

3D Stochastic Completion Fields for Mapping Connectivity in Diffusion MRI

Parya MomayyezSiahkal and Kaleem Siddiqi, *Senior Member, IEEE*

Abstract—The 2D stochastic completion field algorithm, introduced by Williams and Jacobs [1], [2], uses a directional random walk to model the prior probability of completion curves in the plane. This construct has had a powerful impact in computer vision, where it has been used to compute the shapes of likely completion curves between edge fragments in visual imagery. Motivated by these developments, we extend the algorithm to 3D, using a spherical harmonics basis to achieve a rotation invariant computational solution to the Fokker-Planck equation describing the evolution of the probability density function underlying the model. This provides a principled way to compute 3D completion patterns and to derive connectivity measures for orientation data in 3D, as arises in 3D tracking, motion capture, and medical imaging. We demonstrate the utility of the approach for the particular case of diffusion magnetic resonance imaging, where we derive connectivity maps for synthetic data, on a physical phantom and on an in vivo high angular resolution diffusion image of a human brain.

Index Terms—3D directional random walk, Fokker-Planck equation, completion fields, diffusion MRI, probabilistic connectivity, spherical harmonics

1 INTRODUCTION

CONTOUR completion involves “filling in” the missing portions of partly occluded or fragmented outlines in a visual image. This process is fundamental to computer vision, where it constitutes an essential intermediate step in many applications, including the perceptual grouping of edge maps, image segmentation, and object recognition. Its study also has a rich history in human vision, including considerations by Kanizsa [3], whose work emphasizes the importance of both modal and amodal completions in the perception of visual forms. In fact, the notion of illusory contours, where edges and surfaces that have no physical reality in the scene are nonetheless perceived, has been and continues to be a subject of active research [4]. Mathematical contour completion is an ill-posed problem in that an infinite variety of curves can serve as possible completions while being faithful to the edge fragments in an image. As a result, achieving computational solutions to contour completion is challenging and additional assumptions have to be imposed, such as the minimization of total squared curvature [1], [5], [6] or the minimization of total change in curvature [7].

Among the first attempts to solve the problem of contour completion is Ullman’s approach [5], which uses circular arcs to fill the gap between two boundary segments, searching for those which minimize the total curvature squared and meet smoothly, while being tangent to the

given edges. Parent and Zucker [8] take a differential geometric point of view, using co-circularity in a relaxation labeling framework to approximate tangent vectors and curvature fields to be used for curve completion. Similarly, Guy and Medioni [9] implement a tensor voting approach using a co-circularity kernel. Their kernel generates an extension field starting from a given edge segment by assigning a likelihood of connection to preferred directions at every point in space. The superposition of extension fields from two fragments provides a saliency map describing connectivity between the two edges. Shashua and Ullman propose a saliency network for perceptual grouping which incorporates the notion of minimizing curvature squared plus length [10]. The minimization of total curvature variation has been explored by Kimia et al. in [7], where Euler spirals are used to bridge the gap between edges. The 2D completion algorithm of Williams and Jacobs [1], [2] models the prior probability distribution of completion curves by the trajectory of particles following a 2D directional random walk. The maximum likelihood curves derived from this model correspond to curves of least energy, which, like Shashua and Ullman’s approach, minimize a weighted combination of curvature squared plus length. The derived probabilistic representation of possible connections between two boundary segments is dubbed a *stochastic completion field*. Williams and Jacobs argue that the shape, salience, and sharpness of the contours perceived by humans can be represented by the mode, magnitude, and variance of this completion field. In more recent work, Zweck and Williams [11] achieve invariance to translation and rotation in 2D by computing completion fields using shiftable-twistable functions.

Motivated by Williams and Jacobs’ model, the present paper extends this framework to handle curve completion in 3D. While the extension of the directional random walk is straightforward, the computation of the stochastic completion field in 3D requires special treatment to achieve rotation invariance. We accomplish this by developing the

- The authors are with the School of Computer Science and Centre for Intelligent Machines, McGill University, McConnell Engineering, 3480 University Street, Montreal, QC H3A 2A7, Canada.
E-mail: parya.momayyeziahkal@mail.mcgill.ca, siddiqi@cim.mcgill.ca.

Manuscript received 5 Oct. 2011; revised 14 June 2012; accepted 3 Aug. 2012; published online 28 Aug. 2012.

Recommended for acceptance by P. Golland.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMI-2011-10-0716.

Digital Object Identifier no. 10.1109/TPAMI.2012.184.

associated Fokker-Planck equation describing the evolution of the probability density function of particles undergoing the random walk and by using a spherical harmonics (SPH) basis to achieve a rotation invariant solution on the two spheres. We then tailor the 3D completion field algorithm to include additional drift terms to incorporate spatially dense orientation data when available. Whereas the basic method is applicable to data from motion capture or 3D feature tracking systems, we focus on the particular problem of deriving completions in diffusion magnetic resonance imaging (diffusion MRI) data along with probabilistic measures of connectivity.

This paper is organized as follows: We start with a brief introduction to diffusion MRI and review existing models for curve completion in this context. Section 2 develops the theory of 3D stochastic completion fields. In Section 3, we formulate the Fokker-Planck equation, which describes the evolution in time of the probability density function of particles undergoing a 3D directional random walk. Section 4 develops a modification of the model to include additional drift terms in order to take into account spatially dense measurements, should they be available. In Section 5, we extend the model and its underlying Fokker-Planck equation to also include diffusion on the spatial domain variables. Section 6 develops a rotation invariant numerical method for solving the Fokker-Planck equation by using spherical harmonics to represent functions defined on the 2-sphere. In Section 7, we apply the 3D stochastic completion field algorithm to the problem of mapping connectivity in diffusion MRI data. We explain how parameters for the stochastic completion field are selected based on the underlying diffusion MRI data, and develop a measure to quantify connectivity. Qualitative and quantitative validation of this approach is provided in Sections 7.3.2 and 7.3.3 using synthetic multitensor data and the MICCAI 2009 Fibre Cup phantom [12]. We also carry out experiments on some of the major fiber pathways in the human brain. We conclude with a brief discussion of our results and directions for future work in Section 8.

1.1 Curve Completion for Mapping Brain Connectivity

The emergence of diffusion MRI [13] has led to noninvasive ways to explore connectivity in fibrous tissue. It has been used to assist neuroscientists in tasks such as virtual gross anatomical dissections of the brain [14], analysis of the language network [15], [16], human brain parcellation [17], [18], and joint studies of functional/anatomical brain networks [19], [20]. Measurements obtained from diffusion MRI are sensitized to the anisotropic Brownian motion of water molecules in living tissues, including brain white matter. This anisotropic direction-dependent motion of water molecules is primarily based on the higher diffusivity of molecules along fiber tracts than across them, providing information about the local orientation of white matter fiber bundles [21]. A number of fiber tractography algorithms have been developed to infer the pattern of white matter connectivity by exploiting the local orientation measurements in diffusion MRI [22], [23], [24], [25], [26], [27], [28], [29]. The local fiber orientation is estimated using different models of the diffusion signal, including the diffusion tensor (DT) [30], the orientation distribution function (ODF) [31],

[32], and the fiber orientation density (FOD) [33]. The last two models, which have been developed for high angular resolution diffusion imaging (HARDI) data, are the standard methods used for tractography in the recent literature.

The first group of tractography algorithms proposed are purely deterministic schemes which reconstruct the most likely fiber pathways by tracing along the fiber ODF maxima [23], [34], [35]. To improve upon these methods probabilistic tractography techniques incorporate measurements of noise and model uncertainty [27], [29], [36]. Kalman filtering-based methods [37] treat model estimation and tractography simultaneously by recursively updating the local model parameters, estimating their mean and variance, and indicating the most probable direction of heading. Global tractography methods take a different approach by searching for minimal cost pathways between regions of interest (ROI) [25], [38], [39]. The major advantage of global methods is their ability to trace through regions of local perturbation by optimizing a global criterion. Following this, a connectivity measure can be assigned to the reconstructed fiber pathways based on the employed energy metric. Approaches which try to find the steady-state solution of a partial differential equation (PDE) represent another class of probabilistic algorithms. A direct computation of a connectivity map discriminates these PDE-based algorithms from other tractography methods.

Among the PDE-based tractography techniques, O'Donnell et al. [42] work directly with the diffusion equation in an anisotropic medium, using 3D tensors as conductivity tensors. The degree of connectivity is approximated by the steady-state flow along the maximum flow path. Batchelor et al. [41] modify this algorithm by adding a convection term to the PDE, integrating a measure of anisotropy into the concentration flow. In [44], Hageman et al. extend the probabilistic first-order PDE methods by simulating a fluid flow through a diffusion tensor field. In a related approach, Fletcher et al. propose a front evolution scheme to compute the minimal cost of connecting two regions of interest [45]. Pichon et al. [46] generalized PDE-based methods to HARDI by solving the Hamilton-Jacobi-Bellman equation to extract the minimum cost pathways.

Our 3D stochastic completion field algorithm applied to white matter connectivity is a PDE-based scheme that has certain global properties. In particular, the most likely completion curves under this model map to a notion of 3D curves of least energy. The algorithm is linked to the physical process to which diffusion MRI measurements have been sensitized—the anisotropic Brownian motion of water molecules. One of the main contributions of the proposed approach is that global tractography is performed with a true probabilistic nature, where probability values are assigned to every state in the volume. Furthermore, as discussed in Section 7.2, a measure of connectivity between two regions of interest emerges as an inherent characteristic of the algorithm. This measure has the potential to distinguish real connections from false ones. In the following section, we develop the 3D stochastic completion field model.

2 3D STOCHASTIC COMPLETION FIELD

The 3D stochastic completion field, first used in [47] for the purpose of tractography in diffusion MRI, arises as a direct

curvature, and the orientation changes τ_i in the binormal plane with torsion, the maximum likelihood polygonal paths in (4) in a discrete sense can be seen to minimize a weighted sum of length, curvature squared and torsion squared. With a curve of least energy in a continuous sense minimizing

$$\alpha \int_{\Gamma} \kappa(t)^2 dt + \beta \int_{\Gamma} \tau(t)^2 dt + \gamma \int_{\Gamma} dt, \quad (5)$$

which is a natural extension of elastica in 2D (as in [6]) to 3D, the association between the weights in the discrete and the continuous cases is given by $\alpha = \frac{1}{2\sigma_\phi^2}$, $\beta = \frac{1}{2\sigma_\theta^2}$, and $\gamma = \frac{1}{\zeta} + \log(\sigma_\theta\sqrt{2\pi}) + \log(\sigma_\phi\sqrt{2\pi})$.

3 FOKKER-PLANCK EQUATION FOR THE 3D RANDOM WALK

While Monte-Carlo simulation of the stochastic process formulated in (1) is a straightforward approach for computation of the stochastic source and sink fields, a parallelizable and hence more efficient computational model is provided by exploiting the Fokker-Planck equation of the underlying stochastic process of the 3D directional random walk, as in [2]. The Fokker-Planck equation, formulated as a PDE, describes the time evolution of the probability density function of particles following a 3D directional random walk in space. The general form of the Fokker-Planck equation in N-dimensions is given by Risken [48]:

$$\frac{\partial P}{\partial t} + \sum_{i=1}^N \frac{\partial}{\partial X_i} (D_{(i)}^{(1)}(X, t)P) = \sum_{i=1}^N \sum_{j=1}^N \frac{\partial^2}{\partial X_i \partial X_j} (D_{(ij)}^{(2)}(X, t)P), \quad (6)$$

where $D_{(i)}^{(1)}$ for $i = 1, \dots, N$ are the drift terms and $D_{(ij)}^{(2)}$ are the diffusion coefficients for the $i, j = 1, \dots, N$ variables on which the probability density function P depends. As shown in [47], the drift vector and the diffusion matrix associated with the 3D directional random walk are given by

$$D^{(1)} = \begin{bmatrix} D_{(x)}^{(1)} \\ D_{(y)}^{(1)} \\ D_{(z)}^{(1)} \\ D_{(\theta)}^{(1)} \\ D_{(\phi)}^{(1)} \end{bmatrix} = \begin{bmatrix} \sin \theta \cos \phi \\ \sin \theta \sin \phi \\ \cos \theta \\ 0 \\ 0 \end{bmatrix}, \quad (7)$$

$$D^{(2)} = \begin{bmatrix} D_{(xx)}^{(2)} & D_{(xy)}^{(2)} & D_{(xz)}^{(2)} & D_{(x\phi)}^{(2)} & D_{(x\theta)}^{(2)} \\ D_{(yx)}^{(2)} & D_{(yy)}^{(2)} & D_{(yz)}^{(2)} & D_{(y\phi)}^{(2)} & D_{(y\theta)}^{(2)} \\ D_{(zx)}^{(2)} & D_{(zy)}^{(2)} & D_{(zz)}^{(2)} & D_{(z\phi)}^{(2)} & D_{(z\theta)}^{(2)} \\ D_{(\phi x)}^{(2)} & D_{(\phi y)}^{(2)} & D_{(\phi z)}^{(2)} & D_{(\phi\phi)}^{(2)} & D_{(\phi\theta)}^{(2)} \\ D_{(\theta x)}^{(2)} & D_{(\theta y)}^{(2)} & D_{(\theta z)}^{(2)} & D_{(\theta\phi)}^{(2)} & D_{(\theta\theta)}^{(2)} \end{bmatrix} \quad (8)$$

$$= \mathbf{diag} \left[0, 0, 0, \frac{\sigma_\phi^2}{2}, \frac{\sigma_\theta^2}{2} \right].$$

Replacing the above coefficients in (6) leads to the following partial differential equation:

$$\begin{aligned} \frac{\partial P}{\partial t} = & -\sin \theta \cos \phi \frac{\partial P}{\partial x} - \sin \theta \sin \phi \frac{\partial P}{\partial y} - \cos \theta \frac{\partial P}{\partial z} \\ & + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} - \frac{1}{\zeta} P. \end{aligned} \quad (9)$$

Here, P is the probability that a particle passes through state (x, y, z, θ, ϕ) in the course of its random walk from a seed region. The last term $1/\zeta \times P$ is added to the PDE to incorporate the exponential decay of particles in time. The stochastic source and sink fields can now be computed by time integration of the above PDE. We shall achieve this by introducing a rotation invariant numerical scheme in Section 6.

4 EXTRA DRIFT TERMS IN THE PRESENCE OF SPATIALLY DENSE MEASUREMENTS

In many pattern analysis problems, e.g., those involving 3D tracking, motion capture, or diffusion MRI, orientation data are available at dense spatial locations and this must be taken into account into the computation of completion shapes. In the 3D directional random walk developed in Section 2, the orientation change in a particle's motion is given by stochastic terms which are entirely Brownian. In such a setting, a larger orientation change has a lower probability of occurring. This characteristic leads to lower probability values being associated with states reached through a curved pathway, relative to those reached with a straighter one. While this property can be suitable when bridging the gap between two edge segments with no other informative features in between, it is desirable to remove the bias when the underlying data support larger orientation changes.

To resolve this issue and to take into account spatially dense orientation information, we propose to add two extra angular advection terms to the stochastic process. The magnitude of these angular drift terms should incorporate the available local orientation information. An implementation of these terms is studied to reconstruct diffusion MRI connectivity in Section 7.1. The introduction of these extra terms modifies the motion component of the 3D directional random walk as follows:

- The particle moves in 3D in the direction of its current heading:

$$\begin{aligned} dx &= \sin \theta \cos \phi dt; & dy &= \sin \theta \sin \phi dt; \\ dz &= \cos \theta dt. \end{aligned} \quad (10)$$

- The particle's orientation is then changed by applying two deviations, one in the osculating plane and one in the binormal plane. Unlike the original directional random walk, a drift term is added to the orientation changes $d\theta$ and $d\phi$ whose magnitude is proportional to μ_θ and μ_ϕ . Should a local estimate of orientation be available at a voxel during the propagation of completion fields, that information can be exploited to set μ_θ and μ_ϕ . In the context of diffusion MRI, as described in Section 7.1, these values are chosen for each state based on its angular difference from the available orientation data, which provides a local estimate of the direction of fibers. The angular diffusion terms σ_ϕ^2 and σ_θ^2 , on the other

hand, govern the amount of random angular change allowed at each step, contributing to the Brownian component of the directional random walk:

$$\begin{aligned} d\theta &= \mu_\theta dt + \sigma_\theta dW_t(\theta); & d\phi &= \mu_\phi dt + \sigma_\phi dW_t(\phi); \\ dW_t(\theta), dW_t(\phi) &\sim N(0, dt). \end{aligned} \quad (11)$$

As before, we also allow a certain fraction of particles $(1 - e^{-\tau})$ to decay per unit time. The net effect of the stochastic differential equations given in (11) is to push the incoming direction toward the most likely orientation at each voxel (when available) via the associated nonzero drift terms. This encourages the particles to move along the pathways supported by the already available dense orientation data. At the same time, the angular diffusion terms allow the particles to explore other possible trajectories by moving along other directions. This can account for other contributing random factors (such as Brownian motion of molecules in diffusion MRI) and can compensate for noise and other model-dependent inaccuracies.

4.1 Modified Fokker-Planck Equation

The angular drift terms added above modify the original Fokker-Planck equation as follows:

$$\begin{aligned} \frac{\partial P}{\partial t} &= -\sin\theta \cos\phi \frac{\partial P}{\partial x} - \sin\theta \sin\phi \frac{\partial P}{\partial y} - \cos\theta \frac{\partial P}{\partial z} \\ &- \mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} - \frac{1}{\zeta} P. \end{aligned} \quad (12)$$

5 A MODEL WITH DIFFUSION COMPONENTS ON THE ENTIRE STATE SPACE

It is straightforward to generalize the completion field algorithm to model a 3D directional random walk which also includes diffusion components in the spatial domain. To do so, random perturbations are also added to the change in a particle's position (x, y, z) . The differential equations according to which the particle's state is updated at each time step are then given by

$$\begin{aligned} dx &= \sin\theta \cos\phi dt + \sigma_x dW_t(x), \\ dy &= \sin\theta \sin\phi dt + \sigma_y dW_t(y), \\ dz &= \cos\theta dt + \sigma_z dW_t(z), \\ d\theta &= \mu_\theta dt + \sigma_\theta dW_t(\theta), & d\phi &= \mu_\phi dt + \sigma_\phi dW_t(\phi), \\ dW_t(x), dW_t(y), dW_t(z), dW_t(\theta), dW_t(\phi) &\sim N(0, dt). \end{aligned} \quad (13)$$

The corresponding Fokker-Planck equation describing the evolution of the probability density of a particle is then given by

$$\begin{aligned} \frac{\partial P}{\partial t} &= -\sin\theta \cos\phi \frac{\partial P}{\partial x} - \sin\theta \sin\phi \frac{\partial P}{\partial y} - \cos\theta \frac{\partial P}{\partial z} \\ &+ \frac{\sigma_x^2}{2} \frac{\partial^2 P}{\partial x^2} + \frac{\sigma_y^2}{2} \frac{\partial^2 P}{\partial y^2} + \frac{\sigma_z^2}{2} \frac{\partial^2 P}{\partial z^2} \\ &- \mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} - \frac{1}{\zeta} P. \end{aligned} \quad (14)$$

Current diffusion MRI modeling techniques lead to a single diffusion ODF/FOD maximum for a single fiber, a fanning or a branching structure, making it difficult to distinguish between these three cases. Whereas the spatial diffusion coefficients have been set to zero in the experiments in this paper, the use of nonzero values could be helpful in the recovery of fanning/branching configurations by allowing for particles to propagate not only in the direction of ODF maxima but also in a range of possible directions, if there is evidence that such configurations are present in a region.

6 A ROTATIONALLY INVARIANT NUMERICAL SCHEME

Given a region of interest as the initial source state or the final sink state, one needs to integrate the PDE given by (12) over time to compute the stochastic source or sink fields. We accomplish this using operator splitting, which divides (12) into five PDEs, three of which are in the Cartesian space, one in the spherical domain, and the final one incorporates the decay term. Each of these PDEs can be solved separately using an appropriate stable numerical scheme. Since the advection coefficients for the x , y , and z coordinates are always smaller than 1 ($-\sin\theta \cos\phi$, $-\sin\theta \sin\phi$, $-\cos\theta$), a simple explicit numerical method, such as upwind differencing, suffices. An upwind differencing method avoids transport errors by approximating derivatives based on the direction of the underlying physical process.

The main challenge in the numerical estimation of (12) is to achieve rotation invariance in the solution of the PDE involving the spherical terms. Rotation invariance, in the context of the application to diffusion MRI, implies that a measure of connectivity should not depend on the orientation of the diffusion weighted image (DWI) in a given global coordinate system. Although uniform sampling of the θ and ϕ parameters (the colatitude and azimuth in spherical coordinates) offers the possibility of solving the spherical PDE using a simple numerical scheme, such as the Crank-Nicholson method, such a sampling pattern suffers from the inherent problem of having singularities at the poles. In particular, the sampling of the spherical surface ends up being nonuniform since the sampling density varies with distance from the poles. A rotation invariant computation of (12) requires a uniform sampling of the spherical surface, but this in turn makes the application of the finite differencing methods to the spherical domain a challenge since the notion of the nearest neighbor in the θ and ϕ coordinates is no longer well defined.

In the section that follows, we introduce a new methodology for achieving rotation invariance in the solution of the spherical PDE underlying the stochastic completion field algorithm. The key idea is to use a spherical harmonics basis to represent functions defined on a spherical surface [40]. While the development in this paper is tailored to the specific PDE of (12), 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 [50] and can find applications in other computer

vision or medical imaging applications where 3D orientation data are involved.

Notably, this SPH-based formulation is not linked to any particular sampling scheme. Thus, we choose a quasi-uniform point set obtained by an electrostatic repulsion approach on the sphere, known as the minimum energy point distribution [26]. It is important to note that there is no known point set that achieves the analog of uniform sampling on a spherical surface as in euclidean space. However, our goal is to minimize the error associated with the global properties of the sampling points.

6.1 Spherical Harmonics-Based Numerical Solution

SPHs, normally denoted by Y_l^m , are the spherical analog of the Fourier transform basis defined for complex functions on the unit sphere. They form a complete orthonormal system of the space of square integrable functions on the sphere $L^2(\mathbf{S}^2)$. SPHs have been widely used in different application areas dealing with data on a sphere, including geophysics, astrophysics, quantum mechanics, and imaging. Many algorithms have also been introduced in the diffusion MRI community which have adapted the SPH basis to model the diffusion signal [32], [33].

Using this basis, 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 . In this section, we develop an algorithm which uses SPH expansion of the probability function to solve the spherical part of the PDE given in (12), i.e.:

$$\frac{\partial P}{\partial t} = -\mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2}. \quad (15)$$

We use the real form of the spherical harmonics since the function to be expanded is a real-valued probability density function. The real SPHs are formulated in terms of the complex SPHs as follows:

$$\mathbf{Y}_l^m = \begin{cases} \sqrt{2} \text{Re}(Y_l^m) & -l \leq m < 0 \\ Y_l^0 & m = 0 \\ \sqrt{2} \text{Im}(Y_l^m) & 0 < m \leq l. \end{cases} \quad (16)$$

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 $\mathbf{C} = \{c_0^0, c_1^{-1}, c_1^0, c_1^1, \dots, c_L^L\}$ [51]. The linear least-squares approach seeks to find the SPH coefficients which satisfy the following overdetermined linear system of equations while minimizing the 2-norm of the error term at the sample data points:

$$\underbrace{\begin{bmatrix} \mathbf{Y}_0^0(\theta_1, \phi_1) & \dots & \mathbf{Y}_L^L(\theta_1, \phi_1) \\ \vdots & \ddots & \vdots \\ \mathbf{Y}_0^0(\theta_N, \phi_N) & \dots & \mathbf{Y}_L^L(\theta_N, \phi_N) \end{bmatrix}}_{\mathbf{Y}} \times \underbrace{\begin{bmatrix} c_0^0 \\ \vdots \\ c_L^L \end{bmatrix}}_{\mathbf{C}} = \underbrace{\begin{bmatrix} P_1 \\ \vdots \\ P_N \end{bmatrix}}_{\mathbf{P}}. \quad (17)$$

The mean least-squares error estimation of the SPH coefficients is computed using the pseudo-inverse of the SPH matrix:

$$\mathbf{Y}\mathbf{C} = \mathbf{P}, \quad \mathbf{C} = \mathbf{Y}^+\mathbf{P} = (\mathbf{Y}^T\mathbf{Y})^{-1}\mathbf{Y}^T\mathbf{P}. \quad (18)$$

The SPH expansion together with (18) is next used to rewrite (15) in terms of SPHs:

$$\begin{aligned} \frac{\partial P}{\partial t} &= -\mu_\phi \frac{\partial P}{\partial \phi} - \mu_\theta \frac{\partial P}{\partial \theta} + \frac{\sigma_\phi^2}{2} \frac{\partial^2 P}{\partial \phi^2} + \frac{\sigma_\theta^2}{2} \frac{\partial^2 P}{\partial \theta^2} \\ \frac{\partial P}{\partial t} &= \sum_{l=0}^L \sum_{m=-l}^l c_l^m(t) \left[-\mu_\phi(\theta, \phi) \frac{\partial \mathbf{Y}_l^m(\theta, \phi)}{\partial \phi} \right. \\ &\quad - \mu_\theta(\theta, \phi) \frac{\partial \mathbf{Y}_l^m(\theta, \phi)}{\partial \theta} + \frac{\sigma_\phi^2(\theta, \phi)}{2} \frac{\partial^2 \mathbf{Y}_l^m(\theta, \phi)}{\partial \phi^2} \\ &\quad \left. + \frac{\sigma_\theta^2(\theta, \phi)}{2} \frac{\partial^2 \mathbf{Y}_l^m(\theta, \phi)}{\partial \theta^2} \right]. \end{aligned} \quad (19)$$

Note that in (19) the x, y, z -dependence of P is implicit since the spherical PDE is solved locally for each voxel, independently of the neighboring voxels. Equation (19) can be written simultaneously for all the sample points in matrix form, leading to a linear system of ordinary differential equations (ODEs):

$$\begin{aligned} \frac{\partial \mathbf{P}}{\partial t} &= (-\mathbf{M}_\phi \mathbf{Y}_\phi - \mathbf{M}_\theta \mathbf{Y}_\theta + \mathbf{\Sigma}_\phi \mathbf{Y}_{\phi^2} + \mathbf{\Sigma}_\theta \mathbf{Y}_{\theta^2}) \mathbf{C} \\ &= \underbrace{(-\mathbf{M}_\phi \mathbf{Y}_\phi - \mathbf{M}_\theta \mathbf{Y}_\theta + \mathbf{\Sigma}_\phi \mathbf{Y}_{\phi^2} + \mathbf{\Sigma}_\theta \mathbf{Y}_{\theta^2})}_{\mathbf{D}} \mathbf{Y}^+ \mathbf{P}, \end{aligned} \quad (20)$$

$$\mathbf{\Sigma}_\phi = \text{diag}[\sigma_\phi(\theta_1, \phi_1), \dots, \sigma_\phi(\theta_N, \phi_N)],$$

$$\mathbf{\Sigma}_\theta = \text{diag}[\sigma_\theta(\theta_1, \phi_1), \dots, \sigma_\theta(\theta_N, \phi_N)],$$

$$\mathbf{M}_\phi = \text{diag}[\mu_\phi(\theta_1, \phi_1), \dots, \mu_\phi(\theta_N, \phi_N)],$$

$$\mathbf{M}_\theta = \text{diag}[\mu_\theta(\theta_1, \phi_1), \dots, \mu_\theta(\theta_N, \phi_N)],$$

where

$$\begin{aligned} \mathbf{Y}_\phi &= \begin{bmatrix} \frac{\partial \mathbf{Y}_0^0(\theta_1, \phi_1)}{\partial \phi} & \dots & \frac{\partial \mathbf{Y}_L^L(\theta_1, \phi_1)}{\partial \phi} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{Y}_0^0(\theta_N, \phi_N)}{\partial \phi} & \dots & \frac{\partial \mathbf{Y}_L^L(\theta_N, \phi_N)}{\partial \phi} \end{bmatrix}, \\ \mathbf{Y}_\theta &= \begin{bmatrix} \frac{\partial \mathbf{Y}_0^0(\theta_1, \phi_1)}{\partial \theta} & \dots & \frac{\partial \mathbf{Y}_L^L(\theta_1, \phi_1)}{\partial \theta} \\ \vdots & \ddots & \vdots \\ \frac{\partial \mathbf{Y}_0^0(\theta_N, \phi_N)}{\partial \theta} & \dots & \frac{\partial \mathbf{Y}_L^L(\theta_N, \phi_N)}{\partial \theta} \end{bmatrix}, \\ \mathbf{Y}_{\phi^2} &= \begin{bmatrix} \frac{\partial^2 \mathbf{Y}_0^0(\theta_1, \phi_1)}{\partial \phi^2} & \dots & \frac{\partial^2 \mathbf{Y}_L^L(\theta_1, \phi_1)}{\partial \phi^2} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \mathbf{Y}_0^0(\theta_N, \phi_N)}{\partial \phi^2} & \dots & \frac{\partial^2 \mathbf{Y}_L^L(\theta_N, \phi_N)}{\partial \phi^2} \end{bmatrix}, \\ \mathbf{Y}_{\theta^2} &= \begin{bmatrix} \frac{\partial^2 \mathbf{Y}_0^0(\theta_1, \phi_1)}{\partial \theta^2} & \dots & \frac{\partial^2 \mathbf{Y}_L^L(\theta_1, \phi_1)}{\partial \theta^2} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \mathbf{Y}_0^0(\theta_N, \phi_N)}{\partial \theta^2} & \dots & \frac{\partial^2 \mathbf{Y}_L^L(\theta_N, \phi_N)}{\partial \theta^2} \end{bmatrix}. \end{aligned} \quad (21)$$

In the above equations, 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 and of the same order l . The obtained linear system of ODEs updates the probability values of all the states at a voxel simultaneously and thus can be applied to any kind of sampling scheme. We use the Runge-Kutta-Fehlberg (RKF45) method, which is an adaptive step size numerical algorithm, to solve (20) in order to avoid too large or too small time steps.

6.2 The Overall Computational Model

The overall computational model is described in (22). As mentioned before, a time splitting approach is employed to solve the Fokker-Planck equation given by (12) in five steps. The three first steps use an upwind-differencing algorithm to solve the PDE in x , y , and z coordinates. The spherical harmonics based formulation of the previous section is implemented in the fourth step to solve the equation in the spherical domain. And the final step involves the multiplication by an exponential to take the effect of particle's lifetime into account:

$$\begin{aligned}
P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} &= P_{x,y,z,\theta,\phi}^t \\
&- \sin \theta \cos \phi \begin{cases} P_{x,y,z,\theta,\phi}^t - P_{x-1,y,z,\theta,\phi}^t & \text{if } \sin \theta \cos \phi > 0 \\ P_{x+1,y,z,\theta,\phi}^t - P_{x,y,z,\theta,\phi}^t & \text{if } \sin \theta \cos \phi < 0, \end{cases} \\
P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} &= P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} \\
&- \sin \theta \sin \phi \begin{cases} P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} - P_{x,y-1,z,\theta,\phi}^{t+\frac{1}{5}} & \text{if } \sin \theta \sin \phi > 0 \\ P_{x,y+1,z,\theta,\phi}^{t+\frac{1}{5}} - P_{x,y,z,\theta,\phi}^{t+\frac{1}{5}} & \text{if } \sin \theta \sin \phi < 0, \end{cases} \\
P_{x,y,z,\theta,\phi}^{t+\frac{3}{5}} &= P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} \\
&- \cos \theta \begin{cases} P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} - P_{x,y,z-1,\theta,\phi}^{t+\frac{2}{5}} & \text{if } \cos \theta > 0 \\ P_{x,y,z+1,\theta,\phi}^{t+\frac{2}{5}} - P_{x,y,z,\theta,\phi}^{t+\frac{2}{5}} & \text{if } \cos \theta < 0, \end{cases} \\
\mathbf{P}_{x,y,z}^{t+\frac{4}{5}} &= \mathbf{DP}_{x,y,z}^{t+\frac{3}{5}}, \\
P_{x,y,z,\theta,\phi}^{t+1} &= e^{-\frac{1}{\zeta}} P_{x,y,z,\theta,\phi}^{t+\frac{4}{5}}.
\end{aligned} \tag{22}$$

Both the source field and the sink field are computed by time integration of the probability density function, computed by (22), over time starting from the given source and sink regions:

$$\begin{aligned}
P'(x, y, z, \theta, \phi) &= \int_0^\infty dt P^t(x, y, z, \theta, \phi) \\
&\approx \sum_0^t P^t(x, y, z, \theta, \phi).
\end{aligned} \tag{23}$$

Examining the set of (22), it is clear that the computation can be performed locally for each state, which offers the possibility of parallelizing the algorithm, resulting in significant computational savings.

7 APPLICATION TO CONNECTIVITY MAPPING IN DIFFUSION MRI

7.1 Local Selection of Parameters

In the application of the 3D stochastic completion field to fiber tractography, we are interested in those curves of least energy between a source and a sink region that are also supported by the underlying diffusion MRI data. One simple way to incorporate local fiber orientation information obtained from diffusion MRI acquisitions into the completion field framework is to employ a Bayesian approach where the stochastic completion field is calculated independently of the local diffusion MRI FODs/ODFs. The completion field obtained forms the prior probability of

passing through any state in the 3D volume in the process of reconstructing the possible trajectories between two regions of interest. Subsequently, the probability density for the plausible fiber pathways is found by multiplication of this prior distribution by the FOD/ODF probability profiles, considered as the conditional probability. The multiplication outputs the posterior probability of the initial completion field given the diffusion MRI data. Whereas theoretically simple, this approach requires empirical adjustments of angular diffusion and advection parameters, and this emerges as a significant challenge without prior knowledge of the fiber pathways being estimated.

The Fokker-Planck equation developed in Section 6 provides a more principled way to incorporate the local diffusion MRI data. As (20) and (22) demonstrate, the proposed numerical scheme affords the possibility of local specification of the angular advection and diffusion parameters $\mu_\phi, \mu_\theta, \sigma_\phi, \sigma_\theta$ and also the lifetime coefficient ζ , for every state. This selection can be done based on the diffusion MRI data available throughout the volume. Empirical tuning of the parameters is no longer required, and this significantly extends the applicability of the stochastic completion field model to fiber tractography and connectivity measurement.

For the purpose of our experiments, we have used the FODs calculated by the spherical deconvolution method of Anderson [52]. Nevertheless, the completion field algorithm can employ any other type of diffusion profile, irrespective of the actual method used to estimate it. The basic idea for local selection of parameters is to align the heading direction of the particles with the FOD's maxima and to then introduce some perturbation around the heading direction to account for noise and other sources of uncertainty. To accomplish this, for all states whose orientation is within a certain angular difference from the FOD maxima, μ_θ and μ_ϕ are set equal to the angular difference between the closest FOD maximum and their respective orientation. In the current setting, σ_θ and σ_ϕ are set to some fixed value to model the effect of the underlying noise and inaccuracies. Our future goal is to set these angular diffusion parameters according to the actual values of uncertainty measured at each voxel. Finally, the particle's lifetime ζ for each state is also chosen based on the angular difference from the closet fiber maxima. That is, the lifetime coefficient takes a maximum value, ζ_{Max} , for those states which are aligned with the available local fiber maxima. A very small lifetime is assigned to those states far from the fiber maxima which is obtained by scaling ζ_{Max} with a very small value. For other states, the lifetime reduces linearly with the angular difference, reaching its maximum (ζ_{Max}) for angular difference of zero. The value of ζ_{Max} is chosen based on the number of iterations required to get from the source to the sink, and vice versa, which means that to propagate along a longer fiber tract, particles are assigned a larger lifetime. Such a selection criterion for ζ_{Max} reduces the length dependency of the output connectivity measures.

The strategy described above allows parameters to be set automatically from the field of FODs. It supports smoothly varying fiber tracts while avoiding an excessive penalty for curved fiber pathways. To compute the value of ζ_{Max} , the

algorithm is run once using a large enough ζ_{Max} . Following this initial run, the estimated number of required iterations is used to set the value of ζ_{Max} in a final second run.

7.2 A Measure of Connectivity

Most of the methods proposed thus far in the literature for quantification of white matter connectivity are based on tractography. Two of the most commonly used connectivity measures reported in the literature are: 1) *voxel counting*, which enumerates the number of times an ROI has been reached from a seed region employing a Monte-Carlo simulation [27], [29], [53], and 2) the *weakest link* approach [24], [26], [55], where the lowest confidence value of all tract segments along a reconstructed streamline specifies the confidence in it. Confidence in each tract segment is computed based on certain criteria defined by the underlying tracking process or the previously mentioned Monte-Carlo approach. While these connectivity measures have been used extensively in the literature, they have limitations and are heuristic in nature. Furthermore, not all the relevant information along the reconstructed fiber tracts is incorporated in a principled fashion. These approaches also suffer from the issue of being length dependent, where reduced connectivity is observed with distance from the seed region.

We suggest a different notion of connectivity, which is given by the average of the probability values assigned to all possible states within the source and sink ROIs by the stochastic completion field. Denoting the final probability density function by P , the connectivity measure is mathematically defined as:

$$\text{Connectivity} = \frac{\sum_{x,y,z \in \text{Src, Sink}} \sum_{\phi} \sum_{\theta} P(x, y, z, \phi, \theta)}{\#\text{states in all voxels of source \& sink ROIs}}. \quad (24)$$

These probability values have been calculated during the evolution of the probability density from the source and the sink regions and thus reflect the information accumulated along the connectivity pattern. While the local nature of the Fokker-Planck-based algorithm provides the opportunity to tune parameters in a data dependent fashion, the length-dependence issue is addressed by setting the maximum particle lifetime ζ_{max} based on the minimum number of iterations required to bridge the gap between the source and sink regions. In the present work, the parameter selection process is performed by exploring the use of fiber ODF maxima and their consistency with the local direction of heading. However, the local parameter settings can be tailored to incorporate other measures such as the underlying fiber fractional volume, fractional anisotropy (FA), or fiber tract radius if available.

7.3 Experimental Validation

In this section, we evaluate the performance of our algorithm by running a series of validation experiments. Qualitative and quantitative validations of our method are provided using synthetic multitensor data and the MIC-CAI 2009 Fibre Cup phantom, which is made of small-diameter acrylic fibers [12]. We also compute connectivity patterns on FOD reconstructions of HARDI measurements acquired in vivo from a human brain. In the series of validation experiments, we show that the computed

connectivity measures are consistent with the known ground truth. We compare the performance of the completion field algorithm to that of a standard probabilistic fiber tractography algorithm based on residual bootstrapping, similar to that of [29]. Comparative results are also provided between the connectivity measures computed by the completion field algorithm and the weakest link approach. Our results show that the stochastic completion field outperforms the other (state of the art) method by providing more consistent measures across the volume. Throughout our experiments, we use a quasi-uniform minimum energy label set of 100 unit directions, distributed isotropically over a hemisphere, as our set of orientations defining a particle's state at each voxel.

In all the experiments performed, the probability values computed by the completion field algorithm over the entire state space are visualized by representing each state, i.e., every location and orientation (x, y, z, ϕ, θ) , by a stick. To further facilitate the visualization of these probability values, they are first multiplied by a scalar and then transformed by a sigmoid-like function due to their wide range. The transparency of every stick is then made inversely proportional to the transformed probability of the corresponding state. As mentioned in the previous section, the connectivity values for different source and sink pairs are calculated as the average probability of all the states within them.

7.3.1 Probabilistic Fiber Tractography Using the Residual Bootstrap

Residual bootstrapping, in the context of high angular diffusion MRI, is a widely used technique for nonparametric characterization of uncertainty in the estimated diffusion profiles and fiber orientations [29], [54]. This method can approximate probability density functions of the true fiber orientations, which can then be used in a probabilistic framework to perform tractography. The potential of such an approach for reconstruction of white matter fiber tracts has been demonstrated through various tractography studies [28], [29], [36], [54]. We have therefore chosen to compare our method against a probabilistic tractography method similar to the residual bootstrap algorithm of Berman [29], which is explained in the following paragraphs.

First, a residual bootstrapping scheme is implemented to generate 1,000 diffusion weighted images. In a residual bootstrap approach, the acquired diffusion weighted signal is fit to an SPH basis, and for each iteration, the residuals from this SPH fit are added at random without replacement to the SPH profile to generate a new DWI. The random placement of the residuals reflects the noise characteristics of the original acquisition [29]. For each DWI, the fiber orientation distributions are computed using spherical deconvolution and are subsequently used to extract the direction of maximal diffusion at each voxel. The uncertainty profiles are then approximated by a Gaussian distribution whose parameters are computed by matching the FOD maxima across iterations to obtain the mean FOD maxima and their associated variances. We also compute the occurrence rate of each FOD maximum, representing the number of times a given direction has been observed. Tractography is performed by running standard streamline

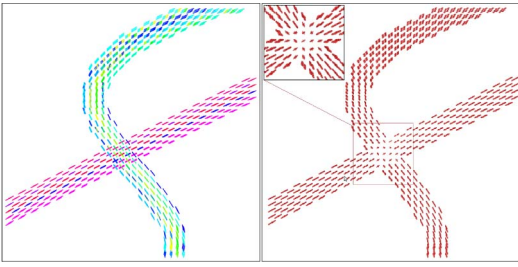


Fig. 2. Left: Crossing fiber configuration used to generate the synthetic data. Right: FOD reconstruction for $FA = 0.84$ and $SNR = 35$.

tracking in a probabilistic setting. Given a user defined seed region, tracking is performed iteratively, starting from the seed ROI, where at each step the direction of propagation is chosen randomly from within the uncertainty profile of the closest FOD maximum.

To measure anatomical connectivity, a confidence value is associated with each tract segment according to the previously computed Gaussian distribution at the corresponding voxel. In a final step, the *weakest link* approach is used to assign a confidence value to each streamline. This confidence value is determined by choosing the support of that segment along the tract that is weakest. The connectivity index for each voxel to the seed region is determined by finding the maximum confidence value over all the streamlines connecting the seed ROI to that voxel. In our experiments, the connectivity measure between a pair of source and sink ROIs is computed using the above approach by running probabilistic tractography starting from both ROIs. The mutual measure of connectivity between the two ROIs is given by taking the average of connectivity indices obtained for sink ROI voxels starting from the source and vice versa.

7.3.2 Experiment on Synthetic Data

In this first series of experiments, we look at the connectivity patterns and indices obtained from running the completion field algorithm on synthetic data. These experiments are designed to evaluate specific features of our algorithm, such as its robustness to noise and curvature.

Synthetic data are generated using Alexander et al.'s multitensor model [51]. For diffusion encoding directions, a quasi-uniform minimum energy point set of 100 directions is chosen, similar to the label set used to define the particle's state in the directional random walk. A crossing fiber configuration is also considered, where a curved pathway crosses a straight one. The synthetic data are created by placing DTs along the fiber pathways where the major eigenvector of tensors aligns with the tangent vector of the curves at a given voxel. Partial volume averaging is also considered for voxels in crossing regions and background voxels are filled with isotropic tensors with eigenvalues equal to $600 * 10^{-6} \frac{\text{mm}^2}{\text{s}}$. 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. To investigate the effect of noise and change in fractional anisotropy, synthetic data were generated for two different values of FA (0.84 and 0.6) and three levels of noise (∞ , 35, and 10). The crossing fiber

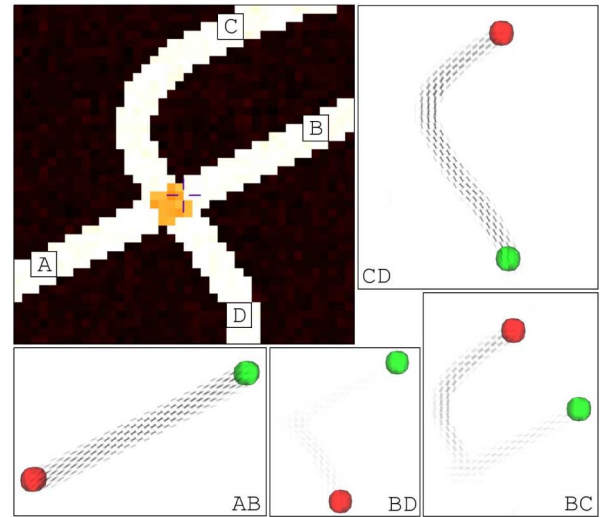


Fig. 3. Steady-state probability distributions obtained by the completion field for the synthetically generated DWI data with $FA = 0.84$ and $SNR = 35$. Top left: The FA image of the synthetic data together with the seed regions used. For visualization, every state is illustrated by a stick, whose transparency is inversely proportional to a sigmoid transformation of that state probability.

configuration and the FODs computed for $FA = 0.84$ and $SNR = 35$ are demonstrated in Fig. 2.

To evaluate the performance of our algorithm, four different ROIs are chosen. The top left image in Fig. 3 shows the chosen ROIs overlaid on the FA image of the synthetic data. The completion field algorithm is run for four pairs of ROIs, i.e., AB, CD, BC, and BD, where the first two are associated with true connections while the latter two are not. As an example, the completion fields obtained for these pairs are illustrated in Fig. 3 for the same FA and SNR as in Fig. 2. One can observe that although our algorithm connects all four ROI pairs, the associated probability values for the BC and BD pairs are much lower than those for the AB and CD connections. The same behavior is observed as one varies the level of noise and FA in the synthetic data.

We have investigated the behavior of the completion field algorithm more thoroughly by looking at how the computed connectivity measures change as the incorporated noise level and the fractional anisotropy value of the synthetic data are modified. Table 1 provides a summary of these connectivity indices, obtained for different ROI pairs with different parameter settings. We observe that the connectivity measures associated with the curved and straight pathways

TABLE 1
Connectivity Indices of Four ROI Pairs
for Different SNR Levels and FA Values

FA	0.84			0.6		
	∞	35	10	∞	35	10
AB	0.2019	0.2019	0.0314	0.2019	0.0754	0.00026
CD	0.1481	0.1459	0.0353	0.1659	0.0704	0.00062
BC	0.0173	0.0171	0.0052	0.0185	0.0082	0.00006
BD	0.0050	0.0043	0.0010	0.0050	0.0017	0.00002

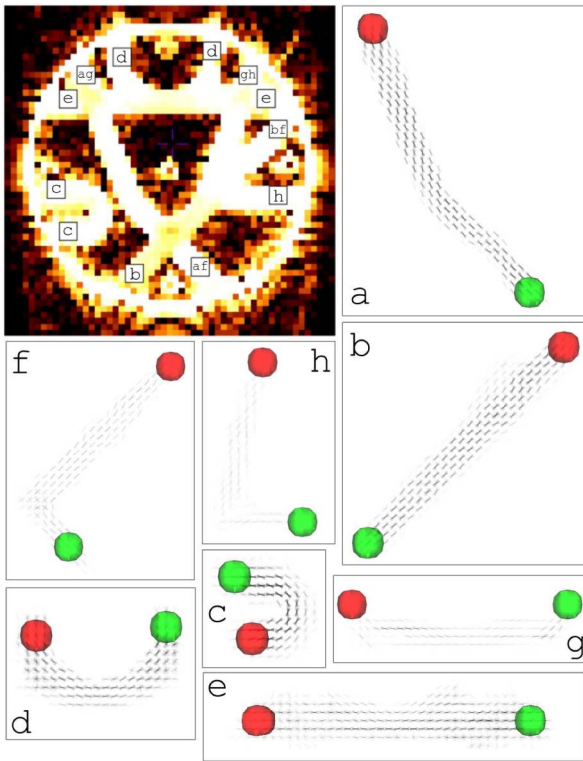


Fig. 4. Visualization of the completion fields obtained for a set of ROI pairs on the physical phantom. Top-Left: Mean diffusivity image and the ROIs used. Every state is illustrated by a stick whose transparency is inversely proportional to a sigmoid transformation of the probability of that state.

are of the same order for the whole range of parameter settings. On the other hand, it can be observed that for high enough SNR ($\text{SNR} = \infty, 35$), that latter being a typical SNR level for in vivo diffusion MRI, the connectivity indices associated with the true connections (AB and CD) are at least an order of magnitude larger than those of the false connections. As the SNR level falls, true and false connections become less distinguishable. Moreover, comparing the values in the last three columns of Table 1 with those given in the first three shows that, for lower values of FA, SNR decrease affects the connectivity measures more considerably since FOD maxima are modified to a greater extent by the added noise. We conclude that as long as the SNR level is within an acceptable range for real data acquisition, the completion field algorithm shows promise in differentiating true connections from false ones.

7.3.3 Experiment on a Physical Phantom

We also test our proposed algorithm on the MICCAI 2009 Fibre cup phantom by Poupon et al. [12], for which the ground truth connections are known. The phantom is made

of hydrophobic acrylic fibers whose diameter is of the same scale as that of myelinated axons. Two diffusion MRI datasets have been provided with different spatial resolutions, of which we have selected the higher resolution image ($3 \times 3 \times 3 \text{ mm}^3$) with an associated b -value of $1,500 \text{ s/mm}^2$.

The steady-state probability densities output by the completion field algorithm for eight different source and sink region pairs are shown in Fig. 4. Based on the ground truth, it is known that the first five pairs (a)-(e) are associated with a true fiber tract, while there is no real tract associated with the latter three (f)-(h). The connectivity measures assigned to each ROI pair are summarized in Table 2. For the purpose of comparison, Table 2 also provides the connectivity indices output by the weakest link approach and the probabilistic tractography algorithm explained in Section 7.3.1.

The connectivity indices presented in Table 2 demonstrate the failure of the probabilistic tractography method to recover all the existing connections and the inability of the weakest link measure to provide consistent connectivity indices across different ROI pairs. The probabilistic bootstrap algorithm did not find any connection between ROI pairs (f), (g), and (h) and consequently the weakest link approach correctly assigns zero connectivity to these cases. However, the tractography fails to recover the nearly straight connectivity pattern between the source and sink ROIs of (b) and (e), which leads to a zero connectivity index for these pairs too. The fact that the connectivity measures obtained are not consistent across true and false connections makes any kind of inference based on these measures inaccurate. It is important to note that this failure is due to deficiencies in both the connectivity measure and the tractography algorithm. Nevertheless, both of these methods are among the state-of-the-art algorithms for tractography and measuring connectivity.

In contrast, the connectivity values computed by the 3D stochastic completion field algorithm are very promising. As can be observed from the first row in Table 2, the completion field algorithm finds a connectivity pattern for all given pairs of ROIs. However, the connectivity values assigned to the (f), (g), and (h) ROI pairs are much lower (at least two orders of magnitude) than those computed for the first five pairs, where the latter represent true connections. Additionally, among the five connectivity measures associated with true connections, those obtained for cases (d) and (e) are an order of magnitude lower than the other three. This is mainly due to the complex configuration of fiber tracts at the top part of the phantom. The computation time for each ROI pair on the phantom, with state space dimension of $64 \times 64 \times 3 \times 200$, was about 15 minutes on average, running on four 2.8 GHz Intel Xeon processors.

TABLE 2
Connectivity Indices Assigned to Eight Different Seed Region Pairs of the Physical Phantom [12]

Algorithm	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)
completion Field	0.3731	0.1827	0.4272	0.020	0.0542	0.0002	0.0009	0
Probabilistic Tractography	0.0231	0	0.3559	0.2397	0	0	0	0




Fig. 5. Visualization of the completion fields obtained for four different sets of ROI pairs. The top left figure shows a pair with no underlying physical connection while the other three correspond to regions with real connections in the brain. Top row, from left to right: False connection, splenium. Bottom row, from left to right: Genu of corpus callosum, cortico-spinal tract.

7.3.4 Experiment on Human Brain Data

We now present experiments in which the 3D completion field algorithm is applied to *in vivo* human brain HARDI data. The diffusion weighted images were acquired *in vivo* from a healthy subject on a Siemens 3T Trio MR scanner using an 8-channel phased-array head coil. Diffusion encoding was achieved using a single-shot spin-echo planar sequence with twice-refocused balanced diffusion encoding gradients. A dataset designed for high-angular resolution reconstruction was acquired with 99 diffusion encoding directions, 2 mm isotropic voxel size, 63 slices, $b = 3,000 \text{ s/mm}^2$, $TE = 121 \text{ ms}$, $TR = 11.1 \text{ s}$, and GRAPPA parallel reconstruction. The series of experiments were performed on FOD reconstructions of the HARDI data of the human brain. We also only considered voxels with a fractional anisotropy value higher than 0.1.

We have examined the performance of the completion field algorithm on a selection of ROI pairs chosen from some of the major fiber tracts in the brain. This choice of ROIs, for which the associated anatomical connectivity is known, allows us to assess the validity of the obtained results. Out of the four chosen ROI pairs, one is not associated with any known connection. (The ROIs are actually located in the medial part of the corpus callosum, one anterior and one posterior. The slice shown below the ROIs is chosen such that the reconstructed completion field can be observed.) The other three pairs are chosen as part of the genu of the corpus callosum, the splenium of the corpus callosum, and finally as part of the cortico-spinal tract. On average, the computation for each ROI pair took about a day, with a state space dimension of $128 * 128 * 63 * 200$, while running on four 2.8 GHz Intel Xeon processors. Fig. 5 shows the computed completion fields for the four sets of source and sink regions. For reference, the results are

TABLE 3
Connectivity Indices for Four ROI Pairs of Human Data

ROIs	Splenium	Genu-CC	CrtSpinal	Medial-CC
Conn*1000	0.016	0.122	587.27	0

overlaid on the T1-weighted image and the source and sink regions are shown with red and green blobs. The top right and the two bottom subfigures in Fig. 5 illustrate true connections, while that on the top left shows the connectivity pattern of the ROI pair for which a connection should not exist. Table 3 also shows the associated connectivity measures for the four ROI pairs. As mentioned earlier, the probability value of every state is multiplied by a scalar and then transformed by a sigmoid function to aid in visualization, but this also leads to a reduction in transparency contrast between high and low probability states. The difference in the degree of connectivity is reflected in the connectivity measures which demonstrate that the completion field algorithm is capable of distinguishing true connections from false ones. The connectivity measure associated with the ROI pair located in the cortico-spinal tract is higher than that for the other two connections. Referring to Fig. 5, this higher connectivity strength is consistent with the presence of more trajectories connecting these two regions. In all cases the connectivity patterns obtained for the true connections are consistent with the known anatomy of the underlying regions.

8 CONCLUSION

The 3D stochastic completion field extends Williams and Jacobs' curve completion model in computer vision [1] to provide a probabilistic representation of the connectivity pattern between two given regions of interest. We have formulated the computation of the 3D stochastic completion in an efficient local parallel computational framework, exploiting the Fokker-Planck equation of the underlying 3D directional random walk. Further, we have introduced a rotation invariant numerical scheme by using a spherical harmonics basis to represent functions defined on the 2-sphere. Adding extra angular drift terms allows for incorporation of dense orientation data into the algorithm. This extension paves the way for applications of the stochastic completion field to computer vision and pattern analysis applications in 3D where dense orientation data arises. For the particular application we have focused on in this paper, the local nature of the proposed computational model allows diffusion MRI measurements to be directly incorporated.

We have demonstrated the potential of this algorithm for distinguishing real connections from false ones in diffusion MRI through experiments performed on synthetic data and on data from a physical phantom. We observe that as long as the SNR of the acquired data is within an acceptable range, the obtained connectivity measures can distinguish true connections from false ones, while the associated completion field renders a reliable probabilistic representation of the underlying connectivity patterns. The same observations hold for the physical

phantom, where ground-truth knowledge confirms the validity of the computed connectivity measures. We have also demonstrated that the proposed measure qualitatively outperforms the weakest link index. Additionally, although no ground truth is available for the human brain dataset, the connectivity patterns obtained are consistent with the known anatomy of the brain.

In future work, we plan to extend this approach to integrate measurements of the underlying noise and uncertainty in the dataset into the framework. In application to diffusion MRI, this can be done through proper selection of angular diffusion parameters at each voxel, based on the associated noise and uncertainty levels. The use of completion fields for probabilistic segmentation of the cortical-subcortical regions in the brain and for the classification of controls versus patients suffering from white matter diseases are additional fruitful directions to explore.

ACKNOWLEDGMENTS

The authors would like to thank Jennifer Campbell and Bruce Pike for many helpful discussions and for providing the fiber ODF data for our experiments. This work was supported by grants from NSERC and FQRNT Québec.

REFERENCES

- [1] L. Williams and D. Jacobs, "Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Saliency," *Neural Computation*, vol. 9, pp. 837-858, 1997.
- [2] L. Williams and D. Jacobs, "Local Parallel Computation of Stochastic Completion Fields," *Neural Computation*, vol. 9, pp. 859-881, 1997.
- [3] G. Kanizsa, *Organization in Vision: Essays on Gestalt Perception*, series Praeger Special Studies. Praeger, 1979.
- [4] S. Petry and G. Meyer, *The Perception of Illusory Contours*. Springer-Verlag, 1987.
- [5] S. Ullman, "Filling-In the Gaps: The Shape of Subjective Contours and a Model for Their Generation," *Biological Cybernetics*, vol. 25, no. 1, pp. 1-6, 1976.
- [6] D. Mumford, "Elastica and Computer Vision," *Algebraic Geometry and Its Applications*. chapter 31, pp. 491-506, Springer-Verlag, 1994.
- [7] B. Kimia, I. Frankel, and A. Popescu, "Euler Spiral for Shape Completion," *Int'l J. Computer Vision*, vol. 54, pp. 159-182, 2003.
- [8] P. Parent and S.W. Zucker, "Trace Inference, Curvature Consistency and Curve Detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 11, no. 8, pp. 823-889, Aug. 1989.
- [9] G. Guy and G. Medioni, "Inferring Global Perceptual Contours from Local Features," *Int'l J. Computer Vision*, vol. 20, pp. 113-133, 1993.
- [10] A. Shashua and S. Ullman, "Structural Saliency: The Detection of Globally Salient Structures Using a Locally Connected Network," *Proc. Second IEEE Int'l Conf. Computer Vision*, pp. 321-327. 1988,
- [11] J. Zweck and L.R. Williams, "Euclidean Group Invariant Computation of Stochastic Completion Fields Using Shiftable-Twistable Functions," *J. Math. Imaging and Vision*, vol. 21, pp. 135-154, 2004.
- [12] C. Poupon, B. Rieul, I. Kezele, M. Perrin, F. Poupon, and J. Mangin, "New Diffusion Phantoms Dedicated to the Study and Validation of High-Angular-Resolution Diffusion Imaging (HARDI) Models," *Magnetic Resonance in Medicine*, vol. 60, pp. 1276-1283, 2008.
- [13] D. LeBihan, E. Breton, D. Lallemand, P. Grenier, E. Cabanis, and M. Laval-Jeantet, "MR Imaging of Intravoxel Incoherent Motions: Application to Diffusion and Perfusion in Neurologic Disorders," *Radiology*, vol. 161, pp. 401-407, Nov. 1986.
- [14] M. Catani, R. Howard, S. Pajevic, and D. DK, "Virtual in Vivo Interactive Dissection of White Matter Fasciculi in the Human Brain," *Neuroimage*, vol. 17, pp. 77-94, 2002.
- [15] S. Frey, J. Campbell, G. Pike, and M. Petrides, "Dissociating the Human Language Pathways with High Angular Resolution Diffusion Fiber Tractography," *J. Neuroscience*, vol. 28, pp. 11435-11444, 2008.
- [16] M. Catani, D. Jones, and D. ffytche, "Perisylvian Language Networks of the Human Brain," *Annals of Neurology*, vol. 57, pp. 8-16, 2005.
- [17] A. Anwander, M. Tittgemeyer, D. von Cramon, A. Friederici, and T. Knösche, "Connectivity-Based Parcellation of Broca's Area," *Cerebral Cortex*, vol. 17, pp. 816-825, 2007.
- [18] M. Rushworth, T. Behrens, and H. Johansen-Berg, "Connection Patterns Distinguish 3 Regions of Human Parietal Cortex," *Cerebral Cortex*, vol. 16, pp. 1418-1430, 2006.
- [19] A. Venkataraman, Y. Rath, M. Kubicki, C. Westin, and P. Golland, "Joint Generative Model for fMRI/DWI and Its Application to Population Studies," *Proc. 13th Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 191-199, Sept. 2010.
- [20] M. Koch, D. Norris, and M. Hund-Georgiadis, "An Investigation of Functional and Anatomical Connectivity Using Magnetric Resonance Imaging," *Neuroimage*, vol. 16, pp. 241-250, 2002.
- [21] D. LeBihan, R. Turner, and P. Douek, "Is Water Diffusion Restricted in Human Brain White Matter? An Echo-Planar NMR Imaging Study," *NeuroReport*, vol. 4, pp. 887-890, July 1993.
- [22] S. Mori, B. Crain, V. Chacko, and P. van Zijl, "Three Dimensional Tracking of Axonal Projections in the Brain by Magnetic Resonance Imaging," *Annals of Neurology*, vol. 45, pp. 265-269, 1999.
- [23] P. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi, "In Vitro Fiber Tractography Using DT-MRI Data," *Magnetic Resonance in Medicine*, vol. 44, pp. 625-632, 2000.
- [24] G. Parker, C. Wheeler-Kingshott, and G. Barker, "Estimating Distributed Anatomical Connectivity Using Fast Marching Methods and Diffusion Tensor Imaging," *Trans. Medical Imaging*, vol. 21, pp. 505-512, 2002.
- [25] C. Poupon, C. Clark, V. Frouin, J. Regis, I. Bloch, D. LeBihan, and J. Mangin, "Regularization of Diffusion-Based Direction Maps for the Tracking of Brain White Matter Fascicles," *NeuroImage*, vol. 12, pp. 184-195, 2000.
- [26] J. Campbell, K. Siddiqi, V. Ryman, A. Sadikot, and B. Pike, "Flow-Based Fiber Tracking with Diffusion Tensor q-Ball Data: Validation and Comparison to Principal Diffusion Direction Techniques," *NeuroImage*, vol. 27, pp. 725-736, 2005.
- [27] P. Behrens et al., "Characterization and Propagation of Uncertainty in Diffusion-Weighted MR Imaging," *Magnetic Resonance in Medicine*, vol. 50, pp. 1077-1088, 2003.
- [28] D. Jones and C. Pierpaoli, "Confidence Mapping in Diffusion Tensor Magnetic Resonance Imaging Tractography Using a Bootstrap Approach," *Magnetic Resonance Medicine*, vol. 53, pp. 1143-1149, 2005.
- [29] J. Berman, S. Chung, P. Mukherjee, C. Hess, E. Han, and R. Henry, "Probabilistic Streamline q-Ball Tractography Using the Residual Bootstrap," *NeuroImage*, vol. 39, pp. 215-222, 2008.
- [30] P. Basser, J. Mattiello, and D.L. Bihan, "MR Diffusion Tensor Spectroscopy and Imaging," *Biophysical J.*, vol. 66, pp. 259-267, Jan. 1994.
- [31] D.S. Tuch, "Q-Ball Imaging," *Magnetic Resonance in Medicine*, vol. 52, pp. 1358-1372, Dec. 2004.
- [32] M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche, "Regularized, Fast, and Robust Analytical q-Ball Imaging," *Magnetic Resonance Medicine*, vol. 58, pp. 497-510, 2007.
- [33] J. Tournier, F. Calamante, D. Gadian, and A. Connelly, "Direct Estimation of the Fiber Orientation Density Function from Diffusion-Weighted MRI Data Using Spherical Deconvolution," *NeuroImage*, vol. 23, pp. 1176-1185, Sept. 2004.
- [34] S. Mori and P.V. Zijl, "Fiber Tracking: Principles and Strategies Technical Review," *NMR in Biomedicine*, vol. 15, pp. 468-480, 2002.
- [35] D. Jones, A. Simmons, S. Williams, and M. Horsfield, "Non-Invasive Assessment of Axonal Fiber Connectivity in the Human Brain Via Diffusion Tensor MRI," *Magnetic Resonance in Medicine*, vol. 42, pp. 37-41, 1999.
- [36] P. MomayyezSiahkal, J. Campbell, P. Savadjiev, G. Pike, and K. Siddiqi, "Beyond Crossing Fibres: Probabilistic Tractography of Complex Subvoxel Fibre Geometries," *Proc. MICCAI Workshop Diffusion Modelling and the Fibre Cup*, pp. 81-92, Sept. 2009.
- [37] J.G. Malcolm, M.E. Shenton, and Y. Rath, "Filtered Multi-Tensor Tractography," *IEEE Trans. Medical Imaging*, vol. 29, no. 9, pp. 1664-1675, Sept. 2010.

- [38] P. Fillard, C. Poupon, and J. Mangin, "A Novel Global Tractography Algorithm Based on an Adaptive Spin Glass Model," *Proc. 12th Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 927-934, Sept. 2009.
- [39] S. Jbabdi, P. Bellec, R. Toro, J. Daunizeau, M. Pelegriani-Issac, and H. Benali, "Accurate Anisotropic Fast Marching for Diffusion-Based Geodesic Tractography," *Int'l J. Biomedical Imaging*, vol. 2008, pp. 1-12, 2008.
- [40] P. Momayyez and K. Siddiqi, "Rotation Invariant Completion Fields for Mapping Diffusion-MRI Connectivity," *Proc. 22nd Int'l Conf. Information Processing in Medical Imaging*, 2011.
- [41] P. Batchelor, D. Hill, F. Calamante, and D. Atkinson, "Study of Connectivity in the Brain Using the Full Diffusion Tensor from MRI," *Proc. 17th Int'l Conf. Information Processing in Medical Imaging*, pp. 121-133, 2001.
- [42] L. O'Donnell, S. Haker, and C. Westin, "New Approaches to Estimation of White Matter Connectivity in Diffusion Tensor MRI: Elliptic PDEs and Geodesics in a Tensor-Warped Space," *Proc. Fifth Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 459-466, Sept. 2002.
- [43] M. Péchaud, M. Descoteaux, and R. Keriven, "Brain Connectivity Using Geodesics in Hardi," *Proc. 12th Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 482-489, Sept. 2009.
- [44] N. Hageman, A. Toga, K. Narr, and D. Shattuck, "A Diffusion Tensor Imaging Tractography Algorithm Based on Navier-Stokes Fluid Mechanics," *IEEE Trans. Medical Imaging*, vol. 28, no. 3, pp. 348-360, Mar. 2009.
- [45] P. Fletcher, R. Tao, W. Jeong, and R. Whitaker, "A Volumetric Approach to Quantifying Region-to-Region White Matter Connectivity in Diffusion Tensor MRI," *Proc. 20th Int'l Conf. Information Processing in Medical Imaging*, pp. 346-358, 2007.
- [46] E. Pichon, C. Westin, and A. Tannenbaum, "A Hamilton-Jacobi-Bellman Approach to High Angular Resolution Diffusion Tractography," *Proc. Eighth Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 180-187, 2005.
- [47] P. Momayyez and K. Siddiqi, "3D Stochastic Completion Fields for Fiber Tractography," *Proc. IEEE Workshop Math. Methods in Biomedical Image Analysis*, pp. 178-185, June 2009.
- [48] H. Risken, *The Fokker-Planck Equation: Methods of Solution and Applications*. Springer, 1996.
- [49] P. Savadjiev, J. Campbell, and B.P.K. Siddiqi, "3D Curve Inference for Diffusion MRI Regularization and Fibre Tractography," *Medical Image Analysis*, vol. 10, pp. 799-813, Aug. 2006.
- [50] B. Fornberg and D. Merrill, "Comparison of Finite Difference and Pseudospectral Methods for Convective Flow over a Sphere," *Geophysical Research Letters*, vol. 24, pp. 3245-3248, 1997.
- [51] D. Alexander, G. Barker, and S. Arridge, "Detection and Modeling of Non-Gaussian Apparent Diffusion Coefficient Profiles in Human Brain Data," *Magnetic Resonance in Medicine*, vol. 48, pp. 331-340, 2002.
- [52] A. Anderson, "Measurement of Fiber Orientation Distributions Using High Angular Resolution Diffusion Imaging," *Magnetic Resonance in Medicine*, vol. 54, no. 5, pp. 1194-1206, 2005.
- [53] T. Yo, A. Anwander, M. Descoteaux, P. Fillard, C. Poupon, and T. Knösch, "Quantifying Brain Connectivity: A Comparative Tractography Study," *Proc. 12th Int'l Conf. Medical Image Computing and Computer-Assisted Intervention*, pp. 886-893, Sept. 2009.
- [54] B. Jeurissen, A. Leemans, D.K. Jones, J.D. Tournier, and J. Sijbers, "Probabilistic Fiber Tracking Using the Residual Bootstrap with Constrained Spherical Deconvolution," *Human Brain Mapping*, vol. 32, pp. 461-479, 2011.
- [55] O. Ciccarelli, A.T. Tossy, G.J. Parker, C.A. Wheeler-Kingshott, G.J. Barker, and P.A. Boulby, "From Diffusion Tractography to Quantitative White Matter Tract Measures: A Reproducibility Study," *NeuroImage*, vol. 18, pp. 348-359, 2003.



Parya MomayyezSiahkal received the BSc degree in electrical engineering from Sharif University of Technology, Iran, in 2004, the MEng degree in signal processing from McGill University in 2007, and the PhD degree in computer science from the School of Computer Science, McGill University, Canada, in 2012. Her fields of interest include shape analysis, medical imaging, computer vision, and pattern recognition.



Kaleem Siddiqi received the BS degree from Lafayette College in 1988 and the MS and PhD degrees from Brown University in 1990 and 1995, respectively, all in the field of electrical engineering. He is currently a professor, William Dawson scholar, and an associate director research of the School of Computer Science at McGill University. He is also a member of McGill's Centre for Intelligent Machines. Before moving to McGill in 1998, he was a postdoctoral associate in the Department of Computer Science at Yale University (1996-1998) and held a visiting position in the Department of Electrical Engineering at McGill University (1995-1996). His research interests include the areas of computer vision, image analysis, and medical imaging. He is a member of the Phi Beta Kappa, Tau Beta Pi, and Eta Kappa Nu. He is a senior member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.