ELSEVIER

NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage xx (2005) xxx – xxx

Flow-based fiber tracking with diffusion tensor and q-ball data: Validation and comparison to principal diffusion direction techniques

Jennifer S.W. Campbell,[a,b,*] Kaleem Siddiqi,[b] Vladimir V. Rymar,[a] Abbas F. Sadikot,[a] and G. Bruce Pike[a]

[a]*McConnell Brain Imaging Centre, McGill University, Canada*
[b]*Centre for Intelligent Machines, McGill University, Canada*

Received 29 June 2004; revised 11 April 2005; accepted 5 May 2005

In this study, we evaluate the performance of a flow-based surface evolution fiber tracking algorithm by means of a physical anisotropic diffusion phantom with known connectivity. We introduce a novel speed function for surface evolution that is derived from either diffusion tensor (DT) data, high angular resolution diffusion (HARD) data, or a combined DT-HARD hybrid approach. We use the model-free q-ball imaging (QBI) approach for HARD reconstruction. The anisotropic diffusion phantom allows us to compare and evaluate the performance of different fiber tracking approaches in the presence of real imaging artifacts, noise, and subvoxel partial volume averaging of fiber directions. The surface evolution approach, using the full diffusion tensor as opposed to the principal diffusion direction (PDD) only, is compared to PDD-based line propagation fiber tracking. Additionally, DT reconstruction is compared to HARD reconstruction for fiber tracking, both using surface evolution. We show the potential for surface evolution using the full diffusion tensor to map connections in regions of subvoxel partial volume averaging of fiber directions, which can be difficult to map with PDD-based methods. We then show that the fiber tracking results can be improved by using high angular resolution reconstruction of the diffusion orientation distribution function in cases where the diffusion tensor model fits the data poorly.

Keywords: Diffusion tensor imaging; Q-ball imaging; Fiber tracking; Level set methods

Introduction

Magnetic resonance diffusion imaging has a unique ability to provide information about the organization of fibrous tissue structures in vivo. It does so via the estimation of the 3D displacement distribution of diffusing water molecules, i.e., the diffusion probability density function (diffusion pdf). Water diffusion in fibrous tissue is anisotropic, with the preferred direction of diffusion lying along the dominant fiber orientation. This is of particular interest in the central nervous system (CNS), where diffusion imaging has the potential to assess neuronal connectivity, an application which has widespread implications for basic neuro-anatomical research and for disease detection and assessment.

The first work on 3D fiber reconstruction in the CNS used diffusion *tensor* (DT) data, so named because it is obtained by modeling the diffusion pdf as an anisotropic 3D Gaussian function, which can be described by a second order tensor (Basser et al., 1994). The reconstruction of tracts was done by line propagation using the principal eigenvector of the diffusion tensor (Basser et al., 2000; Conturo et al., 1999; McGraw et al., 2004; Mori et al., 1999; Vemuri et al., 2002). Such principal diffusion direction (PDD) techniques can be confounded when there is more than one fiber direction within a single imaging voxel. With voxel sizes typical of diffusion acquisitions (10–30 mm^2), there is significant partial volume averaging of fiber directions in anatomical regions of both research and clinical interest, such as the association fibers near the cortex. This partial volume averaging may be due to high curvature, crossing, branching, or splaying of tracts.

A number of solutions have been proposed to deal with the problem of subvoxel partial volume averaging of fiber directions. The diffusion tensor itself contains information about multiple fiber directions: for example, when the fibers are restricted to a plane, a level set of the tensor-described diffusion pdf is a planar ellipsoid. Although the principal eigenvector direction may not reflect the fiber directions in this case, the full tensor may be used for tracking, as has been proposed by several groups including ourselves (Batchelor et al., 2001; Campbell et al., 2002a; Lazar et al., 2003; O'Donnell et al., 2002; Tournier et al., 2003). Alternatively, given sufficient diffusion weighted images (DWIs), bootstrap methods may be used to estimate confidence intervals or a marginal posterior distribution for the direction of

* Corresponding author. McConnell Brain Imaging Centre, Montreal Neurological Institute, 3801 University Street, Room WB325, Montreal, QC, Canada H3A 2B4. Fax: +1 514 398 2975.
E-mail address: jcampbel@bic.mni.mcgill.ca (J.S.W. Campbell).

1053-8119/$ - see front matter © 2005 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2005.05.014

the principal eigenvector (Behrens et al., 2003; Jones, 2003). Although the full tensor does give us more information than its principal eigenvector alone, there has recently been increasing interest in forgoing the classic tensor description of the diffusion pdf in favor of high angular resolution diffusion (HARD) pdf reconstruction to infer multiple fiber directions within single voxels. HARD reconstruction techniques include multi-tensor modeling (Tuch et al., 2002), diffusion spectrum imaging to estimate the full 3D diffusion pdf using q-space methods (Wedeen et al., 2000), and model-free extraction of the radially persistent angular structure using either maximum entropy solutions (Alexander and Jansons, 2002) or modified q-space methods (Tuch et al., 2003). Additionally, regularization techniques can be used to infer multiple fiber directions from a single-tensor field (Cointepas et al., 2002; Ramírez-Manzanares and Rivera, 2003). In simple fiber systems involving partial volume averaging of fiber directions, it has been shown that the directions in which the high angular resolution diffusion pdf is maximal coincide with fiber directions (Lin et al., 2002). We note that there are also methods for measuring the apparent diffusion coefficient (ADC) profile at high angular resolution (Alexander et al., 2002; Frank, 2002; Zhan et al., 2003), however, the maxima of the ADC profile do not necessarily coincide with fiber directions (Tuch et al., 2002; von dem Hagen and Henkelman, 2001), hence further processing is necessary for fiber tracking applications.

In this paper, we extend our previously proposed flow-based fiber tracking approaches (Campbell et al., 2002a,b) in order to use the information in the full diffusion pdf estimated using either HARD or DT techniques. This algorithm is a modification and extension of the Fast Marching Tractography (FMT) technique (Parker et al., 2002a), which uses the principal diffusion direction only, and is similar to several surface evolution approaches that use the full diffusion tensor (Batchelor et al., 2001; O'Donnell et al., 2002; Tournier et al., 2003). Our extension of the FMT technique consists of using all of the information in the diffusion tensor or, if available, HARD measurements. The surface evolution approach has the advantage that it allows tracking to proceed in a continuum of directions, as may be desired in cases where there is uncertainty and/or multiple fiber directions. We show fiber tracking results using both HARD and DT data in the human brain, noting that the diffusion tensor can provide more information than that given by its principal eigenvector only and that HARD reconstruction can give us more information still. Additionally, we quantitatively compare the performance of the flow-based approach to that of line propagation using the PDD. To do so, we designed a physical phantom with known connectivity from excised rat spinal cord. Previous validation studies have been done using a single excised cord (Campbell et al., 2002b; Vemuri et al., 2002) and in the macaque (Parker et al., 2002b): in this study, complex known configurations of subvoxel curvature and fiber crossing are created and scanned at a standard human imaging resolution. While simulated data can provide a gold standard to which tracking results can be compared (Lazar and Alexander, 2003; Lori et al., 2002; Tournier et al., 2002), it is of additional value to evaluate the results of tracking on real MRI data in the presence of normal imaging artifacts, noise characteristics, and voxel size limitations. Doing so allows us to validate the fiber tracking process from acquisition of the MRI data, to estimation of the diffusion displacement distribution, and to the tracking algorithm itself.

Materials and methods

Flow-based fiber tracking: implementation

Flow and assignment of connectivity index

For flow-based fiber tracking, we use the level set surface evolution techniques of Osher and Sethian (1988), which were first applied to fiber tracking by Parker et al. (2002a). Let $S(x, t)$ be a surface with initial value $S_o = S(x, 0)$ given by the boundary of a user-defined seed voxel or region. S is evolved outward with speed F along its normal, $\hat{n} = (1, \theta_n, \phi_n)$:

$$\frac{\partial S}{\partial t} = F\hat{\boldsymbol{n}}. \tag{1}$$

The speed function $F(x \mid \hat{n})$ is set equal to our confidence that a tract exists in the direction $\hat{n}$. Our speed function uses all of the information in the diffusion tensor or, if available, HARD measurements. The definition of F is described in more detail in the next section. The evolution of S is equivalent to simultaneously propagating a continuous field of curves outward from the seed S_o. As S passes through a voxel, a time of arrival T is assigned to that voxel. T is the solution of the Eikonal equation, $|T|F = 1$. To calculate putative tracts as 3D curves, we retrace the path of normals to S that led to each voxel reached: this path is given by gradient descent through the map of the time of arrival of the surface, $T(x)$, from each point x to the seed. The curves obtained through this surface evolution approach should be similar to those obtained by a Monte Carlo or iterative tracking approach such as those described in Behrens et al. (2003), Björnemo et al. (2002), Hagmann et al. (2003), and Koch et al. (2002). With surface evolution, the propagation of all the curves is simultaneous and continuous and therefore potentially faster, especially as the number of putative paths increases. Fig. 1 shows a schematic of the tracking algorithm. We note that this surface evolution scheme is the first to accept either DT or HARD data as input.

In order to quantify our confidence in the existence of each reconstructed tract, we define a connectivity index (CI). We have a discrete array F_n of speed function values at each step along each

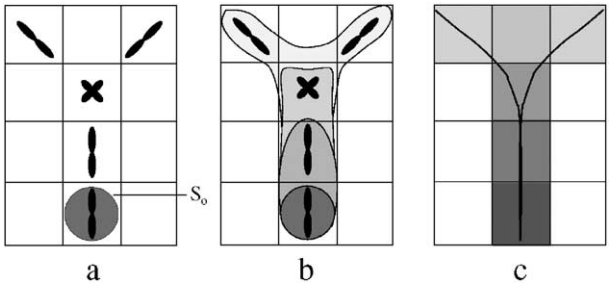

Fig. 1. 2D schematic of flow-based surface evolution tractography. (a) The input data consist of the normalized and thresholded diffusion ODF, Ψ_t, at each voxel (the voxels with no shape shown represent voxels with isotropic diffusion, hence no fiber structure). The user defines a seed region, and a surface S is initialized to be the outer edge of the seed region (S_o). (b) The surface evolves at speed F along its normal $\hat{n}$, with F given by the value of Ψ_t in the direction $\hat{n}$. As S evolves, a scalar map T of the time of arrival of the surface is constructed. (c) Gradient descent through the time of arrival map (shown as grayscale) gives the fiber tract reconstructions as curves.

tract. We define the CI to be the lowest value in the array of speed functions values:

$$\mathrm{CI} = \min(F_n). \tag{2}$$

The CI is a measure of our certainty in a tract's existence given the measurements we have made. It reflects uncertainty due to partial volume averaging of tract directions in cases of crossing, branching, curvature, and tracts forming a "bottleneck"; partial volume averaging of fibers with isotropic material, which leads to low anisotropy; and imaging noise, which can confound the fiber directions. The CI map gives a relative certainty in the existence of all tracts reconstructed through one seed voxel or region.

Speed function for surface evolution

For the speed function F in Eq. (1), we desired a function that reflected our confidence that a fiber is oriented in the direction of the surface propagation. This function should reflect subvoxel partial volume averaging of fiber directions due to branching, crossing, and curvature, as well as measurement errors. Ideally, this function should handle the presence of a *distribution* of fiber directions in each voxel, such that the resulting connectivity indices for the reconstructed tracts can be interpreted as probabilities. The rigorous estimation of such a fiber orientation distribution function (fiber ODF) for both DT and HARD data is beyond the scope of this paper, however, we propose a speed function based on the diffusion orientation distribution function (diffusion ODF). In the case of crossing fibers, the ODF measured using HARD techniques can be expected to have maxima along the fiber directions (Lin et al., 2002), and the ODF measured using the DT technique can be expected to be "flat", i.e., to give some information about the fiber orientations and to give more information than the principal diffusion direction only. The speed function estimated from HARD data will be similar to that from DT data in the case of subvoxel curvature. This case is particularly suited to surface evolution because there are no discrete fiber orientations within the voxel, but a continuum of directions. The diffusion ODF will be high along all of these directions.

We define the diffusion orientation distribution function, Ψ, to be the normalized projection of the 3D diffusion pdf on the surface of the unit sphere:

$$\Psi(\mathbf{u}) = \frac{\Psi^*(\mathbf{u})}{\bar{\Psi}^*}, \tag{3}$$

where

$$\Psi^*(\mathbf{u}) = \int_0^\infty P(\mathbf{ru})\tau_d)d\mathbf{r}. \tag{4}$$

Here, $\mathbf{u} = (1, \theta, \phi)$ is a unit vector, and $P(\mathbf{r}\,|\,\tau_d)$ is the probability that a water molecule displaces by vector $\mathbf{r}$ during the time over which diffusion is observed, τ_d. $\bar{\Psi}^*$ is the mean value of Ψ^* over all values of $\mathbf{u}$. Ψ can be obtained from either diffusion tensor or high angular resolution diffusion measurements. We set the speed function, F, for surface evolution equal to the diffusion ODF value in the direction of the surface normal, $\hat{n}$:

$$F(\hat{\boldsymbol{n}}) = \Psi(\hat{\boldsymbol{n}}). \tag{5}$$

This generalized diffusion ODF formulation for the speed function F is an extension of previous full-tensor approaches for flow-based fiber tracking (Batchelor et al., 2001; Campbell et al., 2002a; O'Donnell et al., 2002). Using the full diffusion ODF to drive the flow, even in the diffusion tensor case, should improve results over those using the principal eigenvector only, as was used in Parker's FMT implementation (Parker et al., 2002a).

Using F, as defined in Eq. (5), allows the evolving surface S to flow through regions where there is ambiguity or low confidence in fiber direction: for example, regions of partial volume averaging of fiber directions due to subvoxel crossing, 'bottlenecks', or bending; regions of partial volume averaging of isotropic gray matter and fiber structure (e.g., the thalamus); or voxels corrupted by noise and/or imaging artifacts (where the pdf maximum may be slightly shifted from the correct direction). The speed of propagation is simply reduced in these regions, and the ambiguity at these points will be reflected in the connectivity index. When diffusion tensor measurements are used, F is suitable even when there is partial volume averaging of fiber directions: when there are multiple fiber directions in one plane, the tensor will be "flat". The fiber direction within this plane is unknown, and the speed of surface evolution will reflect this by being nonzero everywhere in this plane and zero outside of it. However, HARD data can be useful for refining the speed function further by giving it clear maxima along the fiber directions in the case of fibers crossing within a voxel. This refinement can make the flow more controlled and more accurate, especially in cases where the maximum of the tensor-described ODF does not lie along either fiber direction.

In many cases in the human brain, we can expect the diffusion tensor model to be adequate, and the use of HARD data may increase susceptibility to noise-induced artifacts because it is model-free, whereas the tensor ODF can be described with only five parameters. When HARD data are available, we propose to first calculate the diffusion tensor and use the HARD data for surface evolution only in cases where the diffusion tensor model fits the data poorly. We do this by calculating the χ^2 statistic for the linear regression used to calculate the diffusion tensor (Shrager et al., 2002), and using HARD reconstruction when $Q(\chi^2)$ is low (Press et al., 1992) (in our implementation, when $Q(\chi^2) < 0.001$). The standard deviation is estimated by measuring the mean of the signal intensity in a region outside of the object being imaged and multiplying by $\sqrt{2/\pi}$ (Henkelman, 1985). In this paper, we use q-ball imaging (QBI) to obtain high angular resolution estimates of the diffusion ODF (Tuch et al., 2003).We shall refer to the surface evolution algorithm run with diffusion tensor data as surface-DT, the algorithm run with QBI data only surface-QB, and the surface evolution run with diffusion tensor data where the diffusion tensor fit is good and QBI data elsewhere will be called surface-hybrid.

We note that in order to obtain the fiber ODF from the diffusion ODF, we must deconvolve the diffusion ODF with the single-fiber response function. However, the single-fiber response function is not known. The exact mechanisms responsible for the observed white matter diffusion signal remain a topic of research. Simple models are currently being investigated for the estimation of the fiber ODF (Tournier et al., 2004), however, the response to a single fiber can be expected to change throughout the white matter in the brain due to microstructural differences between, for example, association and commissural fibers, and could also change in the presence of pathology. This issue was not addressed in this work, however, the surface evolution algorithm we have described was originally designed to use the fiber ODF, if available. We note that the use of the diffusion ODF to drive the flow results in a connectivity index that is highly weighted by diffusion anisotropy. Additionally, tracts passing through regions of partial volume averaging of fiber directions will necessarily be assigned lower connectivity indices.

The definition of a fiber ODF that assigns equal values to all fiber directions independent of the number of fibers (but dependent on other sources of uncertainty) would improve the interpretability of the connectivity indices. We note also that the relationship between the HARD diffusion ODF and the fiber ODF can be expected to be monotonic, meaning the voxels assigned the highest connectivity indices using the HARD diffusion ODF should also be those assigned the highest connectivity indices using a true fiber ODF.

Implementation details

Using the diffusion ODF to define the speed function F, Eq. (5) can be problematic in regions of high curvature such as the scenario illustrated in Fig. 2. Hence, we first threshold the normalized ODF at its mean in order to restrict flow in regions of isotropic diffusion:

$$\Psi^{**}(\hat{n}) = \Psi(\hat{n}) \quad \text{if } \Psi(\hat{n}) > 1$$
$$= 0 \qquad \text{if } \Psi(\hat{n}) \le 1, \tag{6}$$

and use Ψ^{**} in Eq. (5) instead of Ψ. The threshold is necessary only in regions of high curvature because, in other regions, flow in directions where the diffusion ODF is less than or equal to its mean will result in tracts with low connectivity indices. We wish to be able to track through structures with low anisotropy, such as the thalamus and regions near the cortex, therefore a threshold of significantly greater than 1 is also undesirable. The more anisotropic the underlying diffusion pdf, the smaller the solid angle in which $F(\hat{n})$ will be assigned a nonzero value.

We use stopping criteria based on both the fractional anisotropy (FA) (Pierpaoli and Basser, 1996) and the tract curvature. The speed function is set to zero when FA < 0.1, and the CI is set to zero if the tract curvature from voxel to voxel (calculated after the gradient descent step) exceeds 80. The FA threshold of 0.1 should allow tracking through regions of low anisotropy, such as deep gray matter structures, but should stop the flow in voxels with isotropic diffusion, for instance, the ventricles (Parker et al., 2002b). We note that, given a suitable speed function F, the threshold on anisotropy is not actually necessary but is used in this comparison with line propagation so as to use the same anisotropy threshold for both algorithms.

Before calculating the connectivity index Eq. (2), we first smooth the array of speed function values (F_n) with a one dimensional Gaussian kernel with a full width at half its maximum of three voxels. This step is essentially a regularization step that reduces sensitivity to isolated corrupted voxels, to which this worst-case connectivity index is susceptible, by blurring along the tract direction only.

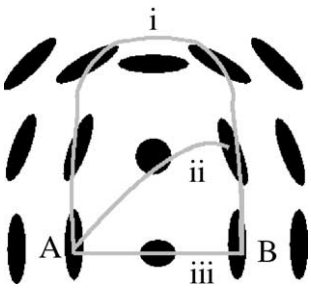

Fig. 2. Thresholding the diffusion ODF: the threshold of $\Psi > 1$ is in place to eliminate the case, possible in regions of high curvature, where the flow from point A reaches point B by path ii or iii before path i.

Line propagation using PDD: implementation

For comparison with the flow-based tractography approach, we chose the FACT (Fiber Assignment with Continuous Tracking) algorithm (Mori et al., 1999). The FACT method is well documented in the literature (Mori et al., 2001, 2002; Stieltjes et al., 2001; Wakana et al., 2003; Xue et al., 1999) and has been compared to other integration methods such as RK4 and Euler's method: it has been shown to have higher precision in synthetic divergent diffusion tensor fields and, along with RK4, to have superior accuracy in synthetic curved fields (Lazar and Alexander, 2003). We expect both divergent and curved fields in the human brain, therefore this integration method was a reasonable choice for comparison. The stopping criteria described above for the flow-based tracking were used. For each voxel in the seed region, we assigned multiple subvoxel seed points: 27 evenly spaced seed points were positioned on a grid within the voxel. This allows for branching of tracts from a single-voxel seed point. From each seed point, the reconstructed tract follows the direction of the principal eigenvector $\mathbf{e}_1$ of the diffusion tensor in that voxel until it enters another voxel, at which point the direction of propagation immediately changes to that of $\mathbf{e}_1$ of the new voxel, and tracking continues until the stopping criteria occur. Voxels are assigned CI = 1 if a reconstructed tract passes through them and CI = 0 if not. In this paper, the FACT algorithm is run only with diffusion tensor data and will be called FACT-DT for clarity.

Anisotropic diffusion phantom experiments

Construction

For evaluation of the surface evolution with DT and HARD data and for comparison to line propagation, a physical phantom with known connectivity was constructed from excised rat spinal cord. We desired a phantom that restricts or hinders water diffusion on a length scale we can probe in a diffusion MR experiment (i.e., on the order of 10 μm) and that has suitable MR relaxation times and signal strength. Phantom construction and scanning were performed on two separate occasions. On each occasion, two Sprague–Dawley rats, aged 4–12 months, were euthanized and their spinal cords surgically excised. The fresh cords were embedded in 2% agar in a configuration designed to have curved, straight, and crossing tracts. The cords were 7–12 cm long and 5 mm in diameter. A third formalin fixed cord was added to one of the phantoms in order to explore the option of using fixed tissue in a phantom for repeated measurements. The fixed cord was not used in this study. The phantom is shown in Fig. 3.

MRI acquisition and diffusion ODF estimation

The cords were scanned 1 h after the surgeries with a Siemens 1.5T Sonata MR scanner (Siemens Medical Systems, Erlangen, Germany) using a knee coil. A single-shot spin-echo echo planar sequence with twice-refocused balanced gradients, designed for minimization of eddy current artifacts, was used (Reese et al., 2003). For diffusion tensor reconstruction, four coregistered datasets were acquired, consisting of 90 diffusion weighted images with isotropically spaced diffusion weighting directions (b = 1300 s/mm^2, TR = 8s, TE = 110 ms, 2.5 mm isotropic voxels, 40 slices), as well as 10 images with b = 0 s/mm^2 and otherwise identical imaging parameters. The scanning time for the diffusion base images was approximately 15 min. The diffusion encoding directions were

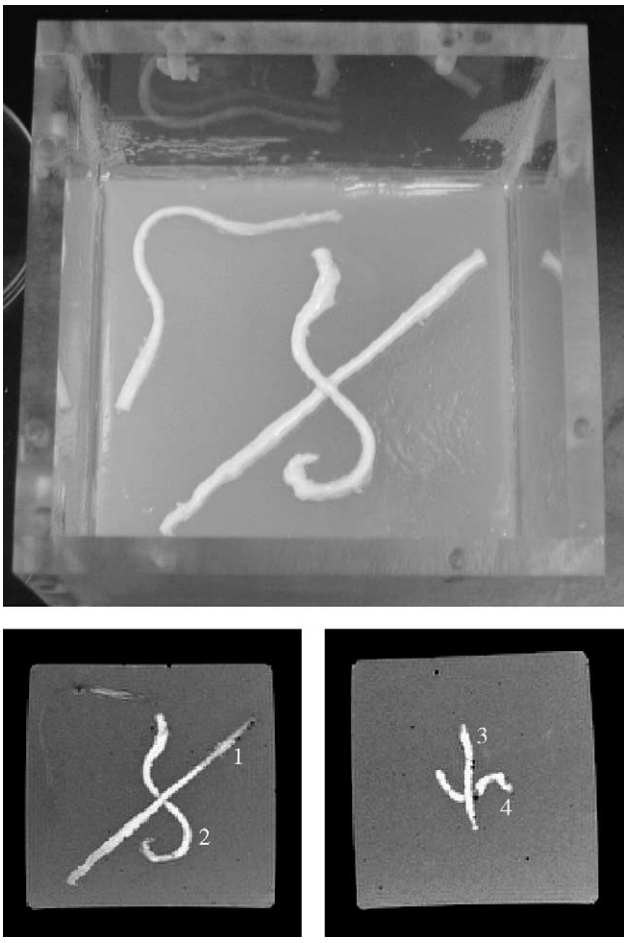

Fig. 3. Anisotropic diffusion phantom: the phantom is constructed from excised rat spinal cords embedded in 2% agar. Top: Photograph of one of the two phantoms. Bottom: T_1 weighted images of both phantoms. The tracts, two per phantom, are numbered for future reference in this paper.

calculated using an electrostatic repulsion algorithm similar to that described in Jones et al. (1999). A 1 mm isotropic resolution T_1-weighted gradient echo scan was also performed (TR = 22 ms, TE = 9.2 ms, $\alpha = 30°$). The data were log-transformed, and the diffusion tensor was fitted using least-squares linear regression.

Q-ball data were acquired only for one of the two phantoms. The same 90 diffusion weighting directions used for the diffusion tensor acquisition were used, with $b = 3000$ s/mm^2, $q = 0.35$ μm^{-1}, TR = 8 s, TE = 110 ms, 2.8 mm isotropic voxels, and 30 slices. Four signal averages were acquired. The value of $\Psi(\mathbf{u})$ is given by the Funk Radon transform of the signal (Tuch et al., 2003):

$$\Psi(\mathbf{u})\left[S(\mathbf{u}_q)\right] = \int_{\mathbf{u}_q \perp \mathbf{u}} S(\mathbf{u}_q) d\mathbf{u}_q, \tag{7}$$

where $S(\mathbf{u}_q)$ represents the diffusion weighted signal strength measured with $q = q\mathbf{u}_q$.

This QBI acquisition is of significantly lower angular resolution than the 492-direction scheme introduced by Tuch et. al., but the 90-direction scheme was deemed sufficient to see multiple fiber crossings and is feasible in a clinically acceptable scanning time (13 min for one signal average). The different voxel sizes used for DTI and QBI were chosen because they are representative of voxel sizes used in the literature (Jones et al., 2002; Tuch et al., 2003) and reflect the inevitable tradeoff between spatial and angular resolution for any fixed scanning time.

Tractography evaluation

For each phantom, digital gold standard maps (GS) of each of the two tracts were defined. GS is a binary map with a value of one in all voxels identified to lie on that tract. The gold standard maps were defined by combined automatic and manual segmentation of the scalar image of the trace of the diffusion tensor, using a threshold of trace (**D**) < 1.0 E-6 mm^2/ms. The trace map was generated from the average of all signal averages of the $b = 1300$ s/mm^2, 2.5 mm isotropic voxel size, dataset. The two tracts crossed in one region, therefore some voxels were identified to lie in both.

In each gold standard tract, five seed regions of interest (ROIs), evenly spaced along the cord, were defined: these consisted of one voxel thick cross-sections of the cord. To evaluate the performance of surface evolution with q-ball reconstruction, the gold standard and seed ROI maps defined for the 2.5 mm isotropic dataset were resampled using trilinear interpolation to match the 2.8 mm isotropic dataset.

For evaluation, the CI map obtained from the surface evolution was converted to a binary map by assigning a value of one to all voxels with CI above an optimized threshold and zero to all voxels with CI below this threshold. For each binary CI map, we calculated an error measure d, given by the sum of the Euclidian distances between each gold standard voxel and the nearest tracked voxel and the Euclidian distances between the remaining tracked voxels and the gold standard. d was then normalized by dividing by the sum of the distances between the gold standard GS and the seed region. Hence, the error measure is equal to 1 if the tracking goes nowhere and is zero in the case of perfect tracking. The error measure penalizes both going off track ("false positives") and incomplete tracking ("false negatives"). The degree to which these errors are weighted depends on the Euclidian distance by which the tracking was incorrect. This error measure is one of many possible error measures: another option would have been the more standard kappa coefficient, but we desired a measure that reflects the Euclidian distance between the gold standard and reconstructed pathways. This is similar to the distance measure used for evaluation of synthetic tracking experiments by Tournier et al. (2002).

The surface-DT and FACT-DT algorithms were evaluated for the following datasets: (i) diffusion tensor reconstruction of the 90 diffusion encoding direction acquisition, $b = 1300$ s/mm^2 (for which we have 4 separate datasets consisting of one signal average each); (ii) diffusion tensor reconstruction of a subset of 30 of the 90 diffusion encoding directions used in (i); (iii) diffusion tensor reconstruction of the average of the four 90 diffusion encoding direction acquisitions, $b = 1300$ s/mm^2; and (iv) diffusion tensor reconstruction of the average of the four 90 diffusion encoding direction acquisitions with $b = 3000$ s/mm^2. Surface-QB and surface-hybrid were then evaluated using the average of the four 90 diffusion encoding direction acquisitions with $b = 3000$ s/mm^2. The region $Q(\chi^2) < 0.001$ consisted of the crossing region, most of the curved parts of tract #2, and some parts of the straight tract #1. The high χ^2 in some parts of the straight cord was most likely due to partial volume averaging of nerve roots with the spinal column.

For all tracking experiments, the raw data were first super-sampled to half the original voxel size in each dimension using trilinear interpolation. For a single seed region, FACT and surface evolution took approximately 10 min each on a 1533 MHz AMD

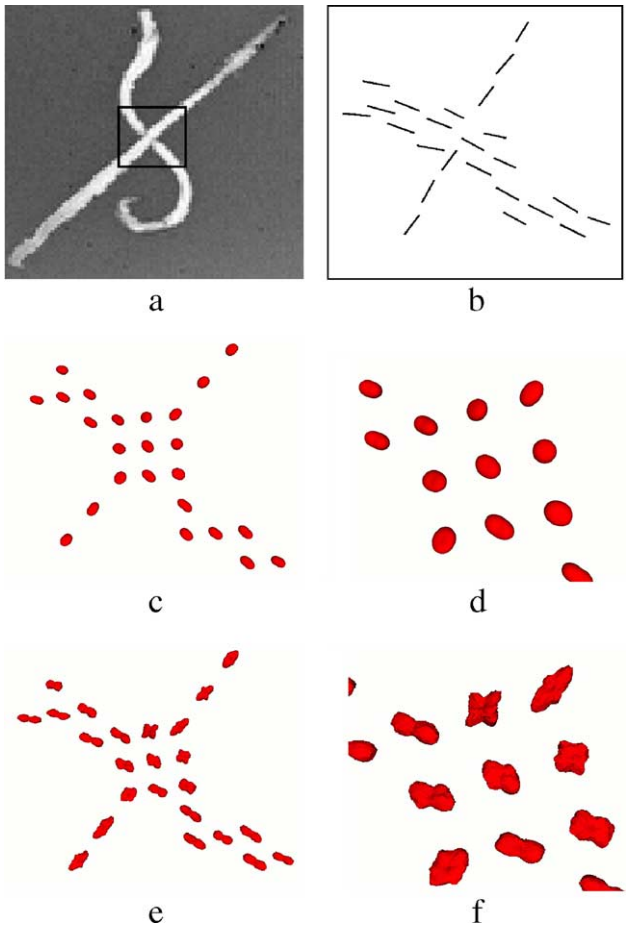

Fig. 4. Evidence of subvoxel partial volume averaging of fiber directions in the anisotropic diffusion phantom. (a) Crossing region. (b) Preferred direction of diffusion as given by the principal eigenvector of the diffusion tensor. Vectors are shown only within the gold standard map. (c) DTI reconstruction of the diffusion ODF. (d) Zoom of central region. (e) QBI reconstruction from the same diffusion weighted images. (f) Zoom of central region. There is support for a fiber running through the crossing region from bottom left to top right in the QBI data, but streamline tracking using the principal eigenvector can easily be confounded by the presence of a perpendicular fiber. The DTI shapes in this region are pancake-like, so flow in all directions above the mean of the ODF will result in both tracts being reconstructed.

Athlon processor, however, neither implementation was optimized. The FACT tracking was implemented in *Matlab*^{R} (The Mathworks, Inc., Natick, MA) and can be sped up by implementation in C. The surface evolution was implemented C but can be sped up by using efficient narrow-band level set methods (Osher and Sethian, 1988).

Human brain experiments

Human brain tractography is demonstrated using data acquired for one healthy subject and the same imaging protocol given above for both diffusion tensor and q-ball reconstruction. One signal average was acquired for both the b = 1300 s/mm^2 and b = 3000 s/mm^2 acquisitions. An eight-channel phased-array head coil equipped with a custom immobilization device was used instead of a knee coil. The b = 1300 s/mm^2 and b = 3000 s/mm^2 datasets were acquired in separate sessions and were subsequently registered using semi-automatic techniques (Collins et al., 1994). Informed consent was obtained prior to participation in the study. Tracking was performed with surface evolution and FACT, and the tracking results from each algorithm were compared to each other and to known human anatomy. Both the FACT-DT and surface-DT tracking were performed with the optimized b = 1300 s/mm^2 acquisition, while surface-hybrid tracking was performed using the b = 3000 s/mm^2 acquisition.

Results

Phantom experiments

Surface evolution compared to line propagation: diffusion tensor model

In our first experiments, we compared the performance of surface evolution using the *tensor* ODF (the ODF is shown in Figs. 4c and d) to the performance of FACT using the principal eigenvector of the diffusion tensor (Fig. 4b). These results are included in Table 1. For a single signal average of diffusion weighted measurements in 90 directions, surface-DT performs significantly better on average ($P < 0.001$), with a mean d of 0.18 ± 0.02 versus 0.34 ± 0.03 for FACT-DT. These results are for a 15 min acquisition time, which is accepted to be a reasonable acquisition time for clinical diffusion scans for fiber tracking.

Figs. 5b–c show examples of the tracking results obtained using FACT-DT and surface-DT. The same data and seed ROI were used for each algorithm. The voxels identified to be connected to the seed are shown as a contiguous volume, rendered as a surface. There was a large amount of variability in tracking performance over seed ROI and over data acquisition. For tracts #1 and #2, the standard deviation of d over acquisition was $\sigma(d)$ = 0.15 ± 0.06 for FACT-DT (overall mean d of 0.38) and $\sigma(d)$ = 0.10 ± 0.09 for surface-DT (overall mean d of 0.17). The variability over seed point was comparable for both algorithms, with $\sigma(d)$ = 0.14 ± 0.05 for FACT-DT and $\sigma(d)$ = 0.09 ± 0.07 for surface-DT.

The noise in the diffusion tensor reconstruction depends on the number of base diffusion weighted images acquired, as well as other factors such as the directions themselves (Jones et al., 1999). We investigated the dependence of each algorithm's performance on the number of base diffusion weighted images by computing the average error measure for two, three, and four averages of the 90 direction diffusion weighting scheme, as well as a subset of 30 directions from one dataset. The results are shown in Fig. 6, where the total number of DWIs is the number of signal averages times the number of diffusion encoding directions.

Table 1
Summary of results of quantitative phantom studies comparing FACT-DT, surface-DT, surface-QB, and surface-hybrid

Algorithm	Number of DWIs	b value (s/mm^2)	Voxel size (mm^3)	Error measure d (dimensionless)
FACT-DT	90	1300	2.5^3	0.34 ± 0.03
Surface-DT	90	1300	2.5^3	0.18 ± 0.02
FACT-DT	4 × 90	1300	2.5^3	0.24 ± 0.05
Surface-DT	4 × 90	1300	2.5^3	0.21 ± 0.02
Surface-QB	4 × 90	3000	2.8^3	0.28 ± 0.07
Surface-DT	4 × 90	3000	2.8^3	0.23 ± 0.05
Surface-hybrid	4 × 90	3000	2.8^3	0.17 ± 0.05

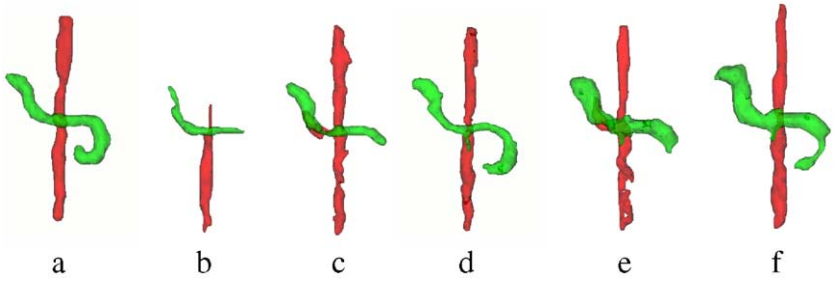

Fig. 5. Surface rendering of connected voxels: (a) Voxels identified as gold standard voxels: tract #1 (red) and tract #2 (green). Voxels reached using (b) FACT-DT and (c) surface-DT using the tensor computed from one signal average of 90 diffusion encoding directions with $b = 1300$ s/mm^2. (d) Results of surface-DT using four signal averages of the protocol used for (b) and (c), (e) surface-QB with the same four-average acquisition except using $b = 3000$ s/mm^2, and (f) surface-hybrid using the same data as in panel (e).

Surface evolution: DTI reconstruction compared to QBI reconstruction

The diffusion ODF calculated using q-ball reconstruction is shown in Figs. 4e and f. These data were used to compare surface-QB to surface-DT using the tensor ODF computed using an equivalent number of diffusion encoding directions and an optimized b value of $b = 1300$ s/mm^2. The surface-QB approach had poorer results (Table 1), with $d = 0.28 \pm 0.07$. To determine whether the difference was due to the acquisition parameters or to the ODF reconstruction, we did tracking with surface evolution using the *tensor* reconstruction of the $b = 3000$ s/mm^2 data. DT reconstruction of these data performed as well as DT reconstruction of the $b = 1300$ data s/mm^2 (Table 1).

Surface-hybrid tracking was done using only the $b = 3000$ s/mm^2 data. The data (Table 1) suggest a trend toward better performance with the hybrid approach, which has a lower mean d (0.17 ± 0.05) than either the q-ball or tensor reconstruction alone. Figs. 5d–f show the best tracking results obtained with surface-DT, surface-QB, and surface-hybrid using four signal averages of 90 diffusion encoding directions.

Human brain experiments

Fig. 7 shows the diffusion ODF calculated using the DTI and QBI reconstruction approaches in the cortical margin. This illustrates the additional information that can be gained from QBI reconstruction: note that in cases where two fibers cross, the tensor-derived ODF is planar, whereas the QBI ODF has clear maxima along the fiber directions. Where the crossing is not orthogonal, the maximum of the tensor-derived ODF does not lie along either direction.

Figs. 8–10 show fiber tracking results in the human brain for surface-DT, surface-hybrid, and FACT-DT. These figures illustrate tracking in major commissural fiber tracts (Figs. 8 and 9) and more

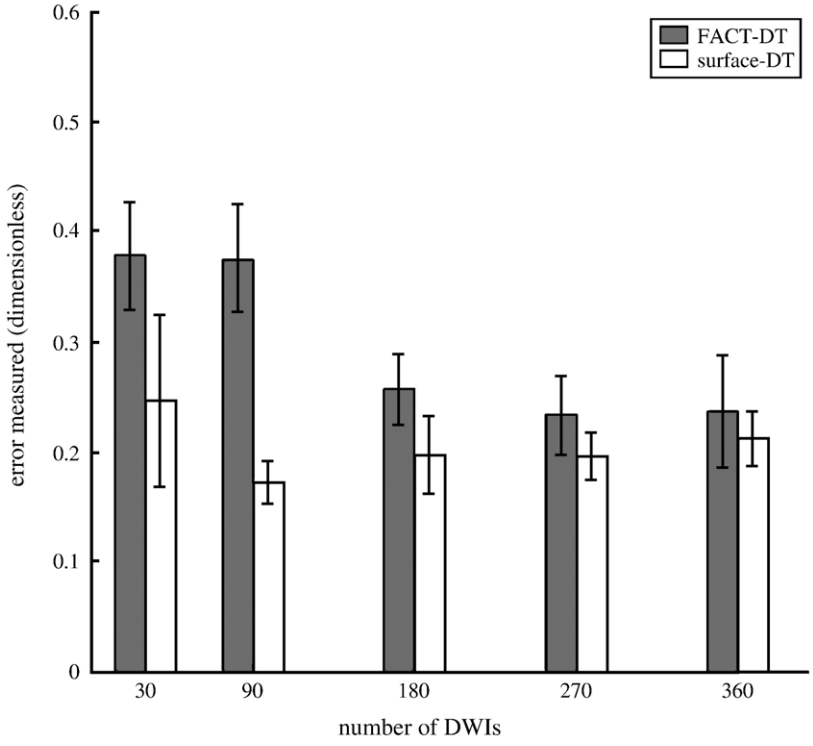

Fig. 6. Comparison of FACT-DT and surface-DT at different noise levels. The error measure d is shown for different numbers of base diffusion weighted images (DWIs) taken from the data acquired at $b = 1300$ s/mm^2.

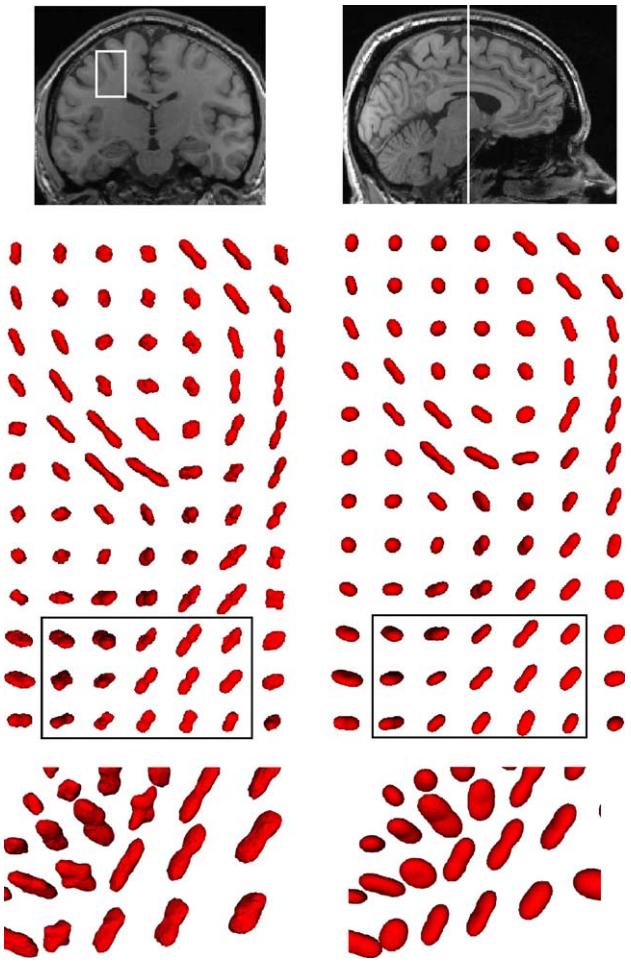

Fig. 7. Diffusion ODF computed using (left) QBI and (right) DTI reconstruction in an area of complex neuronal architecture in the cortical margin. The shapes have been squared and renormalized in order to accentuate anisotropy for visualization purposes, and the insets have been rotated slightly in order to show the directions clearly. Note that, in cases where two fibers cross, the tensor-derived ODF is planar, whereas the QBI ODF has clear maxima along the fiber directions. Where the crossing is not orthogonal, the maximum of the tensor-derived ODF does not lie along either direction.

subtle sub-cortical pathways (Fig. 10). All connectivity index maps are maximum intensity projections (MIPs) through the entire volume. The connections themselves, displayed as curves, are shown as frames from 3D display software. The view is looking straight at one of three orthogonal planes of the underlying anatomical image. The user may interactively adjust the connectivity index threshold determining which tracts are displayed. The tracts are 3D structures and are not restricted to one plane.

Discussion

We have shown that flow-based surface evolution can be used to perform fiber tracking using both DT and HARD data. In order to quantify the performance of this tractography, we have designed a physical phantom with known connectivity. We have shown that the flow-based approach, which allows for tract propagation in multiple directions, can improve fiber tracking results in the presence of partial volume averaging of fiber directions. A hybrid approach using QBI data in regions where the diffusion tensor fit is poor can further improve the tracking results.

In the case of partial volume averaging of fiber directions which we created in our phantom, many of the tracts propagated using the FACT-DT algorithm did not pass through the crossing region (see Fig. 5b), probably because the principal eigenvector direction did not lie along the tract direction. Additionally, the tracts that did pass through the crossing region stopped before reaching the ends of the cord because noise and/or imaging artifacts caused the principal eigenvector direction to be slightly deflected, hence leading the path out of the cord. Surface evolution was better able to overcome these issues: most of the errors in the tracking were due to bleeding down the perpendicular tract, despite the curvature constraint, as can be seen in Fig. 5c. The crossing region in this dataset spans many voxels, and reconstructed tracts can therefore turn slowly from one tract to another without exceeding the curvature constraint. Such scenarios can often occur with typical voxel sizes in vivo. With the tensor approach, the disc-like diffusion ODF makes it difficult to control the flow in these cases. However, the use of q-ball data in the crossing region improved the results, which suggests that the multi-peaked ODFs in this region (see Figs. 4e and f) caused the flow to be fast only along the tract directions, as opposed to flowing at comparable speeds in all directions in the plane of the crossing, as would occur with the tensor approach. In the QB case, if the flow bends down the wrong tract, it is more likely to be stopped by the curvature constraint.

A limitation of both the surface-QB and surface-hybrid approaches is that there is necessarily a tradeoff between angular resolution and spatial resolution. Here, the high b values necessary for QBI necessitated using larger voxels in the base diffusion weighted images. Some widening of the reconstructed tract structures because of the larger voxel size is evident, for example, in the reconstruction of the splenium of the corpus callosum in vivo shown in Fig. 8b. We note that tracking using q-ball data only, as done in the surface-QB approach, performed less well than did surface-DT using the same base diffusion weighted images, which performed comparably to surface-DT using base diffusion weighted images with lower b values. The poor performance of surface-QB was therefore not attributed to the lower SNR of the base DWIs, the larger voxel size, or the increased eddy current induced artifacts due to higher slew rates, but rather to the QBI reconstruction itself. Using the q-ball ODF when the tensor fit is good only increases susceptibility to noise, however, when the tensor fit is poor, the high angular resolution ODF obtained with QBI can help control the flow and assign more reasonable connectivity indices. We note that the larger voxel size used for the QBI acquisition might confound reconstruction of even smaller tracts but was sufficient for the tracts created in this phantom.

Imaging noise can cause tracking to stop or proceed down the wrong path in either tracking approach. The results suggest that surface evolution using the full tensor is more robust than line propagation in the presence of noise, as evidenced by its superior performance with low numbers of base DWIs (Fig. 6). Hence, when tracking using datasets acquired with low SNR, either because of small voxel sizes, short acquisition times, or significant partial volume averaging of fibers with isotropic material, surface evolution should be preferred. However, the two algorithms plateau at the same performance level at higher SNR, meaning the advantages and disadvantages are balanced for the two algorithms at this point. The choice of which algorithm to use

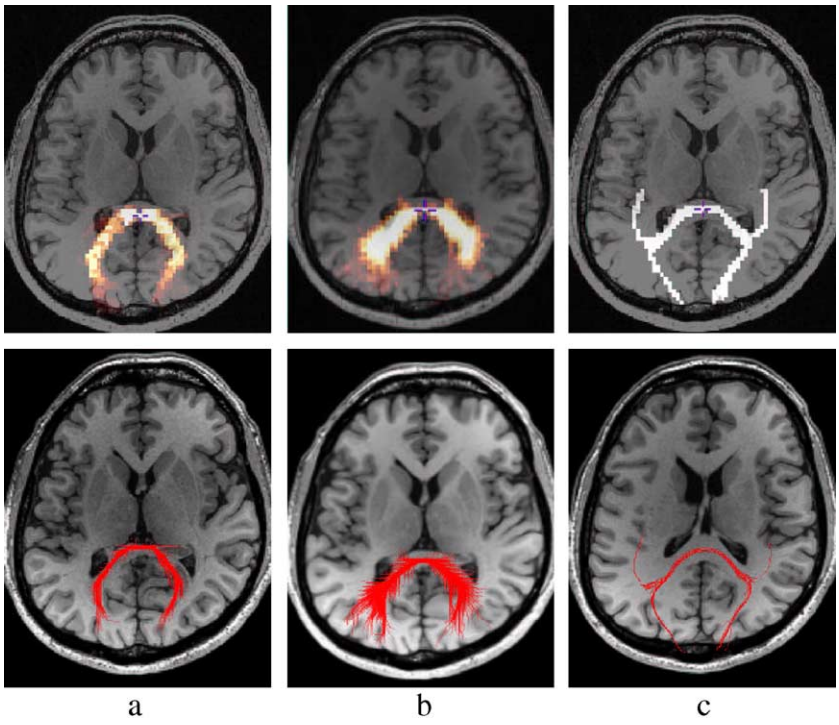

Fig. 8. Tracking results: Top: Maximum intensity projection (MIP) of connectivity index map. Bottom: User-thresholded tracts projected on a 2D plane. From left to right: surface-DT, surface-hybrid with threshold of $Q(\chi^2) = 0.001$, and FACT-DT: the connectivity index map in the case of FACT-DT is binary. The single-voxel seed point (crosshair) was placed in the splenium of the corpus callosum. For surface evolution, the CI maps are windowed to show the most likely connections: there are non-zero CI values elsewhere.

depends on the type of fiber structure expected in any fiber tracking application: for example, if many subvoxel crossings are expected, using the full tensor should be superior. Additionally, it may be possible to reduce the noise sensitivity of line propagation by using regularization techniques (McGraw et al., 2004).

In the human brain experiments, FACT-DT showed a tendency to jump incorrectly from one tract system to another: in Fig. 8c, FACT-DT reconstructs a false connection to the optic radiation, which originates in V1, while in Fig. 10c, the FACT-DT tracking was confounded by the large fibers of the pyramidal tract and turned incorrectly down through the internal capsule. Fig. 10d shows a drawing of the known human anatomy in this region: the corpus callosum (A) and the internal capsule (B) are indicated. It can be seen that the fibers from the corpus callosum and pyramidal tract, although part of distinct pathways, can be expected to cross on a subvoxel scale in the region where FACT-DT fails. This causes a slow curvature to be implied by the PDD, whereas QBI shows multiple fiber directions and the full diffusion tensor is flat. Hence, voxels reached by the incorrect connection had very low connectivity indices in both the surface-hybrid and surface-DT approaches: this is a case where assigning a non-binary connectivity index can help identify uncertain pathways that are reconstructed due to partial volume averaging of fiber directions or to noise. We note that we could assign the same connectivity index to the tracts reconstructed with FACT-DT propagation: we propagate only along the principal diffusion direction (the maximum of the tensor-described diffusion ODF Ψ), note the magnitude of Ψ in this direction, and define the minimum along the tract of this value to be the tract connectivity index. While this approach could help narrow down which of the tracts that *are* reconstructed have high likelihood of existence, it would fail whenever a connection is *not* reconstructed with the PDD approach, for example, in the case of 'flat' tensors, where the principal diffusion direction may not coincide with the tract direction. Using the full tensor, as in the surface-DT approach, can deal with these problems to an extent, however, in cases where the maximum of the tensor-described ODF is incorrect (e.g., the nonorthogonal crossings shown in Fig. 7), the use of q-ball data should be more accurate.

The scalar connectivity index assigned in the surface evolution scheme is useful for assessing our certainty in a reconstructed tract's existence. However, interpreting the connectivity indices in terms of which tracts do and do not exist is left to the user and requires a certain amount of a priori knowledge. For example, we consider the tracking results shown in Fig. 10, where the seed point was placed close to the cortex. This is an area of significant partial volume averaging of fiber directions. Surface-DT and surface-hybrid give slightly different results, whereas in Figs. 8 and 9, the results of these two algorithms are very similar. In Fig. 10, surface-hybrid reconstructs a U-fiber, which is also reconstructed with FACT-DT, with high likelihood, and also reconstructs other lower likelihood association pathways as well as connections to the contralateral hemisphere via the corpus callosum. The surface-DT reconstruction indicates a connection to the descending cortical spinal tract and does not reconstruct the U-fiber as well. We suggest that the surface-hybrid results are more accurate and that it is probable that the seed voxel contains fibers from all of the tracts reconstructed by this method. The tensor model was not as capable of handling the partial volume averaging of directions as was the q-ball approach. However, we emphasize that it is up to the user to choose whether this is the case or whether the connections shown are the result of the tract passing through a voxel containing a "bottleneck" of fibers, including fibers from other fascicles that do

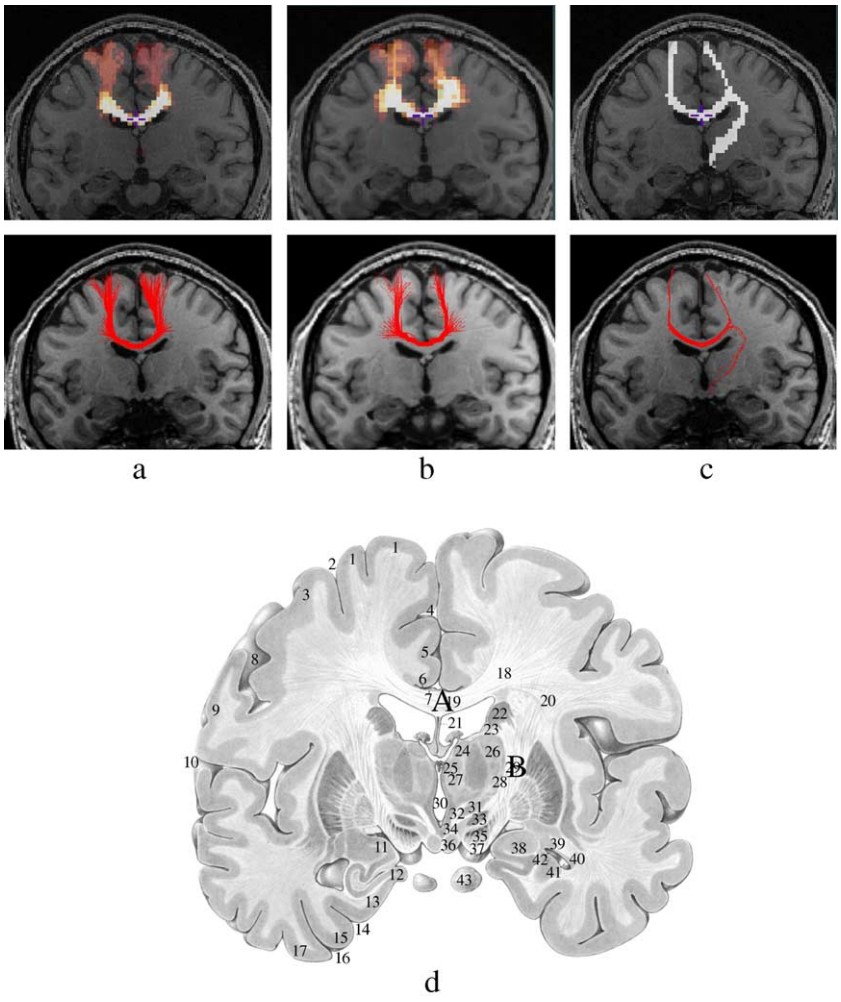

Fig. 9. Tracking results for (left to right): surface-DT, surface-hybrid, and FACT-DT for a single-voxel seed point (crosshair) in the corpus callosum. (d) Drawing of the anatomy in this region from Nieuwenhuys atlas, showing corpus callosum (A) and internal capsule (B) (figure adapted from Nieuwenhuys et al., 1988).

not in fact pass through the seed: jumping from tract to tract can occur in this way.

This study showed a high variability in the tracking performance over seed ROI in the phantom. It is fair to suggest that, if we start tracking iteratively from all voxels in the volume (the "brute force" approach) and ask from which starting voxels the seed ROI is reached and with what likelihood, we will get a different answer from the outward propagation approach. However, the

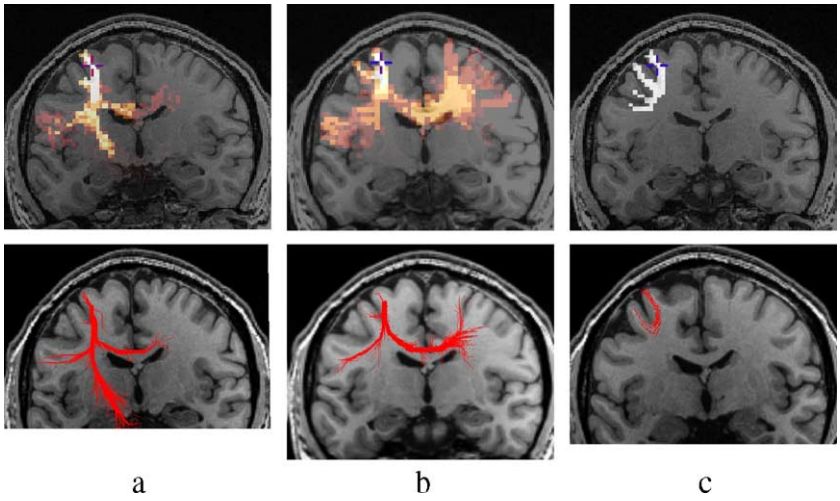

Fig. 10. Tracking results for (left to right): surface-DT, surface-hybrid, and FACT-DT for a single-voxel seed point (crosshair) near the cortex.

tracking performance when starting from the seed ROI only should correlate with the performance when tracking is initiated everywhere and takes significantly less computation time. This study was meant to give results from the forward tracking approach only and should be indicative of how reliable other approaches would be.

In summary, our findings suggest that using the full diffusion tensor for fiber tracking can provide superior results to approaches using only the principal direction of diffusion. Furthermore, HARD reconstruction of the diffusion ODF, illustrated here using the q-ball reconstruction method, can further improve the results when used in voxels in which the tensor model fits the data poorly. We have done extensive validation with a gold standard tract configuration. Tracking with all algorithms is confounded by noise, partial volume averaging of fiber directions, large voxels in the base diffusion weighted images, and limitations of the diffusion pdf measurement. Of the methods considered here, the surface-hybrid approach appears most robust to these confounds.

Acknowledgments

This work was supported by grants from NSERC, FQRNT, and CFI.

References

Alexander, D., Jansons, K., 2002. Spin echo attenuation to diffusion density: a general inversion for measurements on a sphere. International Society for Magnetic Resonance in Medicine Workshop on Diffusion MRI: Biophysical Issues (What can we measure?), p. 206–209.

Alexander, D., Barker, G., Arridge, S., 2002. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. Magn. Reson. Med. 48, 331–340.

Basser, P., Mattiello, J., Le Bihan, D., 1994. Estimation of the effective self-diffusion tensor from the NMR spin echo. J. Magn. Reson. 103, 247–254.

Basser, P., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. Magn. Reson. Med. 44, 625–632.

Batchelor, P., Hill, D., Calamante, F., Atkinson, D., 2001. Study of connectivity in the brain using the full diffusion tensor from MRI. In: Insana, M., Leahy, R. (Eds.), IPMI 2001, Lect. Notes Comput. Sci., vol. 2082, p. 121–133.

Behrens, T.E., Woolrich, M.W., Jenkinson, M., Johansen-Berg, H., Nunes, R.G., Clare, S., Matthews, P.M., Brady, J.M., Smith, S.M., 2003. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magn. Reson. Med. 50 (5), 1077–1088 (Nov.).

Björnemo, M., Brun, A., Kikinis, R., Westin, C.-F., 2002. Regularized stochastic white matter tractography using diffusion tensor MRI. In: Dohi, T., Kikinis, R. (Eds.), MICCAI 2002, Lect. Notes Comput. Sci., vol. 2488, p. 435–442.

Campbell, J., Siddiqi, K., Pike, G., 2002a. White matter fibre tract likelihood evaluated using normalized RMS diffusion distance. Proceedings of the International Society for Magnetic Resonance in Medicine: 10th Scientific Meeting and Exhibition, p. 1130.

Campbell, J., Siddiqi, K., Vemuri, B., Pike, G., 2002b. A geometric flow for white matter fibre tract reconstruction. In: Mercer, B. (Ed.), 2002 IEEE International Symposium on Biomedical Imaging Conference Proceedings. The Institute of Electrical and Electronics Engineers, Inc., Omni Press, p. 505–508.

Cointepas, Y., Poupon, C., Le Bihan, D., Mangin, J.-F., 2002. A spin glass framework to untangle fiber crossing in MR diffusion based tracking. In: Dohi, T., Kikinis, R. (Eds.), MICCAI 2002, Lect. Notes Comput. Sci., vol. 2488, p. 475–482.

Collins, D., Neelin, P., Peters, T., Evans, A., 1994. Automatic 3D inter-subject registration of MR volumetric data in standardized Talairach space. J. Comput. Assist. Tomogr. 18 (2), 192–205.

Conturo, T.E., Lori, N.F., Cull, T.S., Akbudak, E., Snyder, A., Shimony, L.S., McKinstry, R.C., Burton, H., Raichle, M.E., 1999. Tracking neuronal fibre pathways in the living human brain. Proc. Natl. Acad. Sci., (Neurobiol., Appl. Phys. Sci.) 96, 10422–10427.

Frank, L.R., 2002. Characterization of anisotropy in high angular resolution diffusion-weighted MRI. Magn. Reson. Med. 47 (6), 1083–1099 (Jun.).

Hagmann, P., Thiran, J., Jonasson, L., Vandergheynst, P., Clarke, S., Maeder, P., Meuli, R., 2003. DTI mapping of human brain connectivity: statistical fibre tracking and virtual dissection. NeuroImage 19, 545–554.

Henkelman, R.M., 1985. Measurement of signal intensities in the presence of noise in MR images. Med. Phys. 12 (2), 232–233 (Mar.–Apr.).

Jones, D., 2003. Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI. Magn. Reson. Med. 49 (1), 7–12.

Jones, D., Horsfield, M., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magn. Reson. Med. 42, 515–525.

Jones, D.K., Williams, S.C., Gasston, D., Horsfield, M.A., Simmons, A., Howard, R., 2002. Isotropic resolution diffusion tensor imaging with whole brain acquisition in a clinically acceptable time. Hum. Brain Mapp. 15 (4), 216–230.

Koch, M.A., Norris, D.G., Hund-Georgiadis, M., 2002. An investigation of functional and anatomical connectivity using magnetic resonance imaging. NeuroImage 16, 241–250.

Lazar, M., Alexander, A.L., 2003. An error analysis of white matter tractography methods: synthetic diffusion tensor field simulations. NeuroImage 20 (2), 1140–1153.

Lazar, M., Weinstein, D.M., Tsuruda, J.S., Hasan, K.M., Arfanakis, K., Meyerand, M.E., Badie, B., Rowley, H.A., Haughton, V., Field, A., Alexander, A.L., 2003. White matter tractography using diffusion tensor deflection. Hum. Brain Mapp. 18 (4), 306–321.

Lin, C., Wedeen, V., Yao, C., Chen, J., Tseng, W., 2002. Validation of diffusion spectrum magnetic resonance imaging with registered manganese-enhanced optic tracts and phantom. Proceedings of the International Society for Magnetic Resonance in Medicine: 10th Scientific Meeting and Exhibition, p. 442.

Lori, N.F., Akbudak, E., Shimony, J.S., Cull, T.S., Snyder, A.Z., Guillory, R.K., Conturo, T.E., 2002. Diffusion tensor fiber tracking of human brain connectivity: acquisition methods, reliability analysis and biological results. NMR Biomed. 15 (7–8), 494–515.

McGraw, T., Vemuri, B.C., Chen, Y., Rao, M., Mareci, T., 2004. DT-MRI denoising and neuronal fiber tracking. Med. Image Anal. 8 (2), 95–111.

Mori, S., Crain, B., Chacko, V., van Zijl, P., 1999. Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann. Neurol. 45, 265–269.

Mori, S., Davatzikos, C., Xu, D., Amodei, L., Freericksen, K., Solaiyappan, M., Van Zil, P., 2001. A probabilistic map for the analysis of DTI-based fiber tracking and its application to white matter injuries. Proceedings of the International Society for Magnetic Resonance in Medicine: 9th Scientific Meeting and Exhibition, p. 1522.

Mori, S., Kaufmann, W.E., Davatzikos, C., Stieltjes, B., Amodei, L., Fredericksen, K., Pearlson, G., Melhem, E.R., Raymond, G., Moser, H.W., van Zijl, P., 2002. Imaging cortical association tracts in the human brain using diffusion-tensor-based axonal tracking. Magn. Reson. Med. 47 (2), 215–223.

Nieuwenhuys, R., Voogd, J., van Huijzen, C., 1988. The Human Central Nervous System: A Synopsis and Atlas, 3rd ed. Springer Verlag, New York.

O'Donnell, L., Haker, S., Westin, C.-F., 2002. New approaches to estimation of white matter connectivity in diffusion tensor MRI: elliptic PDEs and geodesics in a tensor-warped space. In: Dohi, T., Kikinis, R. (Eds.), MICCAI 2002, Lect. Notes Comput. Sci., vol. 2488, p. 459–466.

Osher, S.J., Sethian, J.A., 1988. Fronts propagating with curvature dependent speed: algorithms based on HamiltonJacobi formulations. J. Comput. Phys. 79, 12–49.

Parker, G., Wheeler-Kingshott, C., Barker, G., 2002a. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. IEEE Trans. Med. Imaging 21 (5), 505–512.

Parker, G.J., Stephan, K.E., Barker, G.J., Rowe, J.B., MacManus, D.G., Wheeler-Kingshott, C.A., Ciccarelli, O., Passingham, R.E., Spinks, R.L., Lemon, R.N., Turner, R., 2002b. Initial demonstration of in vivo tracing of axonal projections in the macaque brain and comparison with the human brain using diffusion tensor imaging and fast marching tractography. NeuroImage 15 (4), 797–809 (Apr.).

Pierpaoli, C., Basser, P., 1996. Toward a quantitative assessment of diffusion anisotropy. Magn. Reson. Med. 36, 893–906.

Press, W.H., Flannery, B., Teukolsky, S.A., Vetterling, W., 1992. Numerical Recipes in C, 2nd ed. Campbridge University Press.

Ramírez-Manzanares, A., Rivera, M., 2003. Brain nerve bundles estimation by restoring and filtering intra-voxel information in diffusion tensor MRI. Proceedings of the 2nd IEEEWorkshop on Variational, Geometric and Level Set Methods in Computer Vision.

Reese, T.G., Heid, O., Weisskoff, R.M., Wedeen, V.J., 2003. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. Magn. Reson. Med. 49 (1), 177–182.

Shrager, R., Jones, D., Pajevic, S., Munson, P., Basser, P., 2002. When is a Gaussian displacement distribution adequate to describe water diffusion in tissues? International Society for Magnetic Resonance in Medicine Workshop on Diffusion MRI: Biophysical Issues (What can we measure?), p. 21–25.

Stieltjes, B., Kaufmann, W., van Zijl, P., Fredericksen, K., Pearlson, G., Solaiyappan, M., Mori, S., 2001. Diffusion tensor imaging and axonal tracking in the human brainstem. NeuroImage, 723–735.

Tournier, J.D., Calamante, F., King, M.D., Gadian, D.G., Connelly, A., 2002. Limitations and requirements of diffusion tensor fiber tracking: an assessment using simulations. Magn. Reson. Med. 47 (4), 701–708.

Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A., 2003. Diffusion-weighted magnetic resonance imaging fibre tracking using a front evolution algorithm. NeuroImage 20 (1), 276–288 (Sep.).

Tournier, J., Calamante, F., Gadian, D., Connelly, A., 2004. Direct estimation of fibre orientations in partial volume contaminated regions using spherical deconvolution. Proceedings of the International Society for Magnetic Resonance in Medicine: 13th Scientific Meeting and Exhibition, p. 88.

Tuch, D., Reese, T., Wiegell, M., Makris, N., Belliveau, J., Wedeen, V., 2002. High angular resolution diffusion imaging reveals intra-voxel white matter fiber heterogeneity. Magn. Reson. Med. 48 (4), 577–582.

Tuch, D.S., Reese, T.G., Wiegell, M.R., Wedeen, V.J., 2003. Diffusion MRI of complex neural architecture. Neuron 40 (5), 885–895.

Vemuri, B., Chen, Y., Rao, M., Wang, Z., Mareci, T., Blackband, S., Reier, P., 2002. Automatic fiber tractography from DTI and its validation. In: Mercer, B. (Ed.), IEEE International Symposium on Biomedical Imaging Conference Proceedings. The Institute of Electrical and Electronis Engineers, Inc., Omni Press, pp. 505–508.

von dem Hagen, E., Henkelman, R., 2001. Orientational diffusion reflects fiber structure within a voxel. Proceedings of the International Society for Magnetic Resonance in Medicine: 9th Scientific Meeting and Exhibition, p. 1528.

Wakana, S., Jiang, H., Nagae-Poetscher, L.M., Van Zijl, P.C., Mori, S., 2003. Fiber tract-based atlas of human white matter anatomy. Radiology, 77–87 (Nov. 26).

Wedeen, V., Reese, T., Tuch, D., Weigel, M., Dou, J.-G., Weiskoff, R., D., C., 2000. Mapping fibre orientation spectra in cerebral white matter with Fourier-transform diffusion MRI. Proceedings of the International Society for Magnetic Resonance in Medicine: 8th Scientific Meeting and Exhibition, p. 82.

Xue, R., van Zilj, P., Crain, B., Solaiyappan, M., Mori, S., 1999. In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. Magn. Reson. Med. 42, 1123–1127.

Zhan, W., Gu, H., Xu, S., Silbersweig, D.A., Stern, E., Yang, Y., 2003. Circular spectrum mapping for intravoxel fiber structures based on high angular resolution apparent diffusion coefficients. Magn. Reson. Med. 49 (6), 1077–1088.