

Rotation Invariant Completion Fields for Mapping Diffusion MRI Connectivity

Parya MomayyezSiahkal and Kaleem Siddiqi

School of Computer Science & Centre for Intelligent Machines, McGill University

Abstract. Partial differential equations have been successfully used for fibre tractography and for mapping connectivity indices in the brain. However, the current implementation of methods which require 3D orientation to be tracked can suffer from serious shortcomings when invariance to 3D rotation is desired. In this paper we focus on the 3D stochastic completion field and introduce a new methodology to solve the underlying PDE in a manner that achieves rotation invariance. The key idea is to use spherical harmonics to solve the Fokker-Planck equation representing the evolution of the probability density function of a 3D directional random walk. We validate the new approach by presenting improved connectivity indices on synthetic data, on the MICCAI 2009 Fibre Cup phantom and on a biological phantom comprised of two rat spinal chords in a crossing configuration.

Keywords: Diffusion MRI, Completion fields, Connectivity, Spherical harmonics, Fokker-Planck equation

1 Introduction

Partial differential equations (PDEs) have been used extensively for white matter tractography in diffusion magnetic resonance imaging (Diffusion MRI) [1], [2], [3], [4]. These methods offer the advantage that a theoretically well-grounded measure of connectivity between two regions can be found via a direct computation of a connectivity map, represented by the steady state solution to the PDE. Additionally, the minimum cost pathways can be easily found via a post-processing step. A further advantage is that they can incorporate relevant information from all directions, not just the maxima of the fibre orientation distributions (FODs), into the framework.

Among the PDE-based tractography/connectivity methods, some try to find the steady state solution of variations of the diffusion equation in an anisotropic medium. O'Donnell et al. [1] work directly with the diffusion equation in an anisotropic medium, using 3-dimensional tensors as conductivity tensors. The degree of connectivity is approximated by the steady state flow along any path. Batchelor et al. [2] modify this algorithm by adding a convection term to the PDE, integrating a measure of anisotropy into the concentration flow. In [3], Hageman et al. extend the probabilistic first-order PDE methods using a model based on fluid-mechanics. A fluid flow is simulated through a diffusion tensor

field such that local anisotropy values are encompassed in the PDE via a viscosity term. In a related approach, Fletcher et al. propose a front evolution scheme to compute the minimal cost of connecting two regions of interest [5]. Pichon et al. [6] have generalized PDE methods to the case of high angular resolution diffusion imaging (HARDI) by solving the Hamilton-Jacobi-Bellman equation for minimum cost curves. Motivated by Williams and Jacobs' 2D completion field model in computer vision [7], Momayyez et al. [4] propose a somewhat different computational model, where the probability density of water molecules in a 3D directional random walk is described by the underlying Fokker-Planck equation. Their approach distinguishes itself from the other PDE-based methods by developing a computational model that is fundamentally linked to the underlying physical process, i.e., the diffusion of water molecules in an anisotropic medium.

Although the algorithm suggested in [4] is numerically stable and offers a convenient way to incorporate local diffusion data by setting diffusion and decay parameters based on the FOD computed at each voxel, the uniform sampling of θ and ϕ (the colatitude and azimuth in spherical coordinates) leads to rotation variance. Rotation variance implies that the computed connectivity indices depend on the orientation of the diffusion weighted image (DWI) in a given global coordinate system, which is clearly unacceptable. The underlying uniform sampling of θ and ϕ introduces singularities at $\theta = 0, \pi$ and also contributes to a nonuniform sampling of the spherical surface such that the density of samples on the sphere varies with the distance from the poles. One might be tempted to simply use a near-uniform sampling scheme on a spherical shell, but this complicates the numerical method since the notion of the nearest neighbour in θ and ϕ coordinates, which plays a role in the discretized PDE, is no longer well-defined.

In the present article we introduce a new methodology for achieving rotation invariance in the solution of the PDE underlying the stochastic completion field algorithm. The key idea is to use a spherical harmonics (SPH) basis to provide a rotation invariant numerical estimation of the evolution of the probability density function associated with a 3D directional random walk. While the development in this paper is tailored to a specific PDE, the basic idea of solving PDEs on domains with a spherical topology has been of significant interest in the fields of astronomy, geophysics, and nuclear physics [8] and can find applications in a wide variety of medical imaging and computer vision applications where 3D orientation data is involved. We show that the proposed framework can be formulated in a manner that is independent of the sampling scheme used for ϕ and θ . Therefore, using a near-uniform sampling on the spherical shell together with the SPH-based algorithm achieves the desired rotation invariance. A second contribution of this article is the modification of the 3D directional random walk such that the Fokker-Planck equation has two additional drift terms, which improve the coherence of the simulated trajectories by better exploiting local fiber orientations. We validate this new approach by comparing the obtained connectivity indices with those computed by the method of [4] on synthetic data. Further results are presented on a rat spinal cord phantom [9] and on the MICCAI 2009 Fibre Cup phantom [10].

2 Background: 3D Stochastic Completion Fields

In the 3D stochastic completion field model introduced in [4] a 3D directional random walk is used to model the Brownian motion of water molecules. The state of a particle (x, y, z, θ, ϕ) in $\mathbb{R}^3 \times S^2$ is updated according to the following set of differential equations:

$$\begin{aligned} dx &= \sin \theta \cos \phi dt; & dy &= \sin \theta \sin \phi dt; & dz &= \cos \theta dt \\ d\theta &= \sigma_\theta dB(\theta); & d\phi &= \sigma_\phi dB(\phi); & dB(\theta), dB(\phi) &\sim N(0, dt). \end{aligned} \quad (1)$$

Under this model particles tend to travel in a straight line along the direction of their current orientation followed by a slight change in the orientation at each step. The deviations in orientation are controlled by the Brownian motion terms $B(\phi)$ and $B(\theta)$ and are proportional to the diffusion parameters σ_ϕ^2 and σ_θ^2 . Additionally, an average lifetime ζ is associated with each particle, thus favouring shorter paths over longer ones. Under this model, the probability of passing through a particular state while bridging the gap between a source and a sink region is given by the product of a *stochastic source field* and a *stochastic sink field* [7], [4]. Furthermore, the Fokker-Planck equation describing the time evolution of the probability density of a particle in the course of its 3D directional random walk [11] is given as:

$$\frac{\partial P}{\partial t} = -\sin \theta \cos \phi \frac{\partial P}{\partial x} - \sin \theta \sin \phi \frac{\partial P}{\partial y} - \cos \theta \frac{\partial P}{\partial z} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} - \frac{1}{\zeta} P. \quad (2)$$

The results presented in [4] demonstrate the ability of a Lax-Wendroff scheme for discretizing the advection terms coupled with a Crank-Nicholson scheme for the diffusion terms to achieve numerically stable connectivity measures between two seed regions. However, this implementation suffers from the serious limitation of having singularities at the poles due to the uniform sampling of ϕ and θ , and of being rotationally variant due to the subsequent nonuniform density of spherical samples. In the following we overcome this limitation by using a SPH formulation of the PDE in the spherical domain to achieve rotation invariance, which is the main contribution of this article. This strategy can be applied in other medical imaging settings where 3D orientations have to be tracked and updated in a PDE.

We also suggest two modifications of the above model to yield improved performance. The first one has to do with the use of the Lax-Wendroff scheme for the advective part of Equation 2. Using the Lax method, e.g., for $\frac{\partial P}{\partial t} = -\sin \theta \cos \phi \frac{\partial P}{\partial x}$, a change in the probability value of a voxel (x, y, z) at time t is propagated to voxels $(x+1, y, z)$ and $(x-1, y, z)$ at time $t+1$. In fact, if $\sin \theta \cos \phi > 0$, then only the value of the voxel $(x+1, y, z)$ must be updated. A similar effect occurs for negative values of $\sin \theta \cos \phi$. In the current implementation, we use an upwind differencing method for the advection terms to resolve this issue. The second modification is the addition of two extra angular advection steps to the algorithm, as described in the following section, to exploit the local fibre orientations at each voxel more carefully. This in turn reduces the penalty for highly curved tracts, when the underlying data supports such pathways.

3 Angular Constraints on 3D Directional Random Walk

In a directional random walk, as described in [4], the stochastic component of the particle's motion is constrained to the orientation change which is entirely Brownian. Adhering to this original definition, the probability of aligning the particle's orientation with an FOD maximum is inversely related to the exponential of the angular difference between the incoming direction and the maxima. Due to this characteristic a lower probability value will be associated with curved fibre tracts compared to straight ones. While this can be a desirable property in situations where one attempts to bridge the gap between two edge segments with no other informative features in between, one would want to remove this bias when the underlying data supports larger orientation changes.

To resolve this issue and to remove the bias towards straight fibre tracts, we propose a variation of the 3D directional random walk where a particle's state is updated according to the following steps:

1. The particle moves in 3D in the direction of its current orientation:

$$dx = \sin \theta \cos \phi dt; \quad dy = \sin \theta \sin \phi dt; \quad dz = \cos \theta dt. \quad (3)$$

2. The particle's orientation is changed according to two deviations in the oscillating and the binormal planes incorporating the stochastic motion of the random walk. Unlike the original directional random walk, a drift term is added to the orientation changes $d\theta$ and $d\phi$ at each step proportional to μ_θ and μ_ϕ which are set based on the angular difference between the current orientation and the FOD maximum. The diffusion terms σ_ϕ^2 and σ_θ^2 , on the other hand, govern the amount of diffusion allowed at each step contributing to the Brownian component of the directional random walk as before.

$$d\theta = \mu_\theta dt + \sigma_\theta dB(\theta); \quad d\phi = \mu_\phi dt + \sigma_\phi dB(\phi); \quad dB(\theta), dB(\phi) \sim N(0, dt). \quad (4)$$

The stochastic differential equations given in Eq. 4 push the incoming direction towards the most likely orientation at each voxel (when available) through the associated nonzero drift terms. This encourages the particles to follow the closest FOD maxima in the course of their directional random walk, while the nonzero diffusion coefficients allow the particles to follow other directions as well to compensate for noise, data and other model-dependent inaccuracies.

Modified Fokker-Planck Equation The introduction of the drift terms into the orientation changes at each step leads to the modification of the underlying Fokker-Planck equation by the addition of extra angular advection terms to the PDE:

$$\begin{aligned} \frac{\partial P}{\partial t} = & -\sin \theta \cos \phi \frac{\partial P}{\partial x} - \sin \theta \sin \phi \frac{\partial P}{\partial y} - \cos \theta \frac{\partial P}{\partial z} \\ & - \mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} - \frac{1}{\zeta} P. \end{aligned} \quad (5)$$

In the rest of this paper, the Fokker-Planck equation will refer to Equation 5.

4 Rotation Invariant Computation

In [4], the solution to the diffusion terms of Eq. 5 is computed using the implicit Crank-Nicholson numerical scheme. The numerical estimation is based upon uniform sampling of ϕ and θ on the spherical surface, which suffers from the inherent problem of rotation variance. Rotation invariant computation of the Fokker-Planck equation requires a uniform sampling of the spherical surface, but this makes the application of finite differencing methods in the spherical domain challenging. Whereas there is no known point set on a spherical surface which can be realized as the analog of uniform sampling in Euclidean space, our goal is to minimize the error associated with the global properties of the sampling points. To accomplish this, we use a quasi-uniform point set obtained by an electrostatic repulsion approach on the sphere, known as the minimum energy point distribution (ME) [9]. A numerical solution based on the SPH expansion of the probability function on the spherical shell is consequently devised. This SPH-based formulation is not linked to any particular sampling scheme, providing the appropriate framework for employing a quasi-uniform spherical point set.

4.1 Spherical Harmonics Based Formulation

Spherical harmonics (SPH), normally denoted by Y_l^m , are the spherical analog of the Fourier transform basis defined for complex functions on the unit sphere. These functions form a complete orthonormal system of the space of square integrable functions on the sphere $L^2(\mathbf{S}^2)$ and are given as:

$$Y_l^m(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_l^m(\cos \theta) \exp(im\phi), \quad (6)$$

where P_l^m is an associated Legendre polynomial¹. SPHs provide a natural framework to work with functions living on the sphere and as such, they have been used widely in different application areas dealing with data on sphere. Many algorithms have also been introduced in the diffusion-MRI community which have adapted spherical harmonics basis to model the diffusion signal [12], [13].

Forming an orthonormal basis on the spherical shell, any spherical function can be written as $P(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l c_l^m Y_l^m(\theta, \phi)$ where the expansion is usually truncated at some order L . Given N data points on the sphere $\mathbf{P} = \{P_1, \dots, P_N\}$ and choosing an approximation order of L , a linear least-squares scheme can be used to solve for the unknown SPH coefficients $\mathbf{c} = \{c_0^0, c_1^{-1}, c_1^0, c_1^1, \dots, c_L^L\}$ [14]. The linear least-squares approach seeks to solve the

¹ In this work, we use the real form of the spherical harmonics since the function to be expanded is a real-valued probability density function.

following linear system of equations using the pseudo-inverse of the SPH matrix:

$$\begin{bmatrix} Y_0^0(\theta_1, \phi_1) & Y_1^{-1}(\theta_1, \phi_1) & \dots & Y_L^L(\theta_1, \phi_1) \\ \vdots & \vdots & \ddots & \vdots \\ Y_0^0(\theta_N, \phi_N) & Y_1^{-1}(\theta_N, \phi_N) & \dots & Y_L^L(\theta_N, \phi_N) \end{bmatrix} \times \begin{bmatrix} c_0^0 \\ \vdots \\ c_L^L \end{bmatrix} = \begin{bmatrix} P_1 \\ \vdots \\ P_N \end{bmatrix}, \quad (7)$$

$$\mathbf{Y}\mathbf{c} = \mathbf{P}, \quad \mathbf{c} = \mathbf{Y}^+\mathbf{P} = (\mathbf{Y}^T\mathbf{Y})^{-1}\mathbf{Y}^T\mathbf{P}.$$

The SPH expansion together with the above equation is next used to develop an algorithm to solve the spherical part of the PDE given in Eq. 5:

$$\begin{aligned} P(t, \theta, \phi) &= \sum_{l=0}^L \sum_{m=-l}^l c_l^m(t) Y_l^m(\theta, \phi) \\ \frac{\partial P}{\partial t} &= -\mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} \\ &= \sum_{l=0}^L \sum_{m=-l}^l c_l^m(t) \left(-\mu_\phi(\theta, \phi) \frac{\partial Y_l^m(\theta, \phi)}{\partial \phi} - \mu_\theta(\theta, \phi) \frac{\partial Y_l^m(\theta, \phi)}{\partial \theta} \right. \\ &\quad \left. + \frac{\sigma_\phi^2(\theta, \phi)}{2} \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \phi^2} + \frac{\sigma_\theta^2(\theta, \phi)}{2} \frac{\partial^2 Y_l^m(\theta, \phi)}{\partial \theta^2} \right). \end{aligned} \quad (8)$$

Note that in Eq. 8 the x -, y -, z - dependence of P is implicit since the spherical PDE is solved locally for each voxel independent of the neighbouring voxels. The partial derivatives of Y_l^m with respect to θ and ϕ at a point can be expressed as the combination of SPHs of different phase factors m (but the same order l) at the given point. Using this property, Eq. 8 can be written simultaneously for all the sample points in matrix form leading to a linear system of ordinary differential equations (ODEs):

$$\begin{aligned} \frac{\partial \mathbf{P}}{\partial t} &= (-\mathbf{M}_\phi \mathbf{Y}_\phi - \mathbf{M}_\theta \mathbf{Y}_\theta + \Sigma_\phi \mathbf{Y}_{\phi^2} + \Sigma_\theta \mathbf{Y}_{\theta^2}) \mathbf{c} \\ &= (-\mathbf{M}_\phi \mathbf{Y}_\phi - \mathbf{M}_\theta \mathbf{Y}_\theta + \Sigma_\phi \mathbf{Y}_{\phi^2} + \Sigma_\theta \mathbf{Y}_{\theta^2}) \mathbf{Y}^+ \mathbf{P} = \mathbf{D}\mathbf{P}, \end{aligned} \quad (9)$$

where:

$$\begin{aligned} \mathbf{M}_\phi &= \mathbf{diag}[\mu_\phi(\theta_1, \phi_1), \dots, \mu_\phi(\theta_N, \phi_N)] & \mathbf{M}_\theta &= \mathbf{diag}[\mu_\theta(\theta_1, \phi_1), \dots, \mu_\theta(\theta_N, \phi_N)], \\ \Sigma_\phi &= \mathbf{diag}[\sigma_\phi(\theta_1, \phi_1), \dots, \sigma_\phi(\theta_N, \phi_N)] & \Sigma_\theta &= \mathbf{diag}[\sigma_\theta(\theta_1, \phi_1), \dots, \sigma_\theta(\theta_N, \phi_N)], \\ \mathbf{Y}_\phi &= \begin{bmatrix} \frac{\partial Y_0^0(\theta_1, \phi_1)}{\partial \phi} & \dots & \frac{\partial Y_L^L(\theta_1, \phi_1)}{\partial \phi} \\ \vdots & \ddots & \vdots \\ \frac{\partial Y_0^0(\theta_N, \phi_N)}{\partial \phi} & \dots & \frac{\partial Y_L^L(\theta_N, \phi_N)}{\partial \phi} \end{bmatrix} & \mathbf{Y}_\theta &= \begin{bmatrix} \frac{\partial Y_0^0(\theta_1, \phi_1)}{\partial \theta} & \dots & \frac{\partial Y_L^L(\theta_1, \phi_1)}{\partial \theta} \\ \vdots & \ddots & \vdots \\ \frac{\partial Y_0^0(\theta_N, \phi_N)}{\partial \theta} & \dots & \frac{\partial Y_L^L(\theta_N, \phi_N)}{\partial \theta} \end{bmatrix}, \\ \mathbf{Y}_{\phi^2} &= \begin{bmatrix} \frac{\partial^2 Y_0^0(\theta_1, \phi_1)}{\partial \phi^2} & \dots & \frac{\partial^2 Y_L^L(\theta_1, \phi_1)}{\partial \phi^2} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 Y_0^0(\theta_N, \phi_N)}{\partial \phi^2} & \dots & \frac{\partial^2 Y_L^L(\theta_N, \phi_N)}{\partial \phi^2} \end{bmatrix} & \mathbf{Y}_{\theta^2} &= \begin{bmatrix} \frac{\partial^2 Y_0^0(\theta_1, \phi_1)}{\partial \theta^2} & \dots & \frac{\partial^2 Y_L^L(\theta_1, \phi_1)}{\partial \theta^2} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 Y_0^0(\theta_N, \phi_N)}{\partial \theta^2} & \dots & \frac{\partial^2 Y_L^L(\theta_N, \phi_N)}{\partial \theta^2} \end{bmatrix}. \end{aligned} \quad (10)$$

It is important to emphasize that the spherical harmonic formulation still allows for local specification of the angular diffusion $\sigma_\phi^2, \sigma_\theta^2$ and advection μ_ϕ, μ_θ coefficients. These parameters are selected locally based on the diffusion-MRI data available throughout the volume. In order to avoid too large or too small time steps, an adaptive step size numerical algorithm, the Runge-Kutta-Fehlberg (RKF45) method, is used to advance the solution of Eq. 9 in time.

5 The Complete Computational Model

A summary of the overall computational method is provided in the following set of equations where, similar to [4], a time splitting approach has been employed. As mentioned in Section 1, we replace the Lax-Wendroff scheme in [4] by an upwind differencing scheme to solve the PDE along the x -, y -, and z -coordinates. The SPH formulation is also used to solve the PDE in the spherical domain:

$$\begin{aligned}
P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} &= P_{x,y,z,\theta,\phi}^t - \sin \theta \cos \phi \begin{cases} P_{x,y,z,\theta,\phi}^t - P_{x-1,y,z,\theta,\phi}^t & \text{if } \sin \theta \cos \phi > 0 \\ P_{x+1,y,z,\theta,\phi}^t - P_{x,y,z,\theta,\phi}^t & \text{if } \sin \theta \cos \phi < 0 \end{cases}, \\
P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} &= P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} - \sin \theta \sin \phi \begin{cases} P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} - P_{x,y-1,z,\theta,\phi}^{t+\frac{1}{5}} & \text{if } \sin \theta \sin \phi > 0 \\ P_{x,y+1,z,\theta,\phi}^{t+\frac{1}{5}} - P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} & \text{if } \sin \theta \sin \phi < 0 \end{cases}, \\
P_{x,y,z,\theta,\phi}^{t+\frac{3}{5}} &= P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} - \cos \theta \begin{cases} P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} - P_{x,y,z-1,\theta,\phi}^{t+\frac{2}{5}} & \text{if } \cos \theta > 0 \\ P_{x,y,z+1,\theta,\phi}^{t+\frac{2}{5}} - P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} & \text{if } \cos \theta < 0 \end{cases}, \\
\mathbf{P}_{x,y,z,\theta,\phi}^{t+\frac{4}{5}} &= \mathbf{D} \mathbf{P}_{x,y,z,\theta,\phi}^{t+\frac{3}{5}}, \\
P_{x,y,z,\theta,\phi}^{t+1} &= e^{-\frac{1}{\zeta}} P_{x,y,z,\theta,\phi}^{t+\frac{4}{5}}.
\end{aligned} \tag{11}$$

Both the source field and the sink field are computed by time integration of the above probability density function representing a particle's state in the course of its directional random walk:

$$P'(x, y, z, \theta, \phi) = \int_0^\infty P^t(x, y, z, \theta, \phi) dt \approx \sum_0^{t'} P^t(x, y, z, \theta, \phi). \tag{12}$$

The angular advection and diffusion coefficients $\mu_\theta, \mu_\phi, \sigma_\theta, \sigma_\phi$ and the life time coefficient ζ are set locally for each state using the fibre orientation distribution (FOD) calculated at each voxel. For all orientations within a certain angular difference from the FOD maxima, μ_θ and μ_ϕ are equal to the angular difference between the closest FOD maximum and the current orientation. σ_θ and σ_ϕ are set to some fixed value to take the effect of the underlying noise and inaccuracies into account. Finally, the particle's life time ζ for each state is also set based on the angular difference, i.e., ζ is chosen to be very small for orientations far from the FOD maxima, and it diminishes linearly with increasing angular difference for the rest of the orientations.

6 Experimental Results

In this section, we evaluate the performance of our algorithm by running a series of validation experiments. Qualitative and quantitative validation of our method is provided using synthetic multi-tensor data, a biological phantom and the MICCAI 2009 Fibre Cup phantom made of small-diameter acrylic fibres. We also provide comparative results between the connectivity measures provided by the 3D completion field algorithm and a standard probabilistic fibre tractography algorithm, similar to that of [15].

6.1 Validation on Synthetic Data

To investigate the rotation invariant property of our new computational model, we generate synthetic data using the multi-tensor model of [14]. For our experiments, a quasi-uniform ME point set of 100 directions on the hemisphere is chosen for diffusion encoding directions. The synthetic data is created by placing diffusion tensors (DTs) along the fibre pathways where the major eigenvector of a DT aligns with the tangent vector of the curves. Partial volume averaging is also considered for voxels with more than one curve passing through them. Background voxels are filled with isotropic tensors with eigenvalues equal to $600 * 10^{-6} \frac{mm^2}{s}$. The eigenvalues for the anisotropic voxels are chosen based on the desired fractional anisotropy. Finally, we simulate the effect of noise by adding complex Gaussian noise to the raw signal.




Fig. 1. Visualizing the effect of rotation on the performance of the SPH-based algorithm and the numerical method of [4]

Figure 1 shows the probability distributions obtained for a synthetic curved tract in two perpendicular spatial orientations, using the numerical method given in [4] and the new rotation invariant method². It is clear that the SPH formula-

² For the purpose of visualization, all results are projected onto the 2D plane. The reader is also encouraged to zoom-in on the figures in the electronic version for a better view of the orientations and the transparency values.

tion recovers the underlying connection more accurately. More significant though is the difference between the probability distributions obtained by the algorithm of [4] for the two spatial orientations. The algorithm even fails to completely recover the connectivity pattern for one of the spatial orientations (Fig. 1(a) left). Conversely, the probability distributions obtained by the SPH-based algorithm clearly demonstrate the robustness of the algorithm to rotation in space. The same conclusion can be drawn from the connectivity measures summarized in Table 1. The small difference still observed for the rotation invariant method is due to the errors introduced by the time splitting approach, which updates the probability distribution function for x -, y - and z - coordinates in separate steps. The effect is more severe for our synthetically generated data with clear jumps from voxel to voxel. We expect this effect to be milder for less coarse real data. The connectivity measures, as described in [4], are computed by taking the average of the probability values of the voxels contained in the source and sink region pairs.

Table 1. Connectivity indices obtained for two spatial orientations of the same synthetic fibre tract

Algorithm	1st Configuration	2nd Configuration
numerical method of [4]	0.0332511	0.271475
SPH-based Algorithm	0.154295	0.197969

6.2 Validation on Phantom Data

Fiber Cup Phantom Data The new SPH-based algorithm was also tested on the MICCAI 2009 Fibre cup phantom [10], with known ground truth. For our experiments, we used the DWIs with a spatial resolution of $3 \times 3 \times 3 \text{ mm}^3$ and a b -value of 1500 s/mm^2 .

Figure 2 provides a probabilistic view of the connectivity patterns obtained for different sink and source region pairs, using our SPH-based algorithm. We have used the same pairs of sink and source regions as in [4], where the first four are associated with true fibre tracts while the last two are not. The connectivity measures provided by the algorithm are also summarized in Table 2 and are compared to the measures obtained by the numerical method of [4] and those of a probabilistic tractography approach similar to [15]. The probabilistic method uses the uncertainty computed from a residual bootstrapping approach to run many iterations of a streamline tractography technique, where the weakest link along the reconstructed fibre tracts represents the connectivity strength from the seed region to every voxel.

Similar to [4], one major observation is the failure of the weakest link measure to provide a consistent connectivity index for true connections. While the algorithm correctly assigns zero connectivity to the connections (e) and (f), no connection is found for the straight connection of (d). Moreover, the connectivity measures for the other three true connections extend over a wide range of values, making any kind of inference based on these measures inaccurate. More

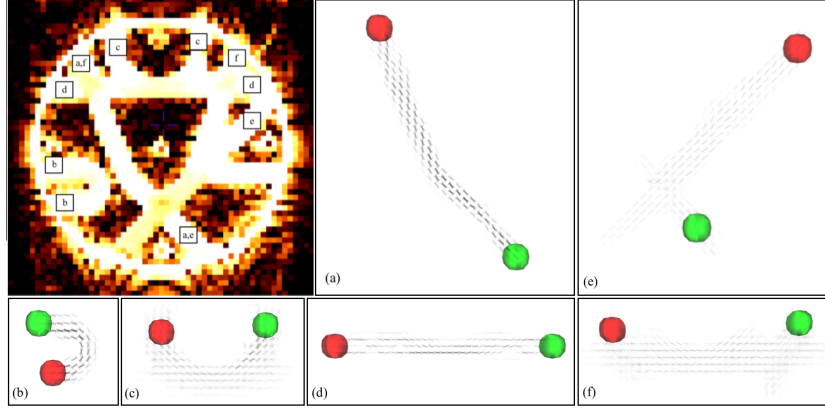


Fig. 2. Results of the SPH-based model. The darker the vectors, the higher the associated probabilities. Top-Left: Mean diffusivity image and the ROIs used

significant though is the improvement gained by using our new SPH-based algorithm over the numerical method of [4]. It should be noted that since this phantom is almost planar, the results obtained mainly show the effect of adding the angular advection terms to the Fokker-Planck equation. While the previous method assigns almost similar measures to the cases (c) and (f), where the first is a true connection and the second is not, the connectivity measure provided by the new algorithm is very low for both cases (e) and (f) and much higher for the other cases which represent true connections. The fact that the connectivity indices of cases (c) and (d) are an order of magnitude lower than those of cases (a) and (b), is mainly due to the complex configuration of fibre tracts at the top part of the phantom, causing the connectivity measure to drop as expected.

Table 2. Connectivity indices obtained for six different seed region pairs

Algorithm	(a)	(b)	(c)	(d)	(e)	(f)
SPH CF	0.327592	0.258567	0.0614615	0.0947138	0.000047583	0.000594305
Num. Mth of [4]	0.17152	0.22445	0.04716	0.24342	0.00551	0.07559
Probabilistic	0.00041	0.239901	0.00792	0	0	0

Biological Phantom Data We also evaluated the performance of our algorithm on a physical phantom with known connectivity built from excised rat spinal cords [9]. The cords were scanned to acquire 90 DWIs as well as 10 T2-weighted images. Unlike the physical phantom used in the experiments of the previous section, the biological phantom is not planar and thus this experiment better explores the 3D performance of our algorithm.

Figure 3 (left) shows the anisotropy image of the phantom and the seed regions used. The table on the right summarizes the connectivity measures obtained between different sink and source regions for our method and the probabilistic tractography algorithm. Based on the ground truth, it is known that ROI (a) is connected to ROI (b) and (c) is connected to (d). The biological phantom has a complex crossing section and low anisotropy values, which make tracking within the phantom a difficult task. However, from the connectivity measures provided by the SPH-based algorithm, it is clear that our algorithm does a better job in distinguishing true connections from false connections when compared to the probabilistic algorithm. The probabilistic tractography approach fails to reconstruct the pathway from ROI (c) to (d). In fact, either starting from ROI (c) or (d), the algorithm cannot pass through the crossing region, providing no connectivity for any of the (cd), (ac) and (bd) ROI pairs. In contrast, the SPH-based completion field algorithm provides us with connectivity measures that are an order of magnitude higher for the true connections.



Algorithm	SPH-Based CF	Probabilistic
ab	0.353533	0.201719
cd	0.174609	0
ac	0.057683	0
bd	0.028613	0

Fig. 3. Anisotropy image and connectivity indices obtained for the biological phantom

7 Conclusion

This paper introduces a mathematical formulation to achieve rotation invariant solutions for partial differential equations where 3D orientation data is involved. The method has been applied to the specific problem of diffusion MRI connectivity measurements using 3D stochastic completion fields. In addition, to make the application of 3D stochastic completion fields more appropriate for brain connectivity measurement in the presence of dense orientation information, additional angular advection terms are incorporated into the 3D directional random walk, leading to a modified Fokker-Planck equation. Using synthetic diffusion MRI data, we have demonstrated that our SPH-based method is more robust to rotation in space. Further experiments on two physical and biological phantoms show that the new formulation outperforms some existing tractography algorithms, a standard probabilistic tractography method and the rotation variant implementation of the stochastic completion field, by providing more reliable connectivity measures across the phantoms and by distinguishing true connections from false connections. In future, we plan to further explore the performance of our algorithm by performing experiments on human brain data and also by setting the

diffusion parameters in a more systematic way, i.e., based on measurements of the underlying noise and data uncertainty.

Acknowledgments. We are grateful to Jennifer Campbell and Bruce Pike for many helpful discussions and for providing the biological phantom data for our experiments. This work was supported by grants from NSERC and FQRNT.

References

1. O'Donnell, L., Haker, S., Westin, C.: New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In: Proceedings of MICCAI'02. (2002) 459–466
2. Batchelor, P., Hill, D., Calamante, F., Atkinson, D.: Study of connectivity in the brain using the full diffusion tensor from MRI. In: Proceedings of IPMI'01. (2001) 121–133
3. Hageman, N., Toga, A., Narr, K., Shattuck, D.: A diffusion tensor imaging tractography algorithm based on Navier-Stokes fluid mechanics. *IEEE Transactions on Medical Imaging* **28** (2009) 348–360
4. MomayyezSiahkal, P., Siddiqi, K.: Probabilistic anatomical connectivity using completion fields. In: Proceedings of MICCAI'10. (2010) 566–573
5. Fletcher, P., Tao, R., Jeong, W., Whitaker, R.: A volumetric approach to quantifying region-to-region white matter connectivity in diffusion tensor MRI. In: Proceedings of IPMI'07. (2007) 346–358
6. Pichon, E., Westin, C., Tannenbaum, A.: A Hamilton-Jacobi-Bellman approach to high angular resolution diffusion tractography. In: Proceedings of MICCAI'05. (2005) 180–187
7. Williams, L., Jacobs, D.: Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation* **9** (1997) 837–858
8. Fornberg, B., Merrill, D.: Comparison of finite difference and pseudo-spectral methods for convective flow over a sphere. *Geoph. Res. Letters* **24** (1997) 3245–3248
9. Campbell, J., Siddiqi, K., Rymar, V., Sadikot, A., Pike, G.: Flow-based fiber tracking with diffusion tensor and Q-ball data: validation and comparison to principal diffusion direction techniques. *NeuroImage* **27** (2005) 725–736
10. Poupon, C., Rieul, B., Kezele, I., Perrin, M., Poupon, F., Mangin, J.: New diffusion phantoms dedicated to the study and validation of high-angular-resolution diffusion imaging (HARDI) models. *Mag. Reson. in Med.* **60** (2008) 1276–1283
11. Mumford, D.: *Elastica and computer vision*. In: *Algebraic Geometry and Its Applications*. Springer-Verlag (1994) 491–506
12. Tournier, J., Yeh, C., Calamante, F., Cho, H., Connelly, A., Lin, P.: Resolving crossing fibres using constrained spherical deconvolution: validation using diffusion-weighted imaging phantom data. *NeuroImage* **41** (2008) 617–625
13. Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R.: Regularized, fast, and robust analytical Q-ball imaging. *Mag. Reson. in Med.* **58** (2007) 497–510
14. Alexander, D., Barker, G., Arridge, S.: Detection and modeling of non-gaussian apparent diffusion coefficient profiles in human brain data. *Mag. Reson. in Med.* **48** (2002) 331–340
15. Berman, J., Chung, S., Mukherjee, P., Hess, C., Han, E., Henry, R.: Probabilistic streamline Q-ball tractography using the residual bootstrap. *NeuroImage* **39** (2008) 215–222