NeuroImage

www.elsevier.com/locate/ynimg
NeuroImage 41 (2008) 58–68

Labeling of ambiguous subvoxel fibre bundle configurations in high angular resolution diffusion MRI ☆

Peter Savadjiev,[a,*] Jennifer S.W. Campbell,[a,b] Maxime Descoteaux,[c] Rachid Deriche,[c] G. Bruce Pike,[b] and Kaleem Siddiqi[a]

[a]*Centre for Intelligent Machines and School of Computer Science, McGill University, Montréal, Canada*
[b]*McConnell Brain Imaging Centre, Montréal Neurological Institute, McGill University, Montréal, Canada*
[c]*Project Team Odyssée, INRIA Sophia-Antipolis, France*

Received 11 October 2007; revised 26 December 2007; accepted 3 January 2008
Available online 1 February 2008

Whereas high angular resolution reconstruction methods for diffusion MRI can estimate multiple dominant fibre orientations within a single imaging voxel, they are fundamentally limited in certain cases of complex subvoxel fibre structures, resulting in ambiguous local orientation distribution functions. In this article we address the important problem of disambiguating such complex subvoxel fibre tract configurations, with the purpose of improving the performance of fibre tractography. We do so by extending a curve inference method to distinguish between the cases of curving and fanning fibre bundles using differential geometric estimates in a local neighbourhood. The key benefit of this method is the inference of curves, instead of only fibre orientations, to model the underlying fibre bundles. This in turn allows distinct fibre geometries that contain nearly identical sets of fibre orientations at a voxel, to be distinguished from one another. Experimental results demonstrate the ability of the method to successfully label voxels into one of the above categories and improve the performance of a fibre-tracking algorithm.

Introduction

Fibre tractography in diffusion MRI is concerned with the problem of reconstructing the shape and geometry of possible white matter pathways in the brain. Diffusion-weighted MRI data can be used to define a probability density function (pdf) for the local Brownian motion of water molecules at each location in the brain (see, e.g., Basser et al., 1994; Le Bihan et al., 2001). Since the fibrous structure of the brain restricts this motion, the resulting diffusion is maximal along the orientation(s) of the underlying fibre tract(s) (Lin et al., 2003) and can thus be used to infer the presence and geometry of underlying fibre populations.

Several algorithms exist in the literature for performing fibre tracking in diffusion MRI data (e.g., Mori et al., 1999; Basser et al., 2000; Parker and Alexander, 2005; Campbell et al., 2005; Friman et al., 2006; Behrens et al., 2007). The earlier tracking methods proceed by following the principal direction of diffusion in each voxel in diffusion tensor (DT) data (see, e.g., Mori et al., 1999). The DT framework (Basser et al., 1994) imposes a 3D Gaussian model of the diffusion pdf. The advantage of such a representation is a typically fast acquisition; however, it is limited in voxels that reflect more than one fibre population. Thus, tracking algorithms that follow the principal direction of diffusion in DT data are likely to fail in cases where there are multiple-fibre directions within a single image voxel (Behrens et al., 2007). In such voxels, DT reconstructions suffer from partial volume averaging effects, and the resulting principal direction of diffusion may provide ambiguous and/or insufficient information about the orientation of the underlying fibre bundle(s). Such effects can occur in voxels where the underlying fibre tracts cross, fan, or curve.

Since the diffusion tensor is a low angular resolution approximation to the diffusion pdf, high angular resolution reconstruction methods have been recently developed, including q-ball (QB) imaging (Tuch, 2004), diffusion spectrum imaging (DSI) (Callaghan, 1993; Wedeen et al., 2000; Tuch, 2002), composite hindered and restricted models of diffusion (CHARMED) (Assaf and Basser, 2005), and the diffusion orientation transform (DOT) (Özarslan et al., 2006). Such methods can be used to estimate the diffusion orientation distribution function (ODF) of water molecules, which is the projection in the radial direction of the diffusion pdf on the surface of the sphere.

High angular resolution diffusion (HARD) MRI presents numerous advantages over the tensor model and can be used to achieve better precision in fibre tractography (Campbell et al., 2005). However, the relationship between the diffusion pdf and tissue structure remains ambiguous (Tuch, 2002). Furthermore, certain subvoxel fibre configurations can result in diffusion ODF shapes that are difficult

☆ A preliminary abstract based on this work appears in the proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM) 2007 conference.

* Corresponding author.
E-mail address: petersv@cim.mcgill.ca (P. Savadjiev).
Available online on ScienceDirect (www.sciencedirect.com).

1053-8119/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.neuroimage.2008.01.028

to interpret during tractography. Ideally, tracking should be performed on fibre ODF data—that is, a function that describes the local distribution of fibre bundles, i.e., their tangent orientations. The diffusion pdf is a blurred version of the fibre ODF, representing (at least in the simple case of one fibre population per voxel) a convolution of the 'true' fibre ODF with a single-fibre response function (e.g., Tournier et al., 2004). Because of this blurring effect, diffusion ODF maxima are often extracted and used as estimates of dominant fibre orientations (Lin et al., 2003; Seunarine et al., 2006). As an alternative to the use of diffusion ODF maxima to approximate fibre ODFs, techniques such as spherical deconvolution (Tournier et al., 2004, 2007; Anderson, 2005; Sakaie and Lowe, 2007; Dell'Acqua et al., 2007; Kaden et al., 2007; Jian and Vemuri, 2007; Descoteaux et al., 2007a,b; Alexander, 2005) have been introduced to extract fibre orientation distribution functions (fibre ODFs) by deconvolving the diffusion-weighted signal or the diffusion ODF with a single-fibre response function. The persistent angular structure (PAS) (Jansons and Alexander, 2003), which has been shown to be a special case of deconvolution, has also been proposed. Despite numerous advantages, fibre ODFs like diffusion ODFs can still be ambiguous and difficult to interpret in the presence of complex subvoxel fibre tract configurations (Le Bihan et al., 2006; Parker and Alexander, 2005) and thus confound fibre-tracking algorithms. Different fibre geometries can yield similar ODFs but require different actions to be taken in tractography.

To illustrate the fact that different fibre configurations, which require different actions to be taken in tractography, can still yield similar ODFs, consider the two types of subvoxel fibre structures schematically depicted in Fig. 1. In this example, both the single curving fibre tract (top row) and the fanning fibre tract (bottom row) are likely to result in an almost undistinguishable ODF with a single broad peak oriented in the vertical direction (middle panel in both rows). This is due to the presence of a relatively wide array of fibre tangent directions within one voxel, which results in a broad fibre ODF profile. Although they yield a similar ODF, each of these structures requires a different action to be taken by a fibre-tracking algorithm. In the case of a curving fibre bundle, only one curve (or *streamline*) should be propagated, in the medial direction of the broad maximum. In the case of a fan configuration, multiple curves should be followed when propagating in one direction (represented by a polarity vector), but only one direction should be followed when propagating in the other. Hence, in order to take the appropriate action, it is crucial for tracking methods to be able to differentiate between these types of subvoxel configurations.

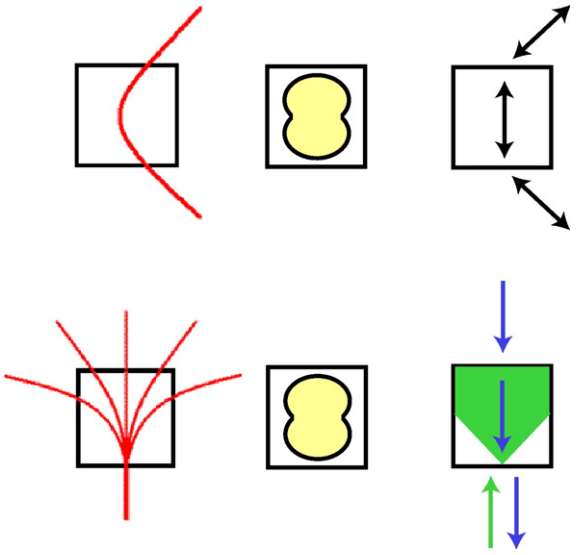

Fig. 1. A schematic illustration of subvoxel fibre configurations that can cause ambiguous fibre ODF shapes and thus confound fibre-tracking algorithms. Both a single curving fibre tract (top row) and a fanning fibre tract (bottom row) can yield a similar ODF with a single broad peak (middle column in each row). In the case of the curving fibre bundle, regardless of the incoming direction, tracking should follow the medial direction in the peak (top row, right). In the case of the fanning fibre bundle, however, when tracking enters the voxel in the fan (polarity) direction (green arrow, bottom row, right panel), it should follow all directions within the extent of the fan (green cone within the voxel). If tracking enters the voxel from the opposite side of the fan (blue arrow), it should only follow the medial direction of the peak. This illustrates the importance of recovering the polarity of the fan in addition to its extent.

The distinction between a fanning fibre tract and a single-fibre tract with high curvature has been known to be particularly challenging (Le Bihan et al., 2006; Parker and Alexander, 2005), and to our knowledge, no previous work has attempted to explicitly disambiguate such structures. Methods such as those of Seunarine et al. (2006), Behrens et al. (2007), and Ramírez-Manzanares and Rivera (2006) attempt in original ways to model multiple-fibre orientations in voxels where they may exist; however, none of these achieves the categorisations treated in this paper.

In this article we relate fibre ODF data, which describe fibre tangents, to the underlying white matter fibre tracts, modeled as 3D curves. We apply the 3D curve inference algorithm, described in Savadjiev et al. (2006), in a labeling scheme that disambiguates fibre ODFs that result from the following possible configurations:

(1) Single-fibre bundles—these may exhibit curvature and torsion at a subvoxel scale.
(2) Fanning fibre tracts—where a single fibre bundle fans into multiple directions at a subvoxel scale.
(3) Crossing fibre tracts—where two or more (possibly curving) fibre tracts cross within a voxel.

The distinction between the first two categories is the main contribution of this work. Cross configurations typically present a multi-peaked ODF, as opposed to the first two categories that often both present a single, broad peak. This fact suggests a relatively straightforward strategy to distinguish them from the other two cases, and thus the labeling of crossings is included here only for completeness.

The labeling scheme we introduce allows these cases to be treated properly in tractography and thus reduces the occurrence of false positive and false negative connections (Campbell et al., 2005). White matter fibres are known to splay from large, tight bundles to the cortex, and they may need to part, for instance, in order to project to multiple gyri. This happens on a scale such that, with a typical imaging voxel size of 2 mm × 2 mm × 2 mm, there can be a continuous range of fibre directions within one voxel, and tractography can be expected to perform better if these directions are all accounted for. For example, the identification of fibre fanning helps to reconstruct the connections that pass through the corpus callosum and project laterally towards the premotor cortex (as shown in Results section), a situation that currently confounds algorithms that use only the ODF maxima. In addition to the identification of such voxels, 3D curve inference also infers the polarity and the extent of the subvoxel fan. This is important since the extent of the fan determines the extent of space in which tracking should propagate, and the polarity determines the appropriate action to be taken depending on whether the incoming direction is the one in which fanning occurs (in which case all directions within the fan should be followed) or the merging one (in

which case only the medial direction of the fan should be followed). See Fig. 1 for an illustration. We provide labeling results and the recovery of fan polarity on *in vivo* human brain data, as well as results that demonstrate improvements in the performance of an algorithm for fibre tracking in the human brain.

We begin with a brief overview of the 3D curve inference algorithm, which infers differential geometric information over a local neighbourhood and thus allows for the resolution of ambiguities that cannot be solved by considering only individual ODF shapes (Parker and Alexander, 2005). We refer the reader to Savadjiev et al. (2006) for its technical details.

3D curve inference

3D curve inference is a differential geometric framework for the estimation of trace, tangent, curvature and torsion fields of 3D space curves from orientation data defined at locations on a 3D discrete lattice. The algorithm infers helical arcs as osculating approximations to arbitrary 3D curves, based on a support measure defined over a local neighbourhood. In the context of diffusion MRI orientation data, this allows for the local curvature and torsion of white matter fibres to be estimated. In this fashion distinct subvoxel fibre bundle configurations that share the same tangents (orientations) at a particular voxel can be distinguished from one another.

As input the algorithm requires a discretised regular (typically rectangular) 3D lattice, with a fibre ODF defined at each location (voxel) in the lattice. Each of these ODFs is then sampled along several orientations. A notion of *co-helicity*, formally defined in Savadjiev et al. (2006), then relates individual orientations defined at distinct voxels through a geometrical constraint. In particular, the conditions under which three orientations defined at three distinct locations in space can be tangent to a helix are determined. Based on this idea, a measure of the support that a given orientation at a given voxel receives from co-helical configurations of neighbouring orientations at neighbouring voxels, weighted by the ODF value along these orientations at these voxels, is calculated. This (local) support function is then incorporated in a global measure of *average local support* which is defined over the entire volume and is maximised (to a local maximum) using an iterative relaxation labeling algorithm (Hummel and Zucker, 1983; Parent and Zucker, 1989). This has the effect of regularising (reshaping) the ODF shapes, while allowing for multiple distinct fibre geometries to be inferred.

The aspect of the 3D curve inference algorithm which is particularly relevant to the current article is that during the regularisation process, each orientation is associated with a *curvature–torsion–normal class* (Savadjiev et al., 2006) which provides a discretisation of the parameter space describing the best fitting helices given the neighbourhood support information. It is these best-fit helical curves that are used in the labeling algorithm developed in the following section. The regularization itself is not the focus of the current work. As described in the Experimental methods section, the experiments presented in this article are carried out on the original (unregularised) fibre ODF data.

Labeling of ambiguous subvoxel fibre tract configurations

We now introduce a labeling scheme based on the analysis of local configurations of helical curve approximations, inferred through the 3D curve inference algorithm in fibre ODF data, to address the problem of disambiguating the cases of single and fanning fibre tracts. To accomplish this, we develop a model that links fibre tangents (a fibre ODF) to the underlying fibre curves.

To motivate the approach, consider Fig. 2, which shows a schematic of the inferred curves in the case of a fanning fibre tract and of a single curving fibre tract. For simplicity, only a 2D case is illustrated, but as mentioned earlier, the technique is applicable to any 3D ODF dataset. In the general 3D case, the inferred curves will be helical, i.e., they will have both curvature and torsion.

Local helix approximations to fibre bundles are constructed by searching for *co-helical* triplets of fibre ODF directions (see Savadjiev et al., 2006) in a local neighbourhood, so that they agree with the curvature–torsion–normal classes inferred through 3D curve inference. A co-helical triplet is interpolated by a helix which is used as a local approximation to the (arbitrary) 3D curve that represents an underlying white matter fibre tract. Thus, a given fibre ODF direction presents evidence for an underlying curve (helix) if it is the central element of a co-helical triplet in a spherical neighbourhood centred on that voxel and its curvature-torsion- normal class inferred through 3D curve inference agrees with the parameters (curvature, torsion) of the cohelical triplet. As an example, the three sampling direction vectors corresponding to the three blue maxima in Fig. 2 (right) form a co-helical triplet of directions. Similarly, in Fig. 2 (left), the groups of red, green and blue maxima all form co-helical triplets with the black maximum at the bottom, which is common to all three groups.

One or more such helices can pass through a given voxel, one per each ODF direction in that voxel. For example, three such helices pass through the central ODF in Fig. 2 (left), associated to the red, green and blue directions, respectively. The number and the configuration of these helices are used to label the voxel as belonging to a fan, cross, or single-fibre tract configuration, as described in Table 1. This labeling uses two types of information: (i) ODF shape information *local* to a voxel. Here the number of local maxima, or "peaks", can be used to distinguish crossings from the other two cases. (ii) A geometric model inferred from a *neighbourhood* of voxels. The helices inferred by the

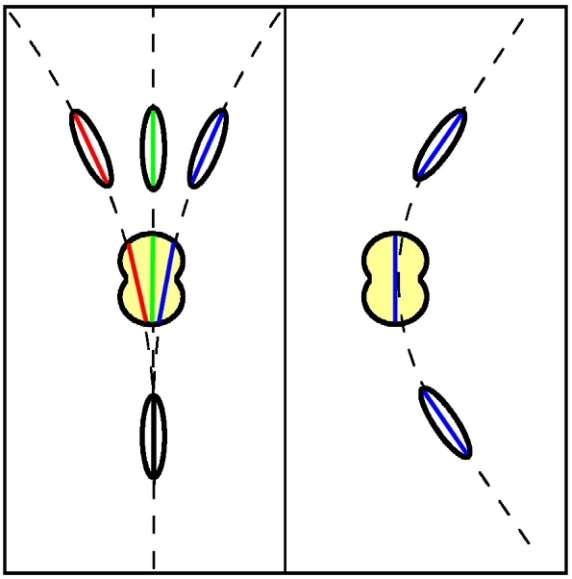

Fig. 2. A schematic illustration of the notion of co-helical triplets inferred through the 3D curve inference algorithm. In the fan case (left), multiple ODFs on the fanning side of the configuration can have their maxima (red, green, blue) cohelical with different directions in the central (ambiguous) ODF and with the same maximum (black) of an ODF on the merging side of the configuration. The corresponding inferred curves are shown with dashed lines. In the single curving fibre tract case (right), only one set of co-helical triplets of ODF maxima (blue) exists, and thus only one curve is inferred.

Table 1
The three classes of labels determined as a function of the number of local maxima, or "peaks," in the ODF, as well as the number and configuration of their corresponding helices

Label	Number of fibre ODF maxima	Configuration of curves in voxel
Single-fibre tract	One	One helix
Fan	One	Two or more distinct and diverging helices
Cross	Two or more	A number of helices equal to the number of local maxima such that they are associated to fibre ODF directions that have an angular separation larger than some threshold

The minimum angular separation threshold for the case of the "cross" label ensures that the multiple helices are not approximations to the same underlying fibre bundle but rather to distinct crossing fibre bundles.

3D curve inference algorithm are used to distinguish fans from single, possibly curving, fibre tracts. We emphasise that the inferred helical curves are local approximations to more global 3D curves, obtained independently at each voxel. Thus the local helical approximations will typically differ from one location (voxel) to the next.

In summary, the approach uses evidence of inferred helix curves and their local configurations in order to disambiguate the three cases outlined in Table 1. Since helices are parametric curves, and since they are represented by co-helical triplets of tangents, it is straightforward to check the number and the local configuration of the helices that pass through any given voxel. These ideas are developed into an algorithm, described in pseudocode in Appendix A, where implementation details are also discussed. We describe the experimental methods in Experimental methods section and then present results demonstrating the use of this algorithm to label human brain ODF data and to improve the performance of fibre tracking in the Results section.

Experimental methods

MRI acquisition

MRI data were acquired for one healthy subject on a Siemens 3-T Trio MR scanner (Siemens Medical Systems, Erlangen, Germany) using an 8-channel phased-array head coil. Diffusion encoding was achieved using a single-shot spin-echo echo planar sequence with twice-refocused balanced diffusion encoding gradients. A dataset designed for high angular resolution reconstruction was acquired with $N=99$ diffusion encoding directions with $b=3000$ s/mm^2, 10 T2-weighted images with $b=_0$ s/mm^2, 2-mm isotropic voxel size, 63 slices, TE=121 ms, TR=11.1 s, and GRAPPA parallel reconstruction with an acceleration factor of 2. A 1-mm isotropic resolution T1-weighted anatomical scan was also acquired (TR=9.7 ms, TE=4 ms, $\alpha=12°$). The T1-weighted scan was coregistered with the diffusion series by registration with the first $b=0$ s/mm^2 image using a mutual information-based registration algorithm (Maes et al., 1997).

Fibre ODF calculation

Fibre ODF estimation is a topic of much current research (e.g., Tournier et al., 2004, 2007; Anderson, 2005; Sakaie and Lowe, 2007; Dell'Acqua et al., 2007; Kaden et al., 2007; Jian and Vemuri, 2007; Descoteaux et al., 2007a,b; Alexander, 2005), and new and improved methods are likely to become available in the future as well. One of the strengths of the 3D curve inference algorithm (Savadjiev et al., 2006) and our proposed labeling algorithm is that they are (and will remain) applicable to any type of fibre ODF data, irrespective of the actual method that has been used to estimate it.

For the purposes of our experiments, we estimated the fibre ODF using spherical deconvolution. The technique of Descoteaux and Deriche (2007) and Descoteaux et al. (2007b) was used, which is a linear and regularised method performed on the diffusion ODF estimated from Q-Ball reconstruction. The regularisation step reduces spurious peaks and negative lobes in the fibre ODF. An $L=4$ spherical harmonic basis was used, with regularisation parameter $\lambda=0.006$ (Descoteaux et al., 2007a). The deconvolution kernel was estimated directly from the dataset being used, as done in Tournier et al. (2004). The maxima of the fibre ODFs, which should correspond either to the curve tangents of single curving tracts, or to the median direction in the case of a fan, were then extracted. These were used in the labeling algorithm and for tractography, as detailed below.

Labeling subvoxel geometries

The labeling methodology described in the Labeling of ambiguous subvoxel fibre tract configurations section was applied. Since the labeling algorithm is not intended to be run in voxels containing predominantly CSF or grey matter, voxels in which the fractional anisotropy (FA) was less than 0.1 or the mean diffusivity was greater than 10^{-6} mm^2/ms were masked out. Whereas it may still be possible for some CSF and/or grey matter voxels to be included, they typically provide only low support values to their neighbours in the 3D curve inference algorithm and thus do not alter the outcome significantly. Voxels in which the labeling is inconclusive, for example due to a lack of reliably inferred curves, e.g., in the CSF or grey matter, are assigned the label 'unknown.'

To avoid noise effects, the fibre ODFs were thresholded so that directions along which an ODF had a value below a threshold were ignored, while the others were considered as potential fibre orientations. The value of this threshold depends on the data and is determined empirically. In our implementation, we obtain good results with a threshold of $\mu+1.5\ \sigma$, where μ is the mean value of the ODF and σ is its standard deviation.

3D curve inference (Savadjiev et al., 2006) was run for one iteration with a small step size (the step size determines the strength of the regularisation) in order to infer the associated curvature–torsion–normal classes (i.e., the best-fitting helices) for each orientation. Since the focus of this work is on the inference of curves to allow for the labeling of subvoxel fibre tract configurations, any regularisation effects on the ODF data (i.e., reshaping) were discarded, and only the inferred classes were kept. The labeling was computed using the inferred helix curves and fibre ODF maxima, as described in Table 1. Further details are presented in Appendix A.

Tractography

Fibre tractography was run both using the labeling information, and without this information, for the sake of comparison. The tractography algorithm used was a deterministic, modified streamline propagation approach that optionally took into account the labeling information, fan polarity, and extent of the fanning fibre tracts. It is a modification of the tracking procedure described in Campbell et al. (2006) and Savadjiev et al. (2006). When using the labeling information, the tracking proceeded as follows. At voxels labeled

single curves, with either zero or nonzero curvature, the direction of propagation was given by the fibre ODF maximum direction. In the case of nonzero, smooth curvature, the fibre ODF actually has a broad maximum, and the centre of this broad range of curve tangents was taken as the direction of propagation. This means that the streamlines follow the tangent to the curve at the centre of the voxel, which effectively takes the streamline path along the (now discretised) curve. At voxels labeled fanning fibre tracts, the direction of propagation depended on whether the incoming direction was in the direction of the fan, or in the direction of the merge, as illustrated schematically in the bottom row of Fig. 1. Hence, the fan polarity vector as determined by the curve inference labeling algorithm was essential. If the dot product of the incoming direction with the fan polarity vector was positive, the decision to fan was taken, and the tracking algorithm followed all fibre directions from the fibre ODF. If the dot product was negative, the decision to merge was taken, and the tracking algorithm followed only the central, maximal fibre ODF value, as in the single curve case. For voxels with multiple-fibre ODF maxima, the maximum closest to the incoming direction was followed. When tracking without the labeling information, the central maximum/maxima of the fibre ODF was used in all cases.

In order to perform streamline tractography through a field of potentially fanning fibres, all possible combinations of directions must be sampled in order to extract all possible pathways. This was accomplished by running the streamline algorithm iteratively, choosing the direction of propagation randomly from the range of fibre directions within the fan whenever a streamline reached a voxel labeled "fan" and made the decision to fan, not to merge. 10,000 Iterations were used per start point. This iterative technique was only used in the tracking experiments that incorporated the labeling information.

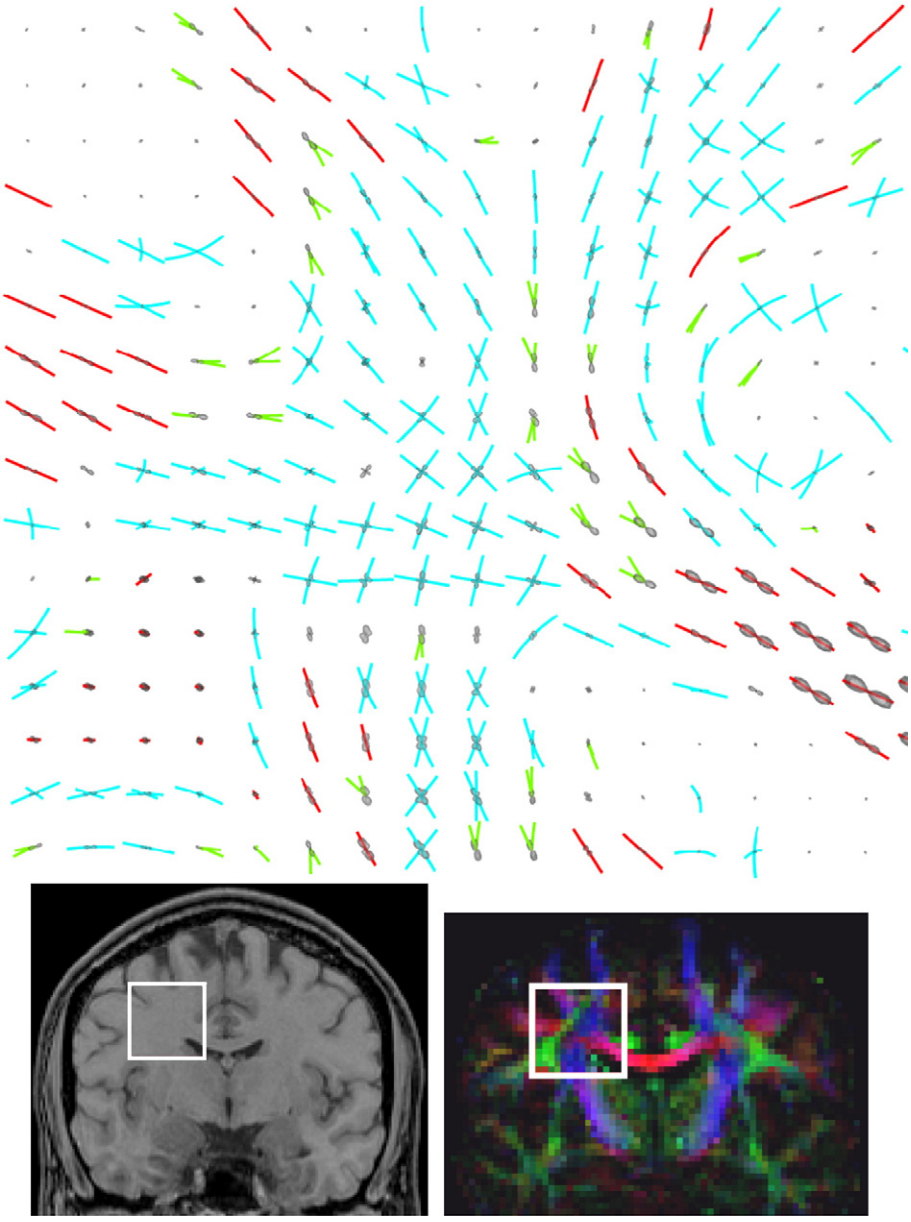

Fig. 3. Labeling of the ODFs (top) in a coronal ROI. The ROI is delimited with a white rectangle in the T1-weighted anatomical image (bottom left), and in the principal diffusion direction RGB map (bottom right). Inferred curves at each voxel are shown overlaid on the ODF. Blue: "cross" label. Red: "single curve" label. Green: "fan" label. In the fan case, the curves are shown oriented along the orientation of the polarity vector and delimit the extent of the fan.

Details of the streamline propagation common in both cases of using the labeling information and not were as follows. For all starting voxels, tracking was initiated on a $6\times6\times6$ rectilinear grid of start points in order to facilitate branching and more adequately reconstruct the dense fibre geometry. Streamlines were propagated using Fibre Assignment by Continuous Tracking (FACT) integration (Mori et al., 1999). The tracking was stopped if the fractional anisotropy (FA) was less than 0.1, the mean diffusivity was greater than 10^{-6} mm^2/ms, or the angular difference in the orientation of the tract from one voxel to the next was greater than 80°.

Tracking was initiated in all voxels in a small region of interest (ROI) in (1) the corpus callosum and (2) the internal capsule. For the ROI in the corpus callosum, tracts that turned erroneously down the cortical–spinal tract were excluded. These two experiments were repeated both using the labeling information and without. In order to assess the exact trajectories of the tracts obtained, the datasets were transformed into Talairach–Tournoux space and compared to the Talairach atlas (Talairach and Tournoux, 1988). The transformation from the T1 images to the Montreal Neurological Institute stereotaxic space was determined using in-house dedicated software (Collins et al., 1994). This transformation was used to transform both the T1 images and the coregistered connectivity profiles to stereotaxic space, in which localisation of the tracking results could be performed.

Results

We begin by visualising labeling results in two different ROIs in the brain dataset. The inferred helical curves at each voxel are shown overlaid on fibre ODF data, with crossing curves colored blue and single curves colored red. In the case of a fan, two curves colored green delimit the extent of the fan and are oriented in the direction of the fan polarity. For clarity, fibre ODFs are visualised after subtraction of the minimal inscribed sphere.

Fig. 3 (top) shows these results in an ROI containing partial volume averaging between the corpus callosum and the cortico-spinal tract. The ROI is delimited with a white rectangle in the T1-weighted anatomical image in Fig. 3 (bottom left). For reference, the red–green–blue (RGB) map of the principal diffusion directions of a coregistered DT reconstruction is shown in Fig. 3 (bottom right), and the ROI is delimited with a white rectangle. In the RGB map, the principal diffusion directions are encoded such that red denotes left–right orientation, blue denotes inferior–superior orientation, and green denotes anterior–posterior orientation.

Fig. 4 (top) shows a second example of voxel labeling and inferred helical curves in a different ROI, located in a sagittal plane. The ROI is delimited with a white rectangle in the T1-weighted anatomical image (Fig. 4 (bottom left)) and in the RGB image of principal diffusion

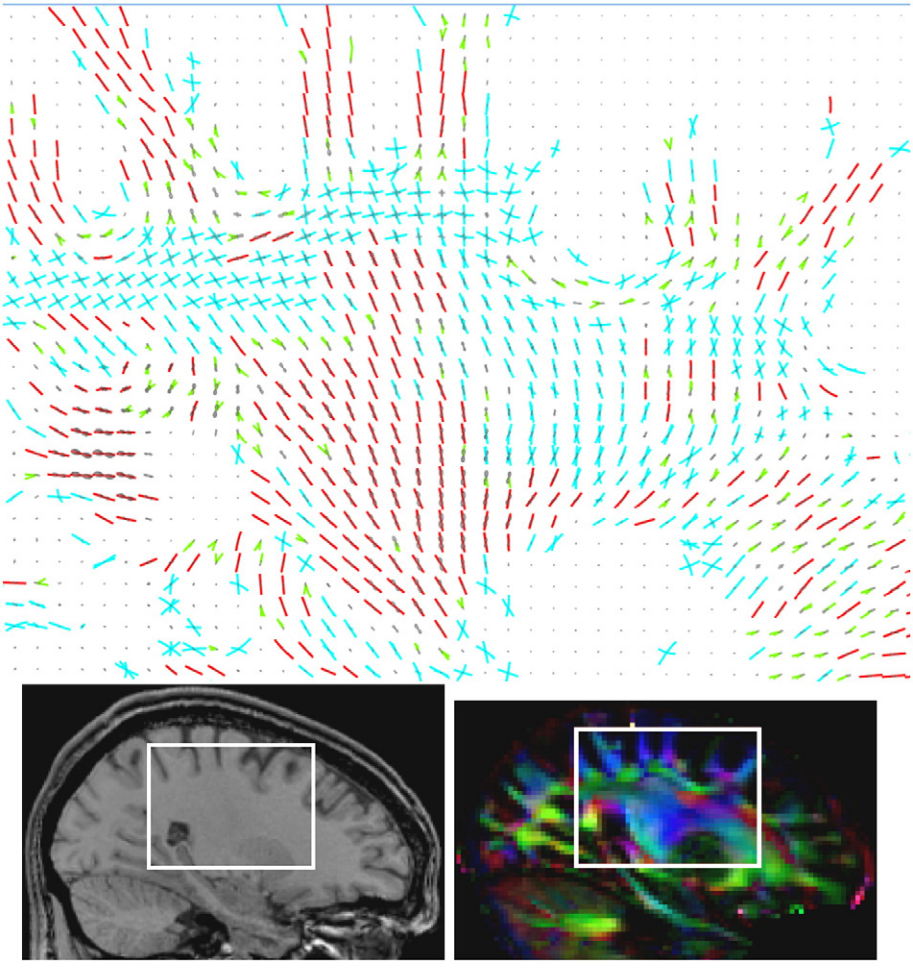

Fig. 4. Labeling of the ODFs (top) in a sagittal ROI. The ROI is delimited with a white rectangle in the T1-weighted anatomical image (bottom left), and in the principal diffusion direction RGB map (bottom right). Inferred curves at each voxel are shown overlaid on the ODF. Blue: “cross” label. Red: “single curve” label. Green: “fan” label. In the fan case, the curves are shown oriented along the orientation of the polarity vector and delimit the extent of the fan.

directions (Fig. 4 (bottom right)). The color coding of the inferred curves is the same as above. In both Figs. 3 and 4, the inferred labelings and curves appear anatomically plausible. Note that these figures present a 2D projection of curves that are three-dimensional, thus there is an effect of perspective.

We then visualise in Fig. 5 (top row) fibre-tracking results with and without the use of labeling information to direct the streamline tracking algorithm, with a seed placed in the corpus callosum. Here, the left-hand side of the top row shows the results of tracking without the labeling information, and the right-hand side gives the results of tracking with labeling. The seed location is colored green, and the surface encompassing all voxels connected to it is visualised in red, with the anatomical image shown for reference. Note that the ROI used for presenting the labeling result in Fig. 3 intersects the recovered tracts and is representative of the labeling results in that area. Note also that any differences in the results between the left and right panels in the top row of Fig. 5 are due to the use of labeling information only. The bottom row in Fig. 5 presents a drawing of the anatomy in that region, adapted from the Nieuwenhuys et al. (1988) atlas.

In Fig. 5 (top row), inclusion of the labeling information allows tractography to follow fibre pathways closer to the cortex and recover their fanning as they part to project to different gyri. The effect illustrated here may appear subtle at first sight but is nevertheless important. Most tracking algorithms normally provide robust results in major fibre pathways such as the corpus callosum, and the labeling information is expected to improve fibre-tracking performance mainly in regions with more complex fibre geometries such as those near the cortex.

The recovered connections are commissural connections between premotor cortex (area 6, Talairach and Tournoux, 1988) in the left hemisphere and area 6 in the right. When labeling is incorporated, almost the entire premotor strip on both sides is reached. The fanning of the callosal projections captured in these fibre-tracking results, as well as in the labeling results in Fig. 3, is consistent with the anatomy illustrated in the Nieuwenhuys et al. (1988) atlas (Fig. 5, bottom row).

A second fibre-tracking example is presented in Fig. 6 (top row), where the tracking algorithm was initiated in a seed region (green) located in the internal capsule. The results are presented in the same manner as in Fig. 5, that is, the surface encompassing all voxels reached from the seed region is colored red, and the anatomical image is shown for reference. Again, when using labeling information (Fig. 6, top row, right) a larger extent of the fan of the tracts towards the cortex is recovered as compared to the case without labeling (Fig. 6, top row, left). In particular, in the case of tracking with labeling, there are connections to areas 6, 4, 3, 5 and 7 (Talairach and Tournoux, 1988). In the case without labeling, the connections to area 6 are absent, and there are less dense connections to the other areas. The ROI used for presenting the labeling result in Fig. 4 intersects the recovered tracts and is representative of the labeling results in that area. Furthermore, any differences in the results between the left and right panels in the top row of Fig. 6 are due to the use of labeling information only. The fanning

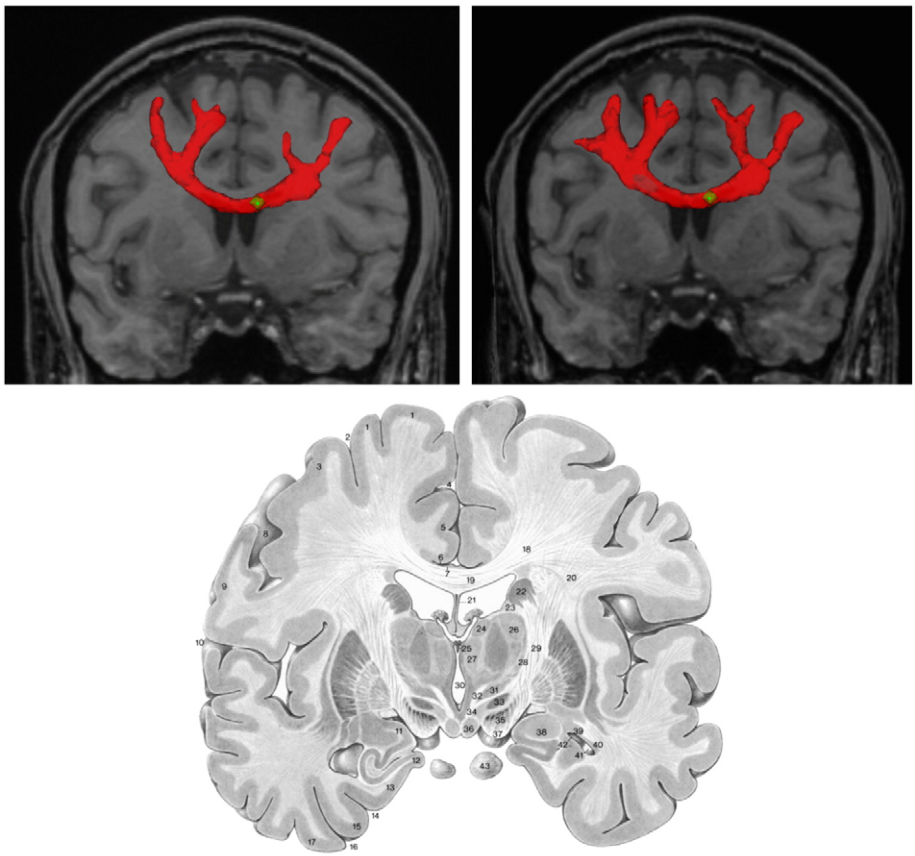

Fig. 5. Top row: examples of fibre-tracking results without using labeling information (left) and with the use of labeling information (right), starting from a seed region (green) placed in the corpus callosum. The ROI in which labeling results were presented in Fig. 3 intersects these tracking results. Bottom row: drawing of the anatomy in a coronal slice in this region, adapted from the Nieuwenhuys et al. (1988) atlas.

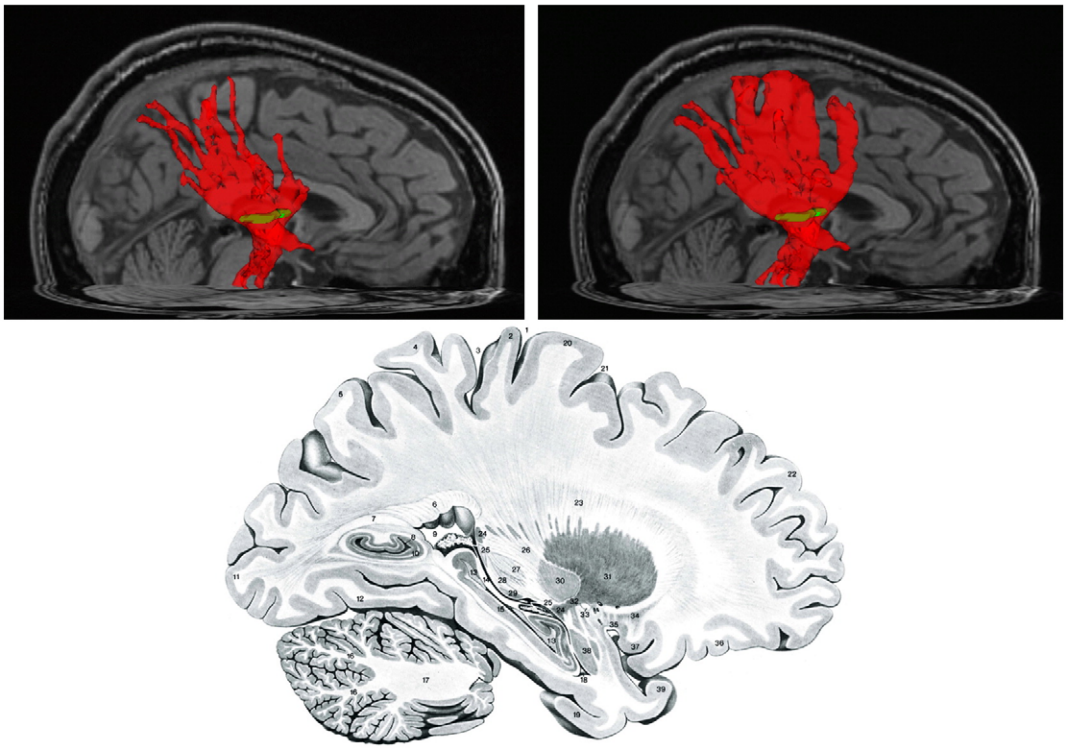

Fig. 6. Top row: tracking results without using labeling information (left) and with the use of labeling information (right). A seed region (green) was placed in the internal capsule. The ROI in which labeling results were presented in Fig. 4 intersects these tracking results. Bottom row: drawing of the anatomy in a sagittal slice in this region, adapted from the Nieuwenhuys et al. (1988) atlas.

observed in these fibre-tracking results, as well as in the labeling results in Fig. 4, is consistent with the anatomy illustrated in Fig. 6 (bottom row), which is a figure adapted from the Nieuwenhuys et al. (1988) atlas.

Discussion and conclusion

In this article, we present an algorithm for the disambiguation of complex subvoxel fibre tract configurations that can confound fibre tractography even when working with high angular resolution fibre ODF data. To our knowledge, this is the first algorithm that attempts to distinguish curving and fanning subvoxel fibre tract configurations. The algorithm is based on the 3D curve inference method (Savadjiev et al., 2006) for inferring local helical approximations to underlying fibre tracts. These local approximations are used in a scheme that labels diffusion ODFs in one of three categories: (1) single-fibre bundles (potentially with subvoxel curvature), (2) fanning fibre tracts, and (3) crossing fibre tracts. The importance of this labeling information is demonstrated through fibre-tracking experiments with and without labeling in human brain data. Although no ground truth is available for the brain dataset, the callosal projections as well as the connections passing through the internal capsule are known to fan and project towards the cortex, and the inclusion of labeling information allows for a better recovery of this fanning, as evidenced through the tractography results. The technique will be valuable in the future for tracing more subtle cortico-cortical association pathways, for instance for inferring the connectivity between different areas of fMRI activation. As the exact course of these pathways is still not well established, we have shown results only in larger pathways. In these pathways, we believe that the benefit of our technique is clear: for instance, the projection fibres of the thalamo-cortical tract reach more cortical areas using the labeling. The ability to see these projections is important, e.g., in assessing how they are affected in white matter disease, and in turn in inferring what functional deficits might arise.

Additionally, we anticipate that the labeling could be used to aid in the interpretation of scalar parameters calculated from diffusion data, such as the anisotropy index. Anisotropy can be artefactually lower in voxels containing non-collinear combinations of fibre orientations, such as fans and crossings. Having voxels labeled as such can help, for instance, in assessing whether changes in anisotropy in disease states are due to changes at the cellular level or changes in subvoxel geometries. These ideas will be the subject of future studies.

Future work will also address bottleneck fibre tract configurations, which are generally not subvoxel configurations but (at typical diffusion imaging spatial resolution) occur rather on a larger scale. Locally, they are indistinguishable from fans, however, they still require a separate action to be taken in tractography. In contrast to the case of fans, bottlenecks require the tracking algorithm to follow only one path, the one with a curvature consistent with the incoming path. Another configuration to be considered in future work is that of a branch, in which a single-fibre tract splits into two distinct tracts. The distinction between a fan and a branch is currently difficult due to limitations in spatial and angular resolution.

We end this discussion by noting that although our labeling methodology was used in conjunction with a deterministic streamline propagation tractography algorithm, it could also be incorporated in a probabilistic tractography framework. Probabilistic tractography should take into account uncertainty due to noise, acquisition limitations such as resolution, and inaccuracy in the model. By providing an explicit model of the underlying subvoxel fibre configurations, the labeling algorithm helps to reduce the uncertainty due to inaccuracy in

the model. For example, in the case of a curving fibre, the probability distribution function (pdf) of the estimated fibre direction should reflect uncertainty factors such as noise and angular resolution but would still be distributed more tightly around the ODF maximal direction than it would be in the case of the fan configuration.

In comparison, current state-of-the-art probabilistic tractography algorithms such as those of Parker and Alexander (2005), Seunarine et al. (2006), and Behrens et al. (2007) do not take into account subvoxel fibre geometries with comparable fibre ODFs, and this can affect their measures of uncertainty and ultimately the tractography results. For example, Parker and Alexander (2005) generate pdfs of the estimated fibre direction(s) at each voxel that are mixtures of Gaussian densities parametrised according to the anisotropy of the ODF data. Seunarine et al. (2006) improve on this model by replacing the Gaussian pdf for fibre directions (which produces circular pdf contours) with the Bingham distribution, which can account for more general elliptic clusters of fibre orientations. Nevertheless, in both cases, a broader ODF peak will necessarily result in a "dispersion" of the probability of connection and thus tractography would follow a broader path and result in a fanning behaviour. This is undesirable in the case of single curving fibre configurations, and this limitation is acknowledged in Parker and Alexander (2005). Behrens et al. (2007) sample from the posterior distribution of principal diffusion directions (PDDs) in a multi-tensor model. In their probabilistic framework, the uncertainty in the data is related to the noise level, which influences the repeatability of the sampling of the PDDs. Thus, in the theoretical case where the data is noise-free, a broad tensor peak may result in a path with high certainty that would not fan. Alternatively, a broad tensor peak in noisy data is likely to produce a path with high uncertainty that would necessarily fan. In either case, an explicit distinction between fanning and curving fibres is not made, and the tractography algorithm may produce incorrect results in the presence of such subvoxel configurations.

As stated in the introduction, an advantage of our labeling method is that it does not rely solely on the shape of ODF peaks to determine the underlying fibre configurations. It can robustly label curving and fan configurations, thus providing a tractography algorithm, whether deterministic or probabilistic, with additional information in the presence of broad ODF peaks that can be used to better interpret the associated uncertainty and to take an informed action.

Acknowledgments

This work was supported by funding from NSERC, FQRNT and an INRIA-FQRNT exchange program.

Appendix A

This Appendix presents implementation details related to the labeling algorithm described in the Labeling of ambiguous subvoxel fibre tract configurations section. It also presents the actual algorithm in pseudocode. In what follows, the term 'fibre ODF' refers to the thresholded fibre ODF discussed in the Labeling subvoxel geometries section. For example, 'fibre ODF directions' refer to those directions in the fibre ODF where the value of the ODF is above the threshold.

Algorithm 1

Voxels are first classified into one group that contains fibre ODFs with a single maximum (peak) and another that contains ODFs with multiple maxima. Algorithm 1 uses the normalised variance of the ODF values in order to further segment single-maximum ODFs. In practice, this normalised variance is computed as the general fractional anisotropy (GFA) index of ODFs, as defined in Tuch (2004). This GFA thresholding is an empirical decision, justified by the consideration that single-maximum ODFs in voxels that contain fanning fibre tracts are expected to have a lower GFA index than those in voxels with single-fibre bundles, as in the corpus callosum for example. To avoid false fan labels, single-maxima ODFs with a high GFA are automatically considered as denoting single curves, provided that there is an underlying curve, i.e., there is a co-helical triplet of tangents, the parameters of which (curvature and torsion) agree with the curvature–torsion–normal class inferred for the central orientation in the triplet through 3D curve inference. In case the method in the Labeling of ambiguous subvoxel fibre tract configurations section fails to infer sufficient curve information to label the voxel into one of the three configuration categories, it is labeled as unknown. In practice, this happens rarely, in particular in the case of single-peaked ODFs with a high GFA, where the label "unknown" is not expected to be assigned. If tracking encounters the label "unknown", it simply follows the ODF maximum which is closest to the incoming direction.

Algorithm 1: The labeling algorithm

Data: A voxelized volume containing one fibre ODF per voxel
Result: A labeling of the input voxels into one of four possible categories: single curve, cross, fan, background/unknown

```
Obtain a list of fibre ODF directions for all voxels;
foreach voxel do
    nummax ← number of local maxima at this voxel;
    if nummax = 1 and GFA ≥ threshold then
        Check if a helical curve is associated with this maximum.;
        if yes then
            Label voxel as "single curve"
        else
            Label voxel as "unknown"
    else
        if nummax = 1 and GFA < threshold then
            Check if ODF is a "fan" (Algorithm 2);
            if yes then
                Label voxel as "fan"
            else
                if a helical curve is associated with it then
                    Label voxel as "single curve"
                else
                    label voxel as "unknown"
    else
        Check if ODF is a "cross" (Algorithm 3);
        if yes then
            Label voxel as "cross"
        else
            Label voxel as "unknown"
```

Algorithm 2

Algorithm 2 is used to check whether a voxel contains a fan configuration of fibre tracts. To do so, one determines the set of co-helical triplets that exist within a certain neighbourhood, e.g., $5 \times 5 \times 5$, centred at the current voxel, that also agree with the class of the central element in the triplet. This process determines the local curves that pass through the current voxel. These curves are then examined to determine if they obey the conditions of a fan, that is, a configuration in which there are two curves that share a common point with a common tangent before entering the voxel, and pass through different points and have non-parallel tangents after leaving the voxel.

Algorithm 2: Sub-algorithm to check if an ODF is a fan

Data: A fibre ODF
Result: 1 if the ODF is a fan, 0 otherwise

```
Initialize an empty list L of curves;
flag ← 0;
foreach ODF direction do
    Find associated curves (see Section 3 for details);
    if a curve C is found then
        if L is not empty then
            foreach curve C̃_L in L do
                Check if C and C̃_L obey the fan configuration;
                if yes then
                    ODF is a fan;
                    flag ← 1;
                    break;
            if flag = 0 then
                add C to L
            else
                break
        else
            add C to L
return flag;
```

Algorithm 3

Algorithm 3 is used to check whether a voxel contains a cross configuration of fibre tracts. The idea is the same as with Algorithm 2, i.e., the algorithm proceeds first by determining in the same manner which curves pass through the current voxel. These curves are then compared among themselves to check if a cross configuration is observed, that is, a configuration in which two or three locally non-parallel curves enter the voxel, intersect in it and then leave the voxel. Whether the curves are locally non-parallel or not is determined by requiring that the angular separation between their associated tangents (ODF directions) in the crossing voxel is above a certain threshold, set to 35 ° in our implementation.

Algorithm 3: Sub-algorithm to verify if an ODF is a cross

Data: A fibre ODF
Result: 1 if the ODF is a crossing, 0 otherwise

```
Initialize an empty list L of curves;
number of crossings ← 0;
foreach ODF direction do
    Find associated curves (see text for details);
    if a curve C is found then
        if L is not empty then
            foreach curve C̃_L in L do
                Verify if C and C̃_L obey the cross configuration;
                if yes then
                    add C to L;
                    number of crossings ← number of crossings + 1;
                    break;
        else
            add C to L
if number of crossings = 1 then
    return 1;
else
    return 0;
```

References

Alexander, D.C., 2005. Maximum entropy spherical deconvolution for diffusion MRI. Image Processing in Medical Imaging, pp. 76–87.

Anderson, A., 2005. Measurements of fiber orientation distributions using high angular resolution diffusion imaging. Magn. Reson. Med. 54, 1194–1206.

Assaf, Y., Basser, P.J., 2005. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. Neuroimage 27 (1), 48–58.

Basser, P.-J., Matiello, J., Le Bihan, D., 1994. MR diffusion tensor spectroscopy and imaging. Biophys. J. 66, 259–267.

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.

Behrens, T.E., Berg, H.J., Jbabdi, S., Rushworth, M.F., Woolrich, M.W., 2007. Probabilistic diffusion tractography with multiple fibre orientations: what can we gain? NeuroImage. 34 (1), 144–155.

Callaghan, P.T., 1993. Principles of Nuclear Magnetic Resonance Microscopy. Oxford University Press, Oxford.

Campbell, J.S.W., Siddiqi, K., Rymar, V.V., Sadikot, A.F., Pike, G.B., 2005. Flow-based fiber tracking with diffusion tensor and Q-ball data: validation and comparison to principal diffusion direction techniques. NeuroImage 27, 725–736.

Campbell, J.S.W., Savadjiev, P., Siddiqi, K., Pike, G.B., 2006. Validation and regularization in diffusion MRI tractography. Proc. International Symposium on Biomedical Imaging (ISBI), pp. 351–354.

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.

Dell'Acqua, F., Rizzo, G., Scifo, P., Clarke, R., Scotti, G., Fazio, F., 2007. A model- based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging. Transactions in Biomedical Engineering 54 (3), 462–472.

Descoteaux, M., Deriche, R., 2007. Sharpening improves clinically feasible Q-Ball imaging reconstructions. Joint Annual Meeting ISMRM-ESMRMB, International Society of Magnetic Resonance in Medicine, Berlin, Germany, p. 906.

Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2007a. Regularized, fast, and robust analytical Q-Ball imaging. Magn. Reson. Med. 58 (3), 497–510.

Descoteaux, M., Deriche, R., Anwander, A., 2007b. Deterministic and Probabilistic Q-Ball Tractography: From Diffusion to Sharp Fiber Distributions, Tech. Rep. 6273, INRIA Sophia Antipolis. July.

Friman, O., Farneback, G., Westin, C.-F., 2006. A bayesian approach for stochastic white matter tractography. IEEE Transactions in Medical Imaging 25 (8), 965–978.

Hummel, R., Zucker, S., 1983. On the foundations of relaxation labeling processes. IEEE Trans. Pattern Anal. Mach. Intell. 5, 267–287.

Jansons, K.M., Alexander, D.C., 2003. Persistent angular structure: new insights fom diffusion magnetic resonance imaging data. Inverse Problems 19, 1031–1046.

Jian, B., Vemuri, B.C., 2007. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. IEEE Trans. Med. Imaging 26 (11), 1464–1471 (Nov).

Kaden, E., Knosche, T.R., Anwander, A., 2007. Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. NeuroImage 37, 474–488.

Le Bihan, D., Mangin, J.-F., Poupon, C., Clark, C., Pappata, S., Molko, N., Chabriat, H., 2001. Diffusion tensor imaging: concepts and applications. J. Magnetic Resonance Imaging 13, 534–546.

Le Bihan, D., Poupon, C., Amadon, A., Lethimonnier, F., 2006. Artifacts and pitfalls in diffusion MRI. J. Magn. Reson. Imaging 24 (3), 478–488.

Lin, C.P., Wedeen, V., Chen, J., Yao, C., Tseng, W., 2003. Validation of diffusion spectrum magnetic resonance axonal fiber imaging with manganese-enhanced optic tracts and ex-vivo phantoms. NeuroImage 19, 482–495.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multimodality image registration by maximization of mutual information. IEEE Trans. Med. Imaging 16 (2), 187–198.

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

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

Özarslan, E., Shepherd, T., Vemuri, B., Blackband, S., Mareci, T., 2006. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). NeuroImage 31 (3), 1086–1103.

Parent, P., Zucker, S., 1989. Trace inference, curvature consistency, and curve detection. IEEE Trans. Pattern Anal. Mach. Intell. 11, 823–839.

Parker, G., Alexander, D., 2005. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. Phil. Trans. R. Soc. B 360, 893–902.

Ramírez-Manzanares, A., Rivera, M., 2006. Basis tensor decomposition for restoring intra-voxel structure and stochastic walks for inferring brain connectivity in DT-MRI. Int. Journal of Computer Vision 69 (1), 77–92.

Sakaie, K.E., Lowe, M.J., 2007. An objective method for regularization of fiber orientation distributions derived from diffusion-weighted MRI. NeuroImage 34, 169–176.

Savadjiev, P., Campbell, J.S.W., Pike, G.B., Siddiqi, K., 2006. 3D curve inference for diffusion MRI regularization and fibre tractography. Med. Image Anal. 10 (5), 799–813.

Seunarine, K.K., Cook, P.A., Embleton, K., Parker, G., Alexander, D., 2006. A general framework for multiple-fibre PICo tractography. Proc. Medical Image Understanding and Analysis (MIUA).

Talairach, J., Tournoux, P., 1988. Co-Planar Stereotaxic Atlas of the Human Brain: 3- Dimensional Proportional System : An Approach to Cerebral Imaging. Thieme Medical Publishers.

Tournier, J.D., Calamante, F., Gadian, D.-G., Connelly, A., 2004. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage 23, 1176–1185.

Tournier, J.-D., Calamante, F., Connelly, A., 2007. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. NeuroImage 35 (4), 1459–1472.

Tuch, D., Diffusion MRI of complex tissue structure, PhD thesis, Harvard-MIT (2002).

Tuch, D., 2004. Q-Ball imaging. Magn. Reson. Med. 52, 1358–1372.

Wedeen, V., Reese, T., Tuch, D., Weigel, M., Dou, J.-G., Weiskoff, R., Chessler, D., 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.