:t}), \tag{10}$$

where $\bar{\boldsymbol{x}}_{t-1}$ is the point reached along the mode of interest, i.e., $\boldsymbol{\mu}_{t-1}^{(c)}$ for some c. After estimating the approximate importance density $\hat{q}(\boldsymbol{x}_t|\cdot)$ at the t-th step and clustering the local orientations, $N = 1000$ samples are drawn from this discrete density if the number of the modes, denoted by C, is equal to 1. However, if $C > 1$, the newly identified branches are treated as separate fibers to be tracked via new SIR schemes. Finally, we compute the importance weight $w_t^{(n)}$ from (2) using (5), (7), and (9) with $\kappa = 3$, and then eliminate the degeneracy described in §2 using *systematic resampling* with $\tau = 2N/3$. We repeat the same procedure by shifting the filter according to the modes of the importance density, i.e., keeping track of the "mean" trajectory $\{\bar{\boldsymbol{x}}_{0:t}\}$, and generate N different trajectories $\{\boldsymbol{x}_{0:T}^{(n)}\}_{n=1}^{N}$ for each branch. The maximum a posteriori (MAP) estimate of the true fiber/branch is identified as the trajectory with the maximal importance weight. In summary, our tracking method proceeds as outlined in Algorithm 1.

5 Method Validation

The performance of our method is evaluated via experiments on synthetic fibers at different noise levels as well as on selected free-running cardiac Purkinje fibers on structural MR images. In order to quantify the spatial tracking error, we compute the symmetrized Chamfer distance (in voxels) between the true fiber trajectory $\mathcal{X}^t \doteq \boldsymbol{x}^t_{0:T} = \{\boldsymbol{x}^t_i\}$ and its MAP estimate $\mathcal{X}^e \doteq \boldsymbol{x}^e_{0:T'} = \{\boldsymbol{x}^e_j\}$ as $\epsilon(\mathcal{X}^t, \mathcal{X}^e) = [d(\mathcal{X}^t, \mathcal{X}^e) + d(\mathcal{X}^e, \mathcal{X}^t)]/2$, where

$$d(\mathcal{X}^t, \mathcal{X}^e) = |\mathcal{X}^t|^{-1} \sum_{\boldsymbol{x}^t_i \in \mathcal{X}^t} \min\{\|\boldsymbol{x}^t_i - \boldsymbol{x}^e_j\| : \boldsymbol{x}^e_j \in \mathcal{X}^e\}. \tag{11}$$

5.1 Experiments on Synthetic Fibers

In order to evaluate the performances of the fiber width estimator (6), orientation detection and tracking methods as a function of image noise, we generate 120 synthetic fibers of radii $r^t \in \{1, 2, 3\}$ by fitting cubic splines through four randomly selected points in a $80\times80\times80$ lattice. In the case of branching fibers, bifurcations are randomly selected on previously generated fibers to add further branches. The centerlines of the resulting fibers constitute the true trajectories $\mathcal{X}^t$. The binary images are then corrupted by Rician noise to obtain image data at five different SNRs. We set the filter parameters $\{l, w, \alpha\} = \{3, 2, 20^\circ\}$ (or $\{4, 4, 20^\circ\}$ depending on the value of r^t), the search radii $\mathcal{R} = \{1, 1.5, 2, \ldots, 5\}$, and subsequently obtain the width estimates $\{r^e_j\}$ and the MAP trajectory $\mathcal{X}^e$ of each fiber.

For performance evaluation, we compute the width estimation error rate $\xi(\mathcal{X}^t, \mathcal{X}^e) \doteq |\mathcal{X}^e|^{-1} \sum_j |r^t - r^e_j|/r^t$ along with tracking error ϵ. In fact, the value of the optimal radius r^e is also critical in order to obtain an accurate observation density. Although a slight underestimation of the fiber width is tolerable unless the resulting filter is totally "buried" into the fiber, it should be noted that the more the fiber width is overestimated, the worse the observation density is modeled. Therefore, we also compute the rate ξ_+ of overestimating the fiber width. Finally, in order to quantify the reliability of our orientation detection strategy at different noise levels, we calculate the orientation detection error rates as follows: In the analysis of non-branching fibers, this rate, denoted by γ_1, is computed as the ratio of the number of points (in $\mathcal{X}^e$) at which the MS locates more than one mode over the total number of points, i.e., the cardinality of $\mathcal{X}^e$. However, in the analysis of branching fibers, we solely focus on the identification of bifurcations. Specifically, we place the filter at the bifurcations, which constitute the set of points at which the MS locates two modes in the noise-free case, and compute the rate γ_2 of misidentified bifurcations at different SNRs.

Analysis of non-branching fibers: Table 1 shows the mean and the standard deviation of the tracking error ϵ and the width estimation error rate ξ along with the average values of the width overestimation rate ξ_+ and the orientation detection error rate γ_1 over 60 non-branching fibers at different SNRs. The symbol "$\boldsymbol{\infty}$" for SNR represents the noise-free case. First, we observe that the

mean width estimation error rates are about 0.09 with $\xi_{\max} = 0.36$ for SNR > 5 dB. These results validate the accuracy of our flux-based estimator at moderate amount of noise. Furthermore, we overestimate the true fiber widths at rates (ξ_+) less than 0.25 for SNR > 10 dB. In terms of tracking accuracy, we observe that for SNR > 10 dB, the mean tracking errors are about 2.80 voxels with $\epsilon_{\min} = 1.00$ and $\epsilon_{\max} = 13.78$. However, the performance starts degrading at higher levels of noise, especially in the case of tracking fiber segments with high curvature. Finally, for SNR > 5 dB, the average values of the rates of erroneously detected orientations (γ_1) are less than 0.07, which demonstrate the reliability of our filter-based orientation detection strategy under noisy conditions.

Table 1. Experimental results on non-branching fibers

Error parameter	**SNR (dB)**					
	∞	**30**	**20**	**15**	**10**	**5**
ξ	0.08±0.07	0.08±0.07	0.08±0.07	0.09±0.07	0.11±0.09	0.20±0.26
$(\xi_{\min}, \xi_{\max})$	(0,0.36)	(0,0.36)	(0,0.31)	(0,0.31)	(0,0.32)	(0,1.42)
ξ_+	0.23	0.24	0.23	0.25	0.31	0.39
ϵ	2.39±1.21	2.82±2.11	3.20±2.39	2.68±1.32	3.96±2.70	5.18±3.83
$(\epsilon_{\min}, \epsilon_{\max})$	(0.87,6.35)	(0.78,13.68)	(0.65,13.78)	(0.90,6.81)	(1.00,11.94)	(1.67,18.92)
γ_1	0.04	0.04	0.04	0.05	0.07	0.17

Analysis of branching fibers: We further perform tracking experiments on 60 remaining branching fibers as well as focus on testing our orientation detection strategy at the bifurcations. At this point, note that extracting such fibers in noisy data may result in misleading tracking errors if a bifurcation point is undetected. Therefore, we estimate the widths of the fibers and obtain their MAP trajectories only in the noise-free case. Table 2(a) shows the mean and the standard deviation of the width estimation error rate ξ and the tracking error ϵ over 60 branching fibers. We observe that the mean tracking error is 4.81 voxels with $\epsilon_{\min} = 2.46$ and $\epsilon_{\max} = 8.21$. Notice that the tracking performance is worse than the one obtained for non-branching fibers by approximately 2.50 voxels. This is mainly due to 1) a number of branches that are generated very close to each other, and 2) differences between the coordinates of the actual bifurcations and the detected ones. Furthermore, we overestimate the true fiber widths in only 18% of the fiber segments and achieve a width estimation error rate of 0.07±0.05 with $\xi_{\min} = 0$ and $\xi_{\max} = 0.22$. Finally, Table 2(b) shows the average rates of misidentified bifurcations (γ_2). Notice that for SNR > 10 dB, less than 20% of the bifurcations are misidentified, but the error rate γ_2 increases at lower SNRs.

Table 2. Experimental results on branching fibers

(a) Tracking and width estimation errors

ϵ	$(\epsilon_{\min}, \epsilon_{\max})$	ξ	$(\xi_{\min}, \xi_{\max})$
4.81±1.42	(2.46,8.21)	0.07±0.05	(0,0.22)

(b) Bifurcation misidentification

SNR (dB)	**30**	**20**	**15**	**10**	**5**
γ_2	0.10	0.15	0.18	0.25	0.33

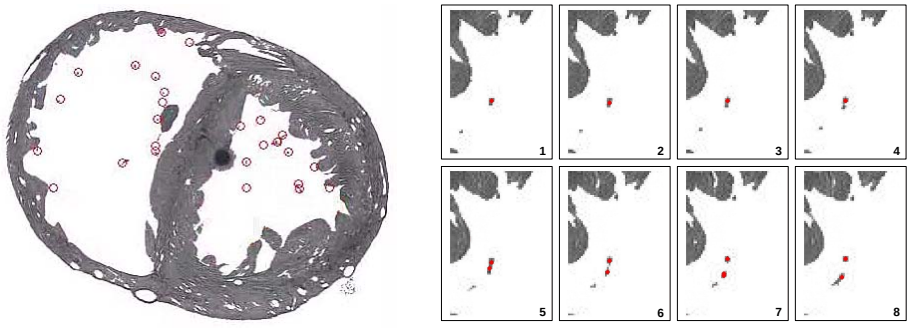

(a) (b)

Fig. 2. (a) An MR image slice with manually extracted Purkinje fibers (red circles), (b) Selected region of interests illustrating the trajectory of a branching Purkinje fiber.

5.2 Experiments on Cardiac Purkinje Fibers

We conduct additional tracking experiments on a structural MR image of a healthy rabbit heart, which was acquired on an 11.7 T MR system at an in-plane resolution of 26.5 μm $\times$ 26.5 μm and an out-of-plane resolution of 24.5 μm [3]. Denoising and rescaling steps are performed to obtain a 3-D image of size 512×512×850 (Fig. 2(a)). Our method is tested on selected subvolumes that contain 53 free-running Purkinje fibers. They are either non-branching fibers running from one Purkinje-myocardial junction (PMJ) to another or branching fibers with at least one Purkinje-Purkinje junction (PPJ), as depicted in Fig. 2(b). In addition, we manually extract the centerlines $\mathcal{X}^t$ of the true fibers and obtain the MAP trajectories $\mathcal{X}^e$ using the filter parameters $\{l, w, \alpha\} = \{4, 4, 20^\circ\}$.

In our experiments, we obtain a tracking error of 3.93±3.15 voxels with $\epsilon_{\min} = 0.78$ and $\epsilon_{\max} = 19.74$. Specifically, our method achieves such promising results that 16 fibers are tracked with errors of less than 2 voxels, and 32 fibers with errors of less than 4 voxels. In addition, all the bifurcations are correctly detected along with the corresponding branch directions to be followed. However, 2 fibers are tracked with errors greater than 10 voxels and we observe that, in those cases, tracking is affected by the presence of nearby fibers/cardiac wall. Specifically, the pivoting filter fails at encapsulating the fibers of interest and this causes an inaccurate estimation of the importance densities. Fig. 3 illustrates our tracking results by showing the surface renderings (green) of selected subvolumes, the seed points (black), and the resulting MAP trajectories (red) of the fibers.

6 Conclusions

The aim of this work is to stochastically track branching fibrous structures in 3-D images. Specifically, we present an SIR-based tracking framework that employs an oriented filter for the estimation of a practical discrete importance density. We also describe a flux-based fiber width estimator to obtain an observation density that models the intensity coherence in the fiber segment of interest.

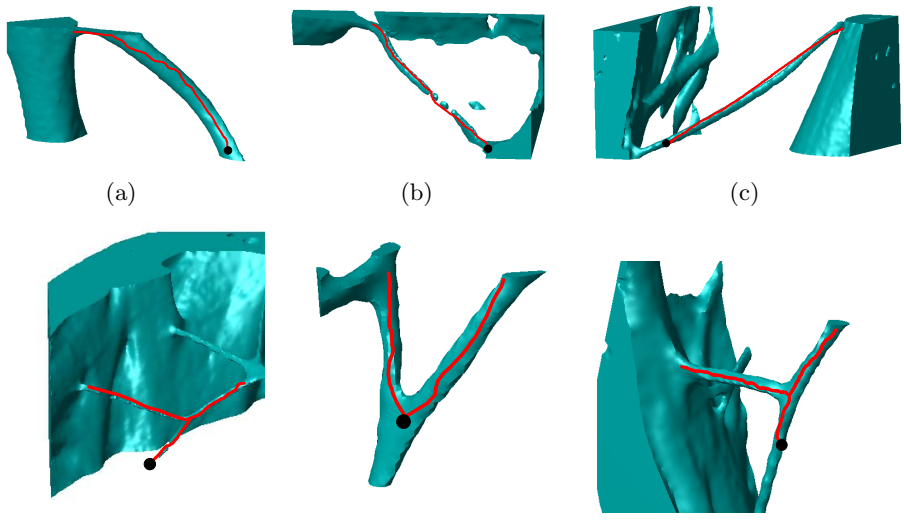

Fig. 3. Visualization of 6 free-running Purkinje fibers and their MAP trajectories (red).

Along with a prior density, which enforces locally smooth fiber trajectories, the aforementioned densities are used to recursively compute the certainties of the trajectories. Furthermore, in order to correctly analyze branching structures, we sample trajectories along all branches by detecting dominant branch directions via spherical mean shift. Comprehensive experiments in synthetic and real data demonstrate that the proposed method achieves promising results in terms of fiber width estimation and tracking accuracy at moderate amount of noise. Nevertheless, tracking fibers in dense fibrous regions as well as fiber segments with high curvature under noisy conditions may be problematic due to the (linear) geometry of the filter. Therefore, as a future work, we will perform a sensitivity analysis on the filter parameters and subsequently use the extracted free-running Purkinje system in cardiac conduction simulations.

Acknowledgments. The authors would like to thank Yu Ouyang for his help with the manual extraction of the Purkinje fibers and Dr. Volkan Cevher for useful discussions. This work has been funded by the grant NIH HL082729 and by startup funds from the JHU Whiting School of Engineering.

References

1. Kirbas, C., Quek, F.: A review of vessel extraction techniques and algorithms. ACM Computing Surveys **36**(2) (2004) 81–121
2. Palágyi, K., Tschirren, J., Hoffman, E., Sonka, M.: Quantitative analysis of pulmonary airway tree structures. Computers in Biology and Medicine **36** (2006) 974–996

3. Vadakkumpadan, F., Rantner, L., Tice, B., Boyle, P., Prassl, A., Vigmond, E., Plank, G., Trayanova, N.: Image-based models of cardiac structure with applications in arrhythmia and defibrillation studies. Journal of Electrocardiology **42**(2) (2009) 157.e1–157.e10
4. Qian, X., Brennan, M., Dione, D., Dobrucki, W., Jackowski, M., Breuer, C., Sinusas, A., Papademetris, X.: A non-parametric vessel detection method for complex vascular structures. Medical Image Analysis **13** (2009) 49–61
5. Mori, S., van Zijl, P.: Fiber tracking: Principles and strategies - A technical review. NMR in Biomedicine **15** (2002) 468–480
6. Florin, C., Paragios, N., Williams, J.: Globally optimal active contours, sequential Monte Carlo and on-line learning for vessel segmentation. In: European Conference on Computer Vision. LNCS 3953 (2006) 476–489
7. Shim, H., Kwon, D., Yun, I., Lee, S.: Robust segmentation of cerebral arterial segments by a sequential Monte Carlo method: Particle filtering. Computer Methods and Programs in Biomedicine **84** (2006) 135–145
8. Schaap, M., Smal, I., Metz, C., van Walsum, T., Niessen, W.: Bayesian tracking of elongated structures in 3D images. In: Information Processing in Medical Imaging. LNCS 4584 (2007) 74–85
9. Lesage, D., Angelini, E., Bloch, I., Funka-Lea, G.: Medial-based Bayesian tracking for vascular segmentation: Application to coronary arteries in 3D CT angiography. In: IEEE International Symposium on Biomedical Imaging. (2008) 268–271
10. Allen, K., Yau, C., Noble, J.: A recursive, stochastic vessel segmentation framework that robustly handles bifurcations. In: Conference on Medical Image Understanding and Analysis. (2008)
11. Lazar, M., Alexander, A.: Bootstrap white matter tractography (BOOT-TRAC). NeuroImage **24**(2) (2005) 524–532
12. Friman, O., Farnebäck, G., Westin, C.F.: A Bayesian approach for stochastic white matter tractography. IEEE Trans. on Medical Imaging **25**(8) (2006) 965–978
13. Jbabdi, S., Woolrich, M., Andersson, J., Behrens, T.: A Bayesian framework for global tractography. NeuroImage **37**(1) (2007) 116–129
14. Zhang, F., Hancock, E., Goodlett, C., Gerig, G.: Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. Medical Image Analysis **13** (2009) 5–18
15. Descoteaux, M., Deriche, R., Knoesche, T., Anwander, A.: Deterministic and probabilistic tractography based on complex fiber orientation distributions. IEEE Trans. on Medical Imaging **28**(2) (2009) 269–286
16. Doucet, A., de Freitas, N., Gordon, N., eds.: Sequential Monte Carlo Methods in Practice. Springer-Verlag (2001)
17. Arulampalam, S., Maskell, S., Gordon, N., Clapp, T.: A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking. IEEE Trans. on Signal Processing **50**(2) (2002) 174–188
18. Geman, D., Jedynak, B.: An active testing model for tracking roads in satellite images. IEEE Trans. on Pattern Analysis and Machine Intelligence **18** (1996) 1–14
19. Çetingül, H., Plank, G., Trayanova, N., Vidal, R.: Estimation of multimodal orientation distribution functions from cardiac MRI for tracking Purkinje fibers through branchings. In: IEEE Int. Symposium on Biomedical Imaging. (2009)
20. Mardia, K., Jupp, P.: Directional Statistics. Wiley Series in Probability and Statistics (2000)
21. Subbarao, R., Meer, P.: Nonlinear mean shift over Riemannian manifolds. International Journal of Computer Vision **84** (2009) 1–20