

## 3D curve inference for diffusion MRI regularization and fibre tractography <sup>☆</sup>

Peter Savadjiev <sup>a,\*</sup>, Jennifer S.W. Campbell <sup>b</sup>, G. Bruce Pike <sup>b</sup>, Kaleem Siddiqi <sup>a</sup>

<sup>a</sup> Centre for Intelligent Machines and School of Computer Science, McGill University,

McConnell Eng. Building, 3480 University Street, Room 410, Montréal, QC, Canada H3A 2A7

<sup>b</sup> McConnell Brain Imaging Centre, Montreal Neurological Institute, McGill University, Montréal, Canada

Received 16 January 2006; received in revised form 3 May 2006; accepted 23 June 2006

Available online 21 August 2006

### Abstract

We develop a differential geometric framework for regularizing diffusion MRI data. The key idea is to model white matter fibres as 3D space curves and to then extend Parent and Zucker's 2D curve inference approach [Parent, P., Zucker, S., 1989. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 823–839] by using a notion of *co-helicity* to indicate compatibility between fibre orientations at each voxel with those in a local neighborhood. We argue that this provides several advantages over earlier regularization methods. We validate the approach quantitatively on a biological phantom and on synthetic data, and qualitatively on data acquired *in vivo* from a human brain. We also demonstrate the use of the technique to improve the performance of a fibre tracking algorithm.

© 2006 Elsevier B.V. All rights reserved.

**Keywords:** Diffusion MRI; Diffusion tensor imaging; High angular resolution diffusion imaging; Regularization; Curve inference; Fibre Tractography

### 1. Introduction and related work

Diffusion magnetic resonance imaging (MRI) is a technology for measuring the random thermal (Brownian) motion of water molecules in live tissues. In the case of fibrous tissues, such as the white matter of the central nervous system, the fibre structures restrict this diffusive motion, and thus the resulting diffusion is maximal along the orientation(s) of the underlying fibre(s) (Lin et al., 2003). The diffusion MRI signal can thus be used to infer the presence and geometry of white matter fibres. Whereas the physical processes that underlie the acquisition of the diffusion MRI signal are beyond the scope of this article (see e.g. Basser et al., 1994; Le Bihan et al., 2001; Tuch

et al., 1999), the essential idea is to first obtain a set of raw diffusion weighted images by applying magnetic gradients in different directions. From these images, one obtains an estimate of the 3D probability distribution function (PDF) describing the diffusion of the water molecules at a discrete set of locations (voxels) in a 3D volume. From the PDF, the orientation distribution function (ODF) at each voxel can then be calculated by computing the integral of the PDF in the radial direction. The 3D ODF dataset can then be used to perform higher-level processing tasks, such as fibre tracking (e.g. Mori et al., 1999; Basser et al., 2000; O'Donnell et al., 2002; Behrens et al., 2003; Jackowski et al., 2005; Parker and Alexander, 2005; Campbell et al., 2005).

In diffusion tensor imaging (DTI), the diffusion PDF is modeled as a zero-mean 3D Gaussian distribution (Basser et al., 1994). Because of the assumption of a Gaussian model, DTI has a limited representational power in voxels that exhibit partial volume averaging effects, such as voxels where distinct fibres cross, or branch, or where fibres are

<sup>☆</sup> A preliminary version of this paper appears in the proceedings of the 8th International Conference On Medical Image Computing and Computer Assisted Intervention, Palm Springs, California, 2005.

\* Corresponding author. Tel.: +1 514 398 6319; fax: +1 514 398 7348.

E-mail address: [petersv@cim.mcgill.ca](mailto:petersv@cim.mcgill.ca) (P. Savadjiev).

locally parallel but have different curvature and torsion, thus forming a “neck”-like configuration. In such voxels, the ODF can be relatively isotropic, and thus provide ambiguous information about the orientations of the crossing fibres. As a consequence there has recently been an increased interest in high angular resolution diffusion imaging (HARDI) (e.g. Tuch et al., 1999), which is used for characterizing non-Gaussian diffusion processes. In particular, it can model ODFs with multiple maxima in voxels where fibres cross or branch, thus overcoming certain limitations inherent in the DTI representation. The disadvantage of HARDI with respect to DTI is that it requires longer acquisition times. Furthermore, both types of imaging methods can suffer from acquisition noise. Different algorithms exist to reconstruct the ODF given HARDI measurements, including the q-ball imaging (QBI) technique of Tuch et al. (2003) and the recent use of spherical harmonic series approximations to provide analytic expressions for the ODF (Descoteaux et al., 2005; Özarslan et al., 2005).

To illustrate aspects of DTI and HARDI, consider a biological phantom created by overlaying two rat spinal cords (Campbell et al., 2005), as shown in Fig. 1 (left). The ODF reconstruction using DTI is shown in Fig. 1 (middle), and the reconstruction using QBI (Tuch et al., 2003) from HARDI measurements is visualized in Fig. 1 (right). In both cases, the ODFs are shown with the minimal inscribed sphere (MIS) subtracted, to better emphasize the locations where there is evidence of orientation preference. The DTI ODFs capture the orientation of each fibre tract, but in the vicinity of the crossing, the ODFs are relatively isotropic and thus ambiguous. The QBI reconstruction is better able to model the multiple orientations present in the area of the crossing.

In this article, we develop a novel technique for regularizing diffusion ODF data, while allowing for multiple fibre orientations at a voxel to be supported when appropriate. Its main advantages are: (1) the possibility of representing multiple ODF maxima at a location, as in the case of fibre crossings and branchings (Fig. 1), (2) applicability to both DTI and QBI ODF reconstructions, (3) estimates of curvature and torsion at each voxel, which can be useful to guide fibre tracking algorithms, and (4) numerical robustness in

the vicinity of sparse data. By “sparse” we refer to data such as the phantom in Fig. 1, where the fibres of interest pass through a relatively small number of voxels with anisotropic ODFs, embedded in a volume with isotropic ODFs that are known to describe only the background. Such data can be challenging for vector diffusion regularization methods such as (Tschumperlé and Deriche, 2002), because the central difference approximations used to compute derivatives numerically can be unreliable.

Our algorithm is fundamentally different from other methods in the literature, e.g. Tschumperlé and Deriche (2002), Wang et al. (2004), and Martin-Fernandez et al. (2004), which do not incorporate an explicit geometric model of the underlying fibre geometry. We model white matter fibres as 3D curves, and we pose the regularization problem as one of 3D curve inference from an initial set of ODFs. We carry out an extension of Parent and Zucker’s 2D curve inference framework (Parent and Zucker, 1989) to 3D, where at each voxel, the value of the ODF in a given orientation is related to the initial confidence in a curve tangent with that orientation being present. In this setting, curve tangents represent local fibre orientations. Those orientations that are supported by other nearby (local) orientations via a geometrical model gain confidence, while others are suppressed. This is accomplished by introducing a notion of *co-helicity* to model the compatibility of each orientation at each voxel with those in a local neighborhood.

The regularized ODFs produced using 3D curve inference are not a quantitative measure of diffusion, since the algorithm sharpens the ODFs and does not preserve mean diffusivity. However, the ODF maxima can be assumed to be tangent to fibres. Thus, in the case of DTI data, the curve inference algorithm not only performs regularization, in the sense of smoothing noisy orientations, but it also allows for the inference of high angular resolution ODFs directly from the diffusion ODF computed using only the single tensor model. This is an attractive feature when the available diffusion MRI acquisition has either sparse diffusion encoding directions, or low  $b$  values (acquisition schemes suitable for HARDI reconstructions typically require 100–500 diffusion encoding directions and  $b$  values  $\geq 3000$  s/mm<sup>2</sup>). Furthermore, the discrete fibre curvature




Fig. 1. A biological phantom created by overlaying two rat cord spinal cords (Campbell et al., 2005) (left) yields a DTI reconstruction (middle) and a QBI reconstruction (Tuch et al., 2003) from HARDI data (right). Both ODF sets are shown with MIS subtraction, for ease of visualization.

and torsion estimates obtained using the algorithm have the potential to improve fibre tracking results in the presence of partial volume averaging effects. For example, they provide valuable information to distinguish between the cases where: (i) a single fibre is branching, (ii) two fibres are crossing one another, or (iii) two distinct fibres are in a “neck”-like configuration. Finally, the 3D curve inference algorithm can take as input not only *diffusion* ODFs (DTI or HARDI reconstructions), but also any ODF dataset, including, for example, fibre ODFs (Tournier et al., 2004).

This article is organized as follows. We begin with a brief review of related work on diffusion MRI regularization in Section 1.1. We then introduce the notion of *co-helicity* to model the local compatibility between fibre orientations in Section 2. In Section 2.3 we show how this can be used to define average local support, which in turn can be maximized using relaxation labeling techniques (Rosenfeld et al., 1976; Hummel and Zucker, 1983). This results in an algorithm that extends Parent and Zucker’s 2D curve inference approach (Parent and Zucker, 1989) to the case of 3D ODF data. We discuss implementation details in Section 3 and present several experimental results on DTI and HARDI data obtained for the rat spinal cord phantom, on in vivo human brain data, as well as on synthetic data in Section 4. We also demonstrate applications of the regularized ODFs to fibre tracking in the phantom dataset as well as in the in vivo brain dataset. We discuss these results in Section 5 and conclude with Section 6.

### 1.1. Related work

In this subsection we provide a brief review of related work on diffusion MRI regularization. Earlier techniques include the spaghetti plate model of Poupon et al. for the regularization of DTI data, where local curvature measures are integrated along the length of a white matter fascicle (Poupon et al., 2000). Martín-Fernandez et al. (2004) regularize the six individual elements of each DTI tensor using a six-dimensional, multivariate Gaussian Markov random field (MRF) framework. Wang et al. (2004) propose the novel idea of simultaneously estimating and regularizing DTI tensors directly from raw data images. Coulon et al. (2001) approach the DTI regularization process in two separate steps. First, they regularize the orientation fields using an adaptation of variational methods for diffusion of vector data on non-flat manifolds (Chan and Shen, 1999; Tang et al., 2000). Second, they use a scalar non-linear anisotropic diffusion process to regularize each tensor eigenvalue separately. Tschumperlé and Deriche (2002) propose a flow algorithm using variational methods for regularizing DTI data, where the smoothing of diffusivities (the eigenvalues of the diffusion tensor) is separated from the smoothing of orientations (the eigenvectors). They introduce a novel extension of the Perona–Malik diffusion scheme (Perona and Malik, 1990) to vector valued data, with the constraint that orthonormality between the eigen-

vectors at each location is preserved. In more recent work, Chef d’Hotel et al. (2004) develop related ideas, where these constraints are handled by restricting the data to lie on suitably defined manifolds.

All of the above algorithms have the inherent limitations that: (1) they were designed explicitly for DTI data and cannot be trivially extended to handle HARDI data; and (2) they assume (explicitly or implicitly) a single fibre orientation at each voxel and thus cannot handle branchings or crossings. To our knowledge the only algorithms which address these concerns, at least in part, are those proposed in Ramírez-Manzanares and Rivera (2003), Chen et al. (2004), and Cointepas et al. (2002). Ramírez-Manzanares and Rivera (2003) express the observed tensors as linear combinations of a set of highly anisotropic basis tensors and minimize a cost function to recover the coefficients of the combination. Chen et al. (2004) propose an algorithm for regularizing HARDI data using a variational method where the diffusion profile at each voxel is approximated with a spherical harmonic series. Cointepas et al. (2002) use a spin glass framework to minimize a global energy of a fibre map, given the diffusion data. Whereas some of these latter methods have been validated on synthetic data, few (or none) have been demonstrated on a biological phantom with known ground truth fibre orientations.

## 2. 3D curve inference

Owing to the importance of visual information carried by curves in images, recovering their trace has been recognized as an important problem by the computer vision community. A prominent example is the 2D curve inference framework introduced by Parent and Zucker (1989), where the problem is cast as one of recovering likely trace, tangent and curvature fields, given initial measurements of curve tangent estimates. This problem is then solved using a relaxation labeling mechanism (Rosenfeld et al., 1976; Hummel and Zucker, 1983) where local curve tangents that are well supported (via a notion of co-circularity between nearby tangents) gain confidence, while others are suppressed.

More recently, extensions have been introduced to create new models for contour-based stereo correspondence, in which orientation disparities are used in conjunction with positional disparities in order to recover 3D space curve structure from two 2D images (Alibhai and Zucker, 2000; Li and Zucker, 2003). Using local Frenet approximations of 3D curves, considerations from projective geometry define the compatibility between pairs of candidate matches (Li and Zucker, 2003).

Related ideas have been introduced towards the inference of texture flow patterns, which play an important role in perceptual organization (Ben-Shahar and Zucker, 2003). Here the authors carry out a geometrical analysis of texture flow, and introduce an algorithm for the regularization of “noisy” 2D orientation fields. They demonstrate many

advantages over traditional orientation diffusion schemes, such as the ability to deal with missing information, and the ability to handle multiple dominant orientations at the same location. Ideas related to curve inference have also been considered within the tensor voting framework (Medioni et al., 2000).

Inspired by these developments, we introduce a differential geometric framework for 3D curve inference. It is motivated by the consideration that contextual information is key to the interpretation of 3D orientation data. This contextual information is integrated locally through geometrical constraints (described below), in order to obtain reliable estimates of likely curves passing through a given location, together with estimates of their curvature and torsion.

We assume as input a 3D orientation distribution function sampled on a regular (typically rectangular) 3D lattice. We obtain estimates of the trace, tangent, curvature and torsion fields of curves in the 3D volume by using a notion of *co-helicity* between a set of three tangents specified at three locations. This is the natural extension to 3D of Parent and Zucker’s *co-circularity* constraint between a pair of tangents in 2D (Parent and Zucker, 1989). In this framework, an *osculating helix* (which has constant curvature and constant torsion) is used to locally approximate a curve passing through a given location.

We now review some basic properties of a helix. We then provide a definition for co-helicity and present an algorithm for determining co-helicity between three orientations, specified at three distinct locations. The notion of co-helicity is then incorporated into a relaxation labeling scheme, which allows for local contextual information to be integrated using geometric constraints.

### 2.1. Properties of a helix and co-helicity

A circular helix is a curve inscribed on the surface of a cylinder, such that at all points on the curve, the associated tangent vector forms a constant angle with the cylinder’s axis. Consider such a helix,  $\alpha(t)$ , parametrized by  $t$ , with its axis coinciding with the  $z$ -axis. Its equations and those of its unit tangent and unit normal are given by

$$\alpha(t) = (x(t), y(t), z(t)) = (r \cos(t), r \sin(t), ct), \tag{1}$$

$$\frac{\alpha'(t)}{\|\alpha'(t)\|} = \frac{(x'(t), y'(t), z'(t))}{\|(x'(t), y'(t), z'(t))\|} = \frac{(-r \sin(t), r \cos(t), c)}{\sqrt{r^2 + c^2}}, \tag{2}$$

$$\frac{\alpha''(t)}{\|\alpha''(t)\|} = \frac{(x''(t), y''(t), z''(t))}{\|(x''(t), y''(t), z''(t))\|} = (-\cos(t), -\sin(t), 0). \tag{3}$$

Here  $r$  is the radius of the helix and  $c$  is a constant defining the vertical separation of the helical loops (measured along the helix axis). Three distinct tangents along the helix can be used to illustrate the notion of co-helicity, as demonstrated in Fig. 2. Let  $\Pi$  be the plane orthogonal to the axis of helix  $\alpha(t)$ . Consider three distinct points on  $\alpha(t)$ , i.e.  $\alpha(t_1)$ ,  $\alpha(t_2)$  and  $\alpha(t_3)$ , with associated tangent vectors  $\alpha'(t_1)$ ,  $\alpha'(t_2)$ , and  $\alpha'(t_3)$ . The orthographic projection of the helix onto  $\Pi$




Fig. 2. An illustration of co-helicity between three helix tangent vectors, specified at three distinct points on helix  $\alpha(t)$ .

is a circle (shown with a dashed curve), and the projections of  $\alpha'(t_1)$ ,  $\alpha'(t_2)$  and  $\alpha'(t_3)$  onto  $\Pi$  are labeled  $\Pi(\alpha'(t_1))$ ,  $\Pi(\alpha'(t_2))$  and  $\Pi(\alpha'(t_3))$ , respectively. These three projections are co-circular (Parent and Zucker, 1989), since they are all tangent to the same circle in plane  $\Pi$ .

**Definition 1 (Co-Helicity).** Let three vectors  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^3$  be specified at three distinct locations  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{R}^3$ . Then  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  are *co-helical* if and only if there is a helix which:

- (1) passes through  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$ ,
- (2) has tangent vectors at these locations that are parallel to  $\pm\mathbf{v}_1, \pm\mathbf{v}_2$  and  $\pm\mathbf{v}_3$ , respectively.

This definition is based on the *orientations*  $\pm\mathbf{v}_1, \pm\mathbf{v}_2$  and  $\pm\mathbf{v}_3$  instead of simply the *directions*  $\mathbf{v}_1, \mathbf{v}_2$  and  $\mathbf{v}_3$  to reflect the fact that water molecule diffusion is an antipodally symmetric processes, i.e. diffusion along direction  $\mathbf{v}$  is the same as diffusion along direction  $-\mathbf{v}$  (and thus diffusion ODFs are always antipodally symmetric). It also reflects the fact that it is necessary to model local fibre orientations as opposed to fibre directions.

We now state the following propositions. **Proposition 1** will be used for increased efficiency in the implementation of our 3D curve inference algorithm (developed in the following subsection), and **Proposition 2** will be used in the proof of **Proposition 3**.

**Proposition 1.** Two tangents to a helix always form an equal angle to the line joining their points of contact with the helix.

**Proposition 2.** Consider two unit tangents to a helix,  $\alpha'(t_1)/\|\alpha'(t_1)\|$  and  $\alpha'(t_2)/\|\alpha'(t_2)\|$ . The unit difference vector between them is equal to the unit normal vector  $\alpha''(\frac{t_1+t_2}{2})/\|\alpha''(\frac{t_1+t_2}{2})\|$  to the helix at  $\alpha(\frac{t_1+t_2}{2})$ . Furthermore, this vector is orthogonal to the helix axis.

**Proof of Proposition 1.** Consider two unit tangent vectors to a helix. They are defined using Eq. (2) as  $\alpha'(t_1)/\|\alpha'(t_1)\|$  and  $\alpha'(t_2)/\|\alpha'(t_2)\|$ . Their locations in space are given by  $\alpha(t_1)$  and  $\alpha(t_2)$ , according to Eq. (1). The vector  $\mathbf{d}_{12}$  between these locations is given by  $\mathbf{d}_{12} = \alpha(t_1) - \alpha(t_2)$ . It is straightforward to check that

$$\left\langle \frac{\alpha'(t_1)}{\|\alpha'(t_1)\|}, \mathbf{d}_{12} \right\rangle = \left\langle \frac{\alpha'(t_2)}{\|\alpha'(t_2)\|}, \mathbf{d}_{12} \right\rangle, \quad (4)$$

where  $\langle \cdot, \cdot \rangle$  denotes the usual dot product in  $\mathbb{R}^3$ .  $\square$

**Proof of Proposition 2.** Again, consider two unit tangent vectors to a helix, defined using Eq. (2) as  $\alpha'(t_1)/\|\alpha'(t_1)\|$  and  $\alpha'(t_2)/\|\alpha'(t_2)\|$ . Their difference is

$$\frac{\alpha'(t_2)}{\|\alpha'(t_2)\|} - \frac{\alpha'(t_1)}{\|\alpha'(t_1)\|} = \frac{1}{\sqrt{r^2 + c^2}} \begin{bmatrix} -2r(\cos(\frac{t_1+t_2}{2}) \sin(\frac{t_2-t_1}{2})) \\ -2r(\sin(\frac{t_1+t_2}{2}) \sin(\frac{t_2-t_1}{2})) \\ 0 \end{bmatrix}. \quad (5)$$

The magnitude of this difference vector is

$$\left\| \frac{\alpha'(t_2)}{\|\alpha'(t_2)\|} - \frac{\alpha'(t_1)}{\|\alpha'(t_1)\|} \right\| = \frac{2r \sin(\frac{t_2-t_1}{2})}{\sqrt{r^2 + c^2}}. \quad (6)$$

Dividing the difference vector (Eq. (5)) by its magnitude (Eq. (6)) in order to normalize it to unit length yields

$$\mathcal{N} = \begin{bmatrix} -\cos(\frac{t_1+t_2}{2}) \\ -\sin(\frac{t_1+t_2}{2}) \\ 0 \end{bmatrix}, \quad (7)$$

which is the unit normal vector at  $\alpha(\frac{t_1+t_2}{2})$ . Eqs. (1)–(3) define a helix whose axis coincides with the  $z$  axis. Clearly,  $\mathcal{N}$  has no  $z$  component, and hence is orthogonal to the helix axis.  $\square$

### 2.2. Determining co-helicity

The use of an osculating helix as a local approximation to an arbitrary (smooth) 3D curve allows for both curvature and torsion to be captured. In this subsection we explain how to fit a helical curve, if one exists, to three orientations defined at three 3D locations, all within some local neighborhood. These local helical curve approximations will later be used within the relaxation labeling framework, described in Section 2.3, in order to infer the likely curves (i.e. white matter fibres) that pass through a given point in the 3D ODF dataset.

**Proposition 3.** *Given three unit vectors  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^3$  specified at three locations  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{R}^3$ , it is possible to determine whether or not  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  are co-helical. Furthermore, if they are co-helical, it is possible to recover the parameters of the helix passing through  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  and having tangent vectors at these locations that are parallel to  $\pm\mathbf{v}_1, \pm\mathbf{v}_2$  and  $\pm\mathbf{v}_3$ , respectively.*

**Proof of Proposition 3.** Although there are several details in this proof, the approach is quite intuitive. Fig. 3 serves as a visual aid to the solution. The key idea is simply to use Proposition 2 to uniquely define a putative helix axis given the three candidate helix tangents and their locations, and then to determine whether these are compatible with the properties of a helix, given the axis and the plane orthogonal to it.




Fig. 3. An illustration of the concepts used in the proof of Proposition 3, as well as in Algorithm 2.2, in order to determine co-helicity between  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$ .

A helix with an axis coinciding with the  $z$ -axis is a parametric curve defined by Eq. (1). In particular, Eq. (1) implies that the orthographic projection of the curve in the direction of the axis onto the plane normal to the axis (which coincides with the  $xy$  plane in this case) is a circle with radius  $r$ , centered at the origin. Furthermore, the parameter  $c$  determines the (constant) angle that the tangent vectors to the curve make with the axis, as well as the translation along the helix axis for a given increment of the value of parameter  $t$ .

To determine whether or not  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  are co-helical (where  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  are distinct), one needs to demonstrate that there is a curve that satisfies the above helix properties, which passes through  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  and which has tangent vectors at these locations that are parallel to  $\pm\mathbf{v}_1, \pm\mathbf{v}_2$  and  $\pm\mathbf{v}_3$ , respectively.

Let us assume that such a helix curve exists. In general, its axis will not be aligned with the  $z$ -axis, nor will it pass through the origin. We use Proposition 2 to recover two (distinct) normal vectors to the helix, from which we can recover the helix axis, which is defined as the cross-product of any two normal vectors of the helix. Once the helix axis is recovered, one can verify the co-helical properties of  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$ . That is, the projections along the axis of  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  (which we call  $\Pi(\mathbf{p}_1)$ ,  $\Pi(\mathbf{p}_2)$  and  $\Pi(\mathbf{p}_3)$ , respectively) onto  $\Pi$  should lie on a circle,  $C_\Pi$ . The center  $O$  of  $C_\Pi$  is a point through which the helix axis passes (see Fig. 3).

Consider three unit-length line segments  $\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3$ , defined to be parallel to  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ , respectively, and to be located at  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$ , respectively. The projections of  $\mathbf{l}_1, \mathbf{l}_2$  and  $\mathbf{l}_3$  onto  $\Pi$  should be tangent to  $C_\Pi$ . As well,  $\mathbf{l}_1, \mathbf{l}_2$  and  $\mathbf{l}_3$  should make a constant angle with the helix axis. Finally, consider a new coordinate system where the  $z$ -axis is defined to correspond to the helix axis. The difference between the  $z$ -coordinates of  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  in this new coordinate system should be consistent with both the (constant) angle between each of  $\mathbf{l}_1, \mathbf{l}_2$  and  $\mathbf{l}_3$  and the helix axis, as well as the angular separation of  $\Pi(\mathbf{p}_1)$ ,  $\Pi(\mathbf{p}_2)$  and  $\Pi(\mathbf{p}_3)$  along  $C_\Pi$  (which is equivalent to the parameter  $t$  in Eq. (1), when  $0 \leq t < 2\pi$ ).

By the definition of co-helicity, if  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  are co-helical then the above properties will be satisfied. Conversely, if they are not co-helical, then one or more of these properties will not hold.  $\square$

We now present an algorithm (Algorithm 2.2) for determining whether or not  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  in  $\mathbb{R}^3$  are co-helical, and if so, for recovering the parameters of the helix. To speed-up the process, one can use Proposition 1 in a preprocessing step to reject non-co-helical triplets.

**Algorithm 2.2.** Determining co-helicity between  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$ .

**Data:** Three vectors  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{R}^3$  at three locations  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3 \in \mathbb{R}^3$ .

**Result:** 1 if  $(\mathbf{v}_1, \mathbf{p}_1)$ ,  $(\mathbf{v}_2, \mathbf{p}_2)$  and  $(\mathbf{v}_3, \mathbf{p}_3)$  are co-helical, 0 otherwise.

- (1) Define  $\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3$  to be unit-length line segments parallel to  $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ , respectively, located at  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$ , respectively.
- (2) (Testing for a special case) If  $\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3, (\mathbf{p}_2 - \mathbf{p}_1), (\mathbf{p}_3 - \mathbf{p}_2)$  are all collinear, **Return 1**.
- (3) For  $i = 1, 2, 3$ , define  $\mathbf{v}_i^{\text{temp}}$  to be either  $\mathbf{v}_i$  or  $-\mathbf{v}_i$ , so that the difference vectors  $\mathbf{v}_{\Delta 1} = \mathbf{v}_2^{\text{temp}} - \mathbf{v}_1^{\text{temp}}$  and  $\mathbf{v}_{\Delta 2} = \mathbf{v}_3^{\text{temp}} - \mathbf{v}_2^{\text{temp}}$  are orthogonal to vectors  $(\mathbf{p}_2 - \mathbf{p}_1)$  and  $(\mathbf{p}_3 - \mathbf{p}_2)$ , respectively. Define a vector  $\mathbf{n}$  as the cross-product of vectors  $\mathbf{v}_{\Delta 1}$  and  $\mathbf{v}_{\Delta 2}$ . By Proposition 2,  $\mathbf{n}$  is parallel to the axis of the putative helix. For convenience, we assume from now on that  $\mathbf{n}$  is normalized to unit length.
- (4) Compute the projections  $\Pi(\mathbf{l}_1), \Pi(\mathbf{l}_2)$  and  $\Pi(\mathbf{l}_3)$  of  $\mathbf{l}_1, \mathbf{l}_2, \mathbf{l}_3$  onto  $\Pi$ , the plane that has  $\mathbf{n}$  as its normal vector, in the direction  $\mathbf{n}$ . Without loss of generality, we fix  $\Pi$  to be the plane that contains  $\mathbf{p}_1$ . Compute the projections  $\Pi(\mathbf{p}_2)$  and  $\Pi(\mathbf{p}_3)$  of  $\mathbf{p}_2$  and  $\mathbf{p}_3$  onto  $\Pi$  in the direction  $\mathbf{n}$ . By construction,  $\Pi(\mathbf{p}_1) = \mathbf{p}_1$ . These projections fix the value of the parameter  $r$  in the helix equation (Eq. (1)). That is,  $r$  is the radius of the circle  $C_\Pi$ , lying in  $\Pi$ , passing through  $\Pi(\mathbf{p}_1), \Pi(\mathbf{p}_2)$ , and  $\Pi(\mathbf{p}_3)$ , and to which  $\Pi(\mathbf{l}_1), \Pi(\mathbf{l}_2)$  and  $\Pi(\mathbf{l}_3)$  are all tangent.
- (5) The center  $O$  of  $C_\Pi$  is determined by intersecting the three lines lying in  $\Pi$  that are normal to  $\Pi(\mathbf{l}_1), \Pi(\mathbf{l}_2)$  and  $\Pi(\mathbf{l}_3)$ . Without loss of generality one can set the parameter  $t$  (see Eq. (1)) corresponding to  $\mathbf{p}_1$  to  $t_1 = 0$ . The parameter  $t_2$  of  $\mathbf{p}_2$  can then be set to the angle between the line  $O - \Pi(\mathbf{p}_1)$  and the line  $O - \Pi(\mathbf{p}_2)$ . The parameter  $t_3$  of  $\mathbf{p}_3$  can be set in a similar way to the angle between lines  $O - \Pi(\mathbf{p}_2)$  and  $O - \Pi(\mathbf{p}_3)$  (see Fig. 3).
- (6) One should then verify that  $\mathbf{l}_1, \mathbf{l}_2$  and  $\mathbf{l}_3$  make the same angle with the helix axis  $\mathbf{n}$ , and that this angle is consistent with the values for  $r, t_2$  and  $t_3$  determined previously. In particular, let

$$\phi_2 = \frac{\pi}{2} - \arccos(\min(\langle \mathbf{v}_2, \mathbf{n} \rangle, \langle -\mathbf{v}_2, \mathbf{n} \rangle)).$$

That is,  $\phi_2$  is the angle between  $\mathbf{l}_2$  and the plane  $\Pi$ .  $\phi_1$  and  $\phi_3$  are defined in a similar manner using  $\mathbf{v}_1$  and  $\mathbf{v}_3$ , respectively. The equality  $\phi_1 = \phi_2 = \phi_3$  should hold. Now let  $\zeta_2 = \langle \mathbf{p}_2, \mathbf{n} \rangle$ .  $\zeta_1$  and  $\zeta_3$  are defined in an equivalent manner using  $\mathbf{p}_1$  and  $\mathbf{p}_3$ , respectively. In other words,  $\zeta_1, \zeta_2$  and  $\zeta_3$  represent the projection of  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  onto the helix axis, or equivalently, the  $z$ -coordinates of  $\mathbf{p}_1, \mathbf{p}_2$  and  $\mathbf{p}_3$  in a new coordinate system where  $\mathbf{n}$  gives the  $z$ -axis. Letting  $\phi = \phi_1 = \phi_2 = \phi_3$ , the following equalities should hold:

$$\tan(\phi) = \frac{\zeta_3 - \zeta_1}{r(t_2 + t_3)},$$

$$\tan(\phi) = \frac{\zeta_3 - \zeta_2}{rt_3},$$

$$\tan(\phi) = \frac{\zeta_2 - \zeta_1}{rt_2}.$$

- (7) If one or more of the above steps fail (e.g. there is no circle to which  $\Pi(\mathbf{l}_1), \Pi(\mathbf{l}_2)$  and  $\Pi(\mathbf{l}_3)$  are all tangent in step (4), or the equations in step (6) do not hold),

**Return 0.**

Else define the value of parameter  $c$  in the helix equation (1) by setting  $|c| = r \tan(\phi)$ . The sign of  $c$  is determined by the direction in which the helix winds around the axis - clockwise or anticlockwise. Equivalently, it can be set to the sign of  $\langle \mathbf{n}, \mathbf{v}_2 \rangle$ . The curvature of the helix segment is given by  $\kappa = r/(r^2 + c^2)$  and the torsion is given by  $\tau = c/(c^2 + r^2)$ ,

**Return 1.**

### 2.3. Relaxation labeling

Following Parent and Zucker (1989), we pose the problem of 3D curve inference as that of assigning confidence in a set of labels associated with nodes in a graph, using relaxation labeling (Rosenfeld et al., 1976; Hummel and Zucker, 1983).

Let the graph consist of  $n$  nodes (i.e. 3D voxels), each with an orientation distribution function  $f(\theta, \phi)$ . Edges are placed between any two nodes that interact, i.e. all such pairs are viewed as neighbors.  $f$  is sampled so as to achieve orientation likelihood estimates in  $m$  orientations, using a uniform sampling of the surface of a sphere. The value of  $m$  will depend on the desired angular resolution in the representation of the data. These  $m$  orientations are associated with the labels  $\lambda$ . A quantity  $p_i(\lambda)$  is assigned to label  $\lambda$  at node  $i$  to indicate the confidence in that particular label existing at node  $i$ . The value of this assignment is determined by

$$p_i(\lambda) = \frac{f(\theta_\lambda, \phi_\lambda)}{\sum_{k=1}^m f(\theta_k, \phi_k)}, \quad (8)$$

so that  $\forall i, \sum_{k=1}^m p_i(\lambda) = 1$  and  $0 \leq p_i(\lambda) \leq 1$ .

Following Rosenfeld et al. (1976) and Hummel and Zucker (1983), this yields a space of weighted labeling assignments  $\mathbf{K}$ , given by

$$\mathbf{K} = \mathbf{P} \in \mathbb{R}^m : \mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_n); \quad (9)$$

$$\mathbf{p}_i = (p_i(1), \dots, p_i(m)) \in \mathbb{R}^m. \quad (10)$$

In Parent and Zucker's 2D curve inference framework, a labeling  $\mathbf{P}$  provides support  $s_i(\lambda)$  for label  $\lambda$  at node  $i$  given by

$$s_i(\lambda) = \sum_{j=1}^n \sum_{\lambda'=1}^m r_{ij}(\lambda, \lambda') p_j(\lambda'), \quad (11)$$

where  $\lambda$  and  $\lambda'$  are labels associated with nodes  $i$  and  $j$ , respectively, and where the  $r_{ij}(\lambda, \lambda')$ 's are the elements of a matrix of compatibilities between pairs of labels, as defined through a co-circularity constraint (Parent and Zucker, 1989). In our work, we use a co-helicity constraint defined between triplets of labels. Thus, the support function must be extended to use a higher-order compatibility (see Hummel and Zucker, 1983), given by

$$s_i(\lambda) = \sum_{j=1}^n \sum_{\lambda'=1}^m \sum_{k=1}^n \sum_{\lambda''=1}^m r_{ijk}(\lambda, \lambda', \lambda'') p_j(\lambda') p_k(\lambda''). \quad (12)$$

Here  $r_{ijk}(\lambda, \lambda', \lambda'')$  represents the compatibility (co-helicity) between orientation  $\lambda$  at node  $i$ , orientation  $\lambda'$  at node  $j$  and orientation  $\lambda''$  at node  $k$ , as determined using Algorithm 2.2.

The goal of relaxation labeling is to find a consistent labeling  $\mathbf{V} \in \mathbf{K}$ , which is equivalent to solving a variational inequality (Hummel and Zucker, 1983). In the case of compatibility coefficients  $r_{ijk}$  measured between triplets of labels  $\lambda$ ,  $\lambda'$ , and  $\lambda''$ , a labeling  $\mathbf{V}$  is consistent if and only if

$$\sum_{i=1}^n \sum_{\lambda=1}^m \sum_{j=1}^n \sum_{\lambda'=1}^m \sum_{k=1}^n \sum_{\lambda''=1}^m r_{ijk}(\lambda, \lambda', \lambda'') p_j(\lambda') p_k(\lambda'') [p_i(\lambda) - v_i(\lambda)] \leq 0, \quad (13)$$

for all  $\mathbf{P} \in \mathbf{K}$  (Hummel and Zucker, 1983). The meaning of this inequality is that a consistent labeling  $\mathbf{V}$  must have a total support greater or equal to that of any other possible labeling  $\mathbf{P}$ . Hummel and Zucker (1983) show that inequality (13) can be solved using gradient ascent on the *average local support function*  $A(\mathbf{P})$

$$A(\mathbf{P}) = \sum_{i=1}^n \sum_{\lambda=1}^m p_i(\lambda) s_i(\lambda). \quad (14)$$

They prove that if the matrix of compatibilities follows certain symmetry conditions (which can be shown to hold for our co-helicity compatibilities, determined through Algorithm 2.2), and if  $A(\mathbf{P})$  attains a local maximum at  $\mathbf{V} \in \mathbf{K}$ , then  $\mathbf{V}$  is a consistent labeling. Hence, for symmetric compatibilities the problem of finding consistent labelings reduces to that of finding local maxima of the objective function  $A(\mathbf{P})$ . The method for finding local maxima of  $A(\mathbf{P})$  follows Algorithm 8.2 in Hummel and Zucker (1983), as well as the radial projection method described in Appendix A of Parent and Zucker (1989). Algorithm 8.2 is an iterative one, and as the iterations progress, the labeling of the network approaches consistency, i.e. the ODFs that the network represents become

more and more regularized. In practice it may not be necessary to achieve a consistent labeling for the ODFs to be sufficiently regularized. In fact, Algorithm 8.2 can be used as an *any-time* algorithm that can be stopped after any number of iterations have completed, once a desired level of regularization is achieved. In Section 5, we address the issue of how to determine appropriate amounts of regularization.

### 3. Implementation issues

#### 3.1. Quantization of location, curvature and torsion

The equations used throughout Section 2 must be applied in the context of a discretized 3D volume, with discretized orientations at each location. In Section 2.3, we examined the discretization of the space of allowable orientations into a set of "labels". Here, we discuss quantization of location, as well as the quantization of curvature and torsion.

We treat each orientation estimate (label) in a given voxel (voxels are assumed to be isotropic, i.e. cubic), as being localized anywhere within a sphere of radius 1/2 the edge of a voxel, centered at that voxel. This has implications for the implementation of various parts of Algorithm 2.2. As one example (there are many others), a pair of orientations, one in voxel  $i$  and one in voxel  $j$  should be considered to satisfy Proposition 1 if the difference in the angle they make with the line joining the centers of the voxels is less than  $\theta = \arcsin(1/d_{ij})$ , where  $d_{ij}$  represents the Euclidean distance between the centers of voxels  $i$  and  $j$ , following an argument very similar to that presented in Parent and Zucker (1989).

Following Parent and Zucker (1989), we implement a discretization of the allowed range of variation of curvature and torsion (the 3D extension of the curvature classes in Parent and Zucker (1989)) in order to be able to distinguish between distinct curves that share a tangent. We quantize the range between the maximum and the minimum admissible curvature values into a set of intervals (typically three to five), and the range of admissible torsion values is discretized in the same manner.

Fig. 4 (left) shows an example of this discretization for a torsion value of 0, and for a certain range of curvature values. It is also clear from this figure that there can be an infinite number of helix curves with the same curvature and torsion that have the central orientation as a tangent, because of the additional degree of freedom given by rotation around that tangent, in the plane normal to the tangent. Thus, in addition to discretization based on curvature and torsion values, it is necessary to perform a discretization in the plane normal to the given orientation, at equal angular intervals, as shown in Fig. 4 (right). Each element of this discretization is referred to as a curvature-torsion-normal class, or simply a class.

An orientation is considered to be a member of a class when it receives maximum support through co-helical




Fig. 4. Left: An example of a discretization of curvature and torsion around an orientation, as defined using a minimum and a maximum curvature value and a torsion value of 0. Right: To achieve a meaningful discretization, one needs to perform an additional subdivision at equal angular intervals in the plane orthogonal to the orientation. In this example, a set of curvature-torsion-normal classes are visualized, for a given range of curvature values and a torsion value of 0.

configurations within that class. To determine class membership for a given orientation, in a manner analogous to the methodology of Parent and Zucker (1989), we measure support independently by curvature-torsion-normal class. To do so, given a total of  $C$  curvature-torsion-normal classes, we modify the support function (12) as follows:

$$s_i(\lambda) = \max_{c=1,\dots,C} \sum_{j=1}^n \sum_{\lambda'=1}^m \sum_{k=1}^n \sum_{\lambda''=1}^m Q_{ijk}^c(\lambda, \lambda', \lambda'') r_{ijk}(\lambda, \lambda', \lambda'') p_j(\lambda') p_k(\lambda''). \quad (15)$$

Here  $Q_{ijk}^c(\lambda, \lambda', \lambda'')$  is a binary-valued partition function such that its value is 1 if label  $\lambda$  at node  $i$ , label  $\lambda'$  at node  $j$  and label  $\lambda''$  at node  $k$  are co-helical with a helix that has parameters (as determined through Algorithm 2.2) that lie within the boundaries of curvature-torsion-normal class  $c$ , and is 0 otherwise. The class number  $c$  that maximizes (15) is recorded as the class which label  $\lambda$  at node  $i$  is a member of.

This notion of class membership allows the extension to 3D of the *curvature consistency principle* of Parent and Zucker (1989). Since in principle it is possible for three labels that belong to separate curves to be in a co-helical configuration, it is desirable that only those co-helical labels that belong to the same curve provide support to each other. Stated in another way, one would like to ensure that co-helical labels that belong in fact to distinct curves do not provide support to each other. To do so, we apply the relaxation labeling algorithm discussed in Section 2.3, with the constraint that *each label  $\lambda$  will accumulate support according to Eq. (12) only if  $\lambda'$  and  $\lambda''$  are members of the same class as  $\lambda$* . This is the 3D extension of the curvature consistency principle introduced in Parent and Zucker (1989). We refer to this extension as the *3D curve consistency principle*.

### 3.2. Computational efficiency

Determining co-helical configurations of triplets is computationally expensive, even in small volumes. We obtain a significant increase in efficiency by precomputing co-helical configurations of triplets within a spherical neighborhood of a given radius, and storing these in a look-up table. This operation is carried out only once, offline, for a particular choice of parameters. The parameters, which influence the running time of the algorithm, include the number of orientations modeled at each location (the size of the label set), the choice of the discretization parameters, i.e. the number of different classes, and the diameter of the spherical neighborhood in which all pairs of voxels are considered neighbors. As an example, with a label set of 100 orientations, a neighborhood diameter of nine voxels, and a discretization using five curvature intervals, three torsion intervals and four normal rotation intervals, precomputing co-helical configurations takes about 8 h on a 3.6 GHz Intel processor running Linux. Once co-helical configurations have been computed, the running time of the algorithm is linear in the number of voxels in the volume. With volume grids of  $128 \times 128 \times 63$  voxels, using the same processor we typically obtain satisfactory regularization results within 6 h of dedicated running time.

## 4. Experimental validation

In this section, we present a series of validation experiments to demonstrate the performance of 3D curve inference. We provide qualitative and quantitative validation of our algorithm using a biological phantom, as well as a synthetic DTI dataset. We also provide regularization results using DTI data and QBI reconstructions from HARDI measurements acquired in vivo from a human brain. In particular, we show that regularizing the brain DTI data results in ODFs that are qualitatively very similar to those obtained with a QBI reconstruction (Tuch et al., 2003) from the HARDI data. We also demonstrate that 3D curve inference, as a preprocessing step, significantly improves the performance of fibre tracking in the phantom as well as in the brain.

The 3D curve inference algorithm can use an arbitrarily large label set. In our implementation we use a label set of 100 unit direction vectors distributed isotropically over a hemisphere, obtained using an electrostatic charge repulsion algorithm (Jones et al., 1999). We use a spherical neighborhood with a diameter of nine voxels in all our experiments.

### 4.1. Quantitative validation on a biological phantom

A biological phantom (introduced in Campbell et al., 2005) was created from two excised Sprague–Dawley rat spinal cords embedded in 2% agar. The cords were approximately 12.5 cm in length and 5 mm in diameter. Two diffusion-weighted datasets were acquired using this phantom, with 90 diffusion encoding directions, with  $b$  values of



3000 s/mm<sup>2</sup> and 1300 s/mm<sup>2</sup>, respectively, as described in Campbell et al. (2005). The first was used for high angular resolution reconstruction with the QBI technique (Tuch et al., 2003). The second was used for diffusion tensor reconstruction of the diffusion ODF, using the standard 3D Gaussian model for the diffusion PDF (Basser et al., 1994):

$$P(r|\tau_d) = \frac{1}{\sqrt{|\mathbf{D}|}(4\pi\tau_d)^3} \exp\left(\frac{-\mathbf{r}^T\mathbf{D}^{-1}\mathbf{r}}{4\tau_d}\right). \quad (16)$$

Here  $P(r|\tau_d)$  is the probability that a water molecule displaces by vector  $\mathbf{r}$  during the time  $\tau_d$  over which diffusion is observed.  $\mathbf{D}$  is the diffusion tensor, a  $3 \times 3$  positive semi-definite matrix whose eigenvalues give the variances of the 3D Gaussian model. The diffusion ODF is then defined as the normalized projection of the 3D diffusion PDF on the surface of the unit sphere:

$$\text{ODF}(\mathbf{u}) = \frac{\text{ODF}^*(\mathbf{u})}{\text{ODF}_{\text{mean}}^*}, \quad (17)$$

$$\text{ODF}^*(\mathbf{u}) = \int_0^\infty P(r\mathbf{u}|\tau_d) dr, \quad (18)$$

where  $\mathbf{u} = (1, \theta, \phi)$  is a unit vector from the origin to the surface of the unit sphere and  $\text{ODF}_{\text{mean}}^*$  is the mean value of  $\text{ODF}^*(\mathbf{u})$  over all values of  $\mathbf{u}$ .

A T1-weighted image of this phantom is shown in Fig. 5 (top left). The ground truth orientations were determined by extracting the centerlines of each cord using the technique of Bouix et al. (2005) and then smoothly extending the orientations in the centre to the boundary of the cord, for each cord (shown separately in Fig. 5, top right). For numerical accuracy we carried this out on a super-sampled version of the original data and then sub-sampled the result. The second row of Fig. 5 shows the DTI reconstruction of the ODFs in the region of the crossing, indicated with a white box in Fig. 5 (top left), prior to regularization (left) and after regularization through the 3D curve inference algorithm (right). For ease of visualization, the unregularized DTI reconstruction is shown after subtraction of the MIS (all subsequent figures showing unregularized ODF sets will also be visualized using MIS subtraction). One can see that the ODFs in the crossing region are fairly isotropic (and thus MIS subtraction leaves little volume). After the application of 3D curve inference, however, the ODFs in the voxels of the crossing region show well-defined maxima oriented along both of the crossing fibres. The third row of Fig. 5 shows, in a similar fashion, the result of applying 3D curve inference (right) on the QBI reconstruction of the ODFs, shown prior to regularization (left). Although there is some evidence for multiple maxima in the QBI reconstruction in the voxels of the crossing, the regularization result brings them out in a much clearer way. Finally, the fourth row presents results obtained using the orthonormal vector regularization algorithm of Tschumperlé and Deriche (2002). The unregularized principal directions of the DTI ODFs are shown on the left and

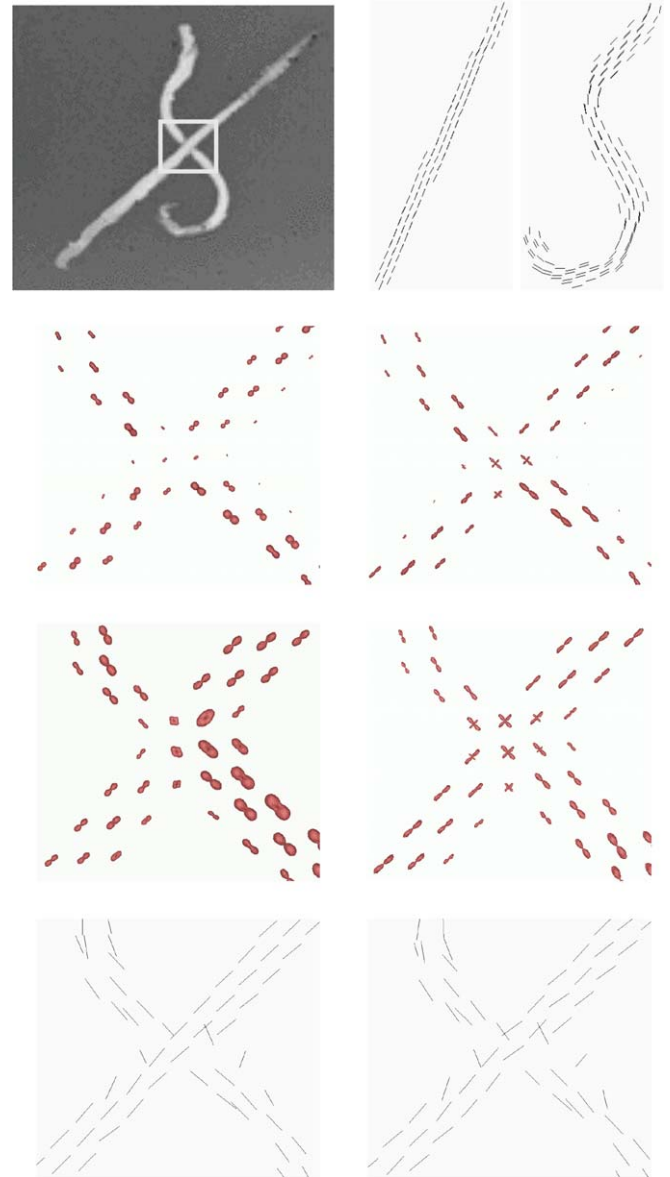


Fig. 5. First row: A biological phantom created by overlaying two rat cord spinal cords (left) and ground truth local fibre orientations shown separately for the two cords (right). Second row: Unregularized DTI reconstruction of the ODFs in the area of the crossing (left) and the same set of ODFs regularized using 3D curve inference (right). Third row: The unregularized QBI reconstruction from HARDI measurements in the vicinity of the crossing (left) and the regularized ODFs using 3D curve inference (right). Fourth row: The principal eigenvector orientations of the DTI dataset in the vicinity of the crossing (left) and the regularized orientations obtained using the technique of Tschumperlé and Deriche (2002) (right).

the regularized results are shown on the right. Such an algorithm cannot handle more than one ODF maximum at each voxel, and is furthermore hindered by the sparseness of the data.

In the regularization experiments on the phantom data, 3D curve inference was run with five curvature intervals, covering the range of curvature radii between  $\infty$  (for 0 curvature) and 1.05 voxel units. In the QBI dataset, one voxel

unit corresponds to 2.8 mm, and in the DTI dataset, to 2.5 mm. There were three torsion quantizations, covering torsion radii between  $\infty$  and 4.6 voxel units (as well as between  $-\infty$  and  $-4.6$ , since torsion can be negative) and the normal plane (see Fig. 4 and Section 3.1) was quantized into four quadrants, each covering  $\pi/2$  rad. Thus, there were a total of  $5 \times 3 \times 4 = 60$  curvature-torsion-normal classes.

To perform a quantitative evaluation of the 3D curve inference algorithm, we use the ground truth orientations described above. For each ODF dataset, regularized or not, the ODF maxima are extracted. Then, at each voxel, the smallest angular difference (error) between the available maximum (or maxima) and ground truth orientation(s) at that voxel is recorded. The median and mean ( $\pm 1$  standard deviation) orientation errors in degrees are shown for the ground truth dataset and the unregularized as well as regularized datasets in the left column of the table in Fig. 6 (top). Observe that 3D curve inference yields results with significantly lower mean errors, in particular when applied to QBI data. It is important to note that due to the discrete sampling of the sphere, even perfect ODF data will be expected to have some error with respect to the ground truth orientations. This minimal expected error is related to the solid angle subtended by one facet of the sphere tessellation induced by the sampling. For example, it can be found to equal  $7.2^\circ$  for the uniform sampling of the hemisphere by 100 directions used in our experiments.

	Phantom Median, Mean $\pm$ Std. Dev.	Synthetic Median, Mean $\pm$ Std. Dev.
Unreg. DTI	15.2°, 19.4 $\pm$ 16.2°	22.0°, 23.2 $\pm$ 12.0°
CI-DTI	12.5°, 16.0 $\pm$ 15.3°	8.5°, 11.1 $\pm$ 9.4°
OVR-DTI	13.2°, 18.2 $\pm$ 16.3°	19.4°, 20.7 $\pm$ 11.9.°
Unreg. QBI	12.8°, 17.2 $\pm$ 16.0°	N/A
CI-QBI	11.4°, 12.2 $\pm$ 7.8°	N/A




Fig. 6. Top: Table of validation results showing median and mean  $\pm 1$  standard deviation orientation errors in degrees for the biological phantom data set and the synthetic data set. The row headings read as follows. Unreg. DTI: unregularized DTI. CI-DTI: 3D curve inference applied to the DTI dataset. OVR-DTI: the orthonormal vector regularization of Tschumperlé and Deriche (2002) applied to the DTI dataset. Unreg. QBI: unregularized QBI reconstructions (Tuch et al., 2003). CI-QBI: 3D curve inference applied to the QBI dataset. Bottom: A snapshot of the noisy synthetic data set, prior to regularization.

#### 4.2. Quantitative validation on synthetic data

A synthetic DTI dataset in a  $100 \times 50 \times 100$  voxel grid was created by placing anisotropic diffusion tensors with their principal direction vector aligned with one of three curves: a planar sine wave, and two helices with different curvature and torsion. Partial volume averaging effects were simulated in voxels where the helices intersected the sine wave, and background voxels were filled with isotropic (spherical) tensors. Since 3D curve inference depends only on the relative shape of the ODFs, the mean eigenvalue of all tensors was arbitrarily set to 3. In anisotropic regions the principal eigenvalue was set to 7 and the others to 1. In voxels with crossings, the two principal directions had eigenvalues of 4, and the other eigenvalue was set to 1. The two angles describing the orientation of each tensor were independently perturbed by adding Gaussian noise, with mean 0 and a standard deviation of  $\pm 22.9^\circ$  (0.4 rad). The original noiseless dataset was treated as the ground truth. One view of the noisy dataset is shown in Fig. 6 (bottom); it is important to note that the helices are non-planar curves. Validation results are shown in the right column of the table in Fig. 6 (top). Observe that once again, 3D curve inference achieves a significant reduction in orientation error, compared to both the noisy unregularized data, as well as the result obtained with the regularization method of Tschumperlé and Deriche (2002).

In this experiment, 3D curve inference was run with a total of 48 curvature-torsion-normal classes, with four curvature quantizations covering curvature radii between  $\infty$  and 2.5 (in voxel units), three torsion intervals covering torsion radii between  $\infty$  and 4.4 voxel units (as well as between  $-\infty$  and  $-4.4$ ) and the normal plane was quantized into four quadrants of  $\pi/2$  rad each (see Section 3.1).

#### 4.3. Qualitative validation on human brain data

We now present regularization results using in vivo human brain data. Human brain diffusion weighted images were acquired using a Siemens 3T Trio MR scanner. The acquisition used 99 diffusion encoding directions, 2 mm isotropic voxel size, 63 slices, and  $b = 3000$  s/mm<sup>2</sup>. Both the DTI and the QBI ODF datasets were calculated using all of the data. Again, the QBI reconstruction was carried out according to the technique of Tuch et al. (2003) and the DTI reconstruction was carried out with the Gaussian model described in Section 4.1. A T1-weighted anatomical scan was also acquired. Fig. 7 shows the principal diffusion direction RGB map overlaid on the T1-weighted anatomical image for two different regions of interest (ROIs) indicated with white contours, for which regularization results are presented below.

Fig. 8 shows the DTI ODFs, the regularized DTI ODFs, the QBI ODFs, and the regularized QBI ODFs in an ROI containing fibres from the cortical spinal tract and the corpus callosum. The ROI is indicated with a white square in Fig. 7 (left). This region has fibres that are in a crossing




Fig. 7. The principal diffusion direction RGB map shown overlaid on the T1-weighted anatomical image, together with a white box indicating the ROIs used for Fig. 8 (left) and Fig. 9 (right).

and/or a “neck”-like configuration. Another ODF comparison is shown in Fig. 9, in a region of partial volume averaging of the cingulum and the corpus callosum, indicated with a white box in Fig. 7 (right). In this area it is known from anatomy (and can also be seen from the RGB directions map in Fig. 7 (right)) that the fibres of the cingulum are approximately orthogonal to those of the corpus callosum. Thus, at the interface of the two fibre tract systems, where partial volume averaging occurs, ODFs in the shape of a cross can be seen in the QBI reconstruction in Fig. 9, third row.

In Figs. 8 and 9, the unregularized DTI reconstruction is ambiguous in voxels where partial volume averaging of directions occurs. In contrast, both the QBI and the regularized DTI ODFs exhibit multiple maxima in these voxels that are qualitatively very similar to each other. Finally,




Fig. 8. ODFs obtained using DTI (top left), regularized DTI (top right), QBI (bottom left), and regularized QBI (bottom right) in a region of partial volume averaging of the cortical spinal tract and the corpus callosum. The ROI is indicated by the white box in Fig. 7 (left).




Fig. 9. ODFs obtained using DTI (top row), regularized DTI (second row), QBI (third row), and regularized QBI (bottom row) in the region of partial volume averaging of the cingulum and the corpus callosum indicated with a white box in Fig. 7 (right). The ODFs are displayed from an oblique point of view, so that crossing configurations can be clearly seen. Thus, ODF maxima that have an up-down orientation are associated with the corpus callosum.

one can see that as with the phantom data, 3D curve inference sharpens the QBI ODFs and makes the maxima salient.

In all experiments with brain data, 3D curve inference was run with a total of 48 curvature-torsion-normal classes, with four curvature quantizations covering curvature radii between  $\infty$  and 2.5 (in voxel units), three torsion intervals covering torsion radii between  $\infty$  and 4.4 voxel units (as well as between  $-\infty$  and  $-4.4$ ) and the normal plane was quantized into four quadrants of  $\pi/2$  rad each (see Section 3.1). In our brain data, one voxel unit corresponds to 2 mm.

#### 4.4. Applications of 3D curve inference to fibre tracking

We now present experiments which apply 3D curve inference as a preprocessing step to improve the performance of a fibre tracking algorithm. As stated previously in Section 1, the output of 3D curve inference is not a diffusion ODF but rather an ODF that encodes a confidence in fibre tracts being present with a given orientation. In this setting, the resulting ODF maxima are the best estimates available to guide a fibre tracking algorithm. Thus a streamline tracking method on the ODF maxima is used, as opposed to a more general diffusion ODF tracking technique. Recent work by Lazar and Alexander (2003) compares different streamline tracking techniques and shows that the fibre assignment by continuous tracking (FACT) method of Mori et al. (1999) has higher precision in synthetic divergent diffusion tensor fields, as well as superior accuracy in synthetic curved fields. Given that both divergent and curved fields are expected in the brain, streamline fibre tracking with FACT integration is a reasonable choice for a representative line propagation technique.

We begin with tracking results on the biological phantom. The FACT method (Mori et al., 1999) was initiated in a seed ROI spanning a cross-section of the curved cord near its end. All maxima of the diffusion ODF that produced curves with radius of curvature 2.5 mm or greater were followed (directions that generated paths with a smaller radius of curvature were assumed to be crossing fibres and were not followed). Tracking was constrained to voxels with fractional anisotropy (FA) of 0.05 or greater. Fig. 10 shows tracking results on unregularized data (top row) and regularized data (bottom row), for the case of DTI (left column) and QBI (right column). Transparent surfaces corresponding to the two different cords are shown for reference. In the unregularized DTI case, tracking cannot continue past the crossing. In the unregularized QBI case, the tracking passes through the crossing, but has difficulties following the cord to its end. Tracking on regularized DTI not only passes the crossing, but also reaches the end of the cord. Tracking on regularized QBI gives the result that agrees most closely with the curvature of the cord.

We conclude this section with a fibre tracking experiment in the brain, once again using streamline fibre tracking with FACT integration (Mori et al., 1999) but starting from a small seed region near the cortex. The results are

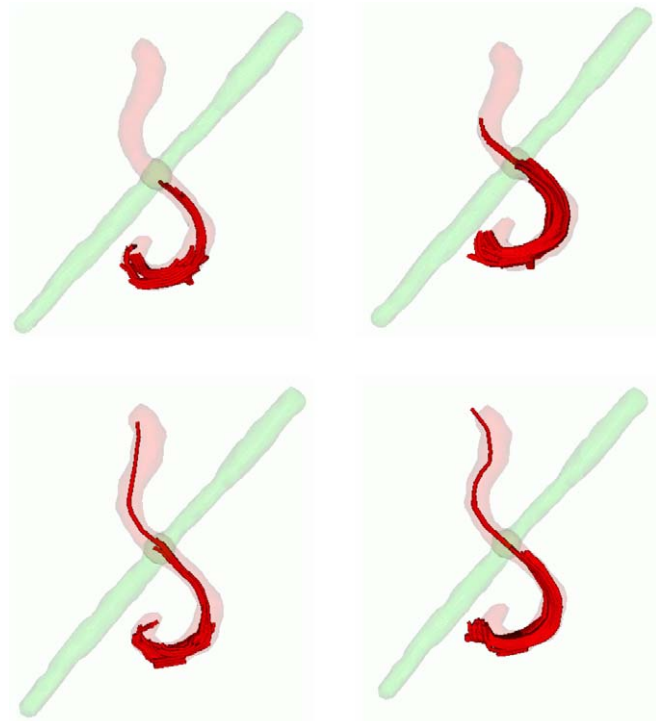


Fig. 10. Tractography results on the phantom, using unregularized DTI (top left), unregularized QBI (top right), regularized DTI (bottom left) and regularized QBI (bottom right). A transparent surface indicating the segmented cords is shown for reference.

presented in Fig. 11. If the QBI tract reconstructions are used as a qualitative “ground truth”, the advantage of using regularized DTI over DTI is obvious. The DTI, QBI and regularized DTI tract reconstructions indicate U-fibre structure near the cortex. The QBI and regularized DTI tract reconstructions also indicate pathways from the seed region to the cortical spinal tract and corpus callosum. It is probable that fibres from all of these pathways pass through the seed region. The high angular resolution of the QBI reconstruction and the inferred high angular resolution of the regularized DTI reconstruction allow the tracking algorithm to branch in regions where multiple fibre orientations exist, giving more accurate information about the connectivity.

## 5. Discussion

We now address certain aspects of our experimental results and propose directions in which to extend the 3D curve inference algorithm.

### 5.1. Fibre tracking experiments

The results presented in the previous section demonstrate the ability of 3D curve inference to infer high-angular resolution ODFs from DTI data, which has a significant impact on the performance of fibre tracking algorithms. In fact, in the case of the phantom data, regularization using




Fig. 11. Tractography results showing connections that pass through a small seed region near the cortex, obtained using streamline fibre tracking with FACT integration (Mori et al., 1999). The ODFs used to drive the tractography were top: DTI, centre: QBI, and bottom: regularized DTI. A sagittal view is shown at left and an axial view at right. The T1-weighted anatomical image is shown for reference.

3D curve inference improves tracking even on the QBI dataset. It is quite possible that streamline fibre tracking with FACT integration (Mori et al., 1999) achieves better results on both the regularized DTI and regularized QBI phantom datasets than on the (unregularized) QBI phantom dataset, because in a voxel where a single curving fibre passes, the unregularized QBI ODF may not have a well-defined maximum due to subvoxel averaging of directions. As a consequence, the ODF maxima selection process may be ambiguous, particularly if noise is present, and may cause the streamline propagation procedure to stop. In this situation regularization would have the effect of sharpening the ODF and of enhancing a maximum in the direction tangent to the curving fibre, thus helping the tracking algorithm.

The tracking experiments in the unregularized brain DTI data (Fig. 11, top row) can recover only the U-fibre, whereas tracking in the QBI reconstruction also recovers branches down the cortico-spinal tract, as well as a branch

along the corpus callosum to the cortex of the other hemisphere. The improvement brought by 3D curve inference is that in the regularized DTI, tracking also recovers branches down the cortico-spinal tract and a branch along the corpus callosum to the other hemisphere. Of course, the result is not identical to the one obtained on QBI data. If one looks closely at Fig. 11 (bottom row), it appears that the U-fibre has been replaced by a fibre that takes a sharp (and thus improbable) turn that nevertheless follows a trajectory similar to it. Furthermore, in comparison to the QBI tracking result, there seem to be a few spurious paths that have been introduced in the tracking result on regularized DTI. These are probably effects of over-regularization of the DTI data. We have empirically determined that for improved results, different regions of the brain require different amounts of regularization, as well as different parameter values, for example different neighborhood sizes. For instance, the amount of regularization that is necessary to recover multiple maxima in a crossing or

branching region may be too much for a region with a U-fibre. Furthermore, the neighborhood size should ideally be chosen so as to be at the same scale as the structures to be recovered. For example, a larger neighborhood size would be more appropriate for the corpus callosum region than for regions near the cortex, where U-fibres occur. The neighborhood size determines how “local” the helix model for fibres is. For example, the helical approximations should be more localized for a U-fibre with rapidly changing curvature than for the corpus callosum. In our experiments we have applied a uniform level of regularization all over the brain, using the same neighborhood size. In future work we plan to use priors that vary throughout the brain so that the appropriate level of regularization can be applied to each region.

### 5.2. ODF sharpening

It may seem surprising at first that the regularized ODFs in the phantom data (Fig. 5) look as sharp as they do. This ODF sharpening in the regularized results is due to the fact that in the phantom, support for orientations aligned with the cords is so strong that support for all other orientations essentially collapses to 0. This is simply a result of the formulation of the relaxation labeling algorithm. Furthermore, it is important to note that regularization is performed after MIS subtraction from the original ODFs, in all experiments.

### 5.3. Synthetic data experiment

Our method for generating noisy synthetic tensor data as described in Section 4.2 is a very basic one and does not reflect the reality of MRI noise. However, it is important to remember that 3D curve inference is applicable to any ODF dataset, not just actual diffusion MRI datasets. In this sense, this experiment is still a valid demonstration of the regularization properties of 3D curve inference and is well suited for a comparison with the variational flow technique of Tschumperlé and Deriche (2002).

### 5.4. Future work

In the experiments presented in this article, we have used only the ODF shapes and their maxima for qualitative and quantitative validation, as well as for fibre tracking. In future work we will examine how discrete local curvature and torsion estimates for curves can impact the performance of fibre tracking algorithms. In particular, we will investigate the problem of distinguishing cases of complex subvoxel fibre configurations, which are known to confound fibre tracking algorithms even in QBI data. Such subvoxel configurations include curving fibres, crossings, neck-like configurations, branchings, fibre fanning and fibre splaying. Such configurations can yield ODFs with ambiguous shapes, even in QBI reconstructions. Disambiguating such configurations requires integrating information

over a local neighborhood. Knowledge of the curvature(s) and torsion(s) of the underlying curve(s) can potentially help to identify these configurations, and thus improve the performance of fibre tracking algorithms. Another foreseeable application of 3D curve inference is in the interpolation of ODF fields. Again, knowledge of the differential geometry of the underlying curves should provide constraints that will guide the interpolation of ODF data. Finally, since 3D curve inference can be applied to any ODF dataset, we intend to investigate its application to fibre ODFs (Tournier et al., 2004) in future work.

## 6. Conclusion

We have presented a differential geometric framework for regularizing diffusion MRI data, where a notion of *co-helicity* is used to compute support for orientations given other neighboring orientations. In this setting, white matter fibres are modeled as 3D curves, and the use of geometric constraints allows for the inference of local tangent, curvature and torsion estimates of curves, given an initial 3D ODF dataset.

This algorithm is applicable to DTI data, as well as to high angular resolution ODF reconstructions from HARDI measurements. We have demonstrated through experimental results several of the advantages of our algorithm. In particular, the 3D curve inference framework allows for the inference and representation of multiple ODF maxima at one location, when appropriate. As a result, regularized DTI ODFs can exhibit multiple maxima that agree qualitatively with the maxima of the QBI reconstruction (Tuch et al., 2003) from HARDI measurements. These maxima are also closer to the ground truth than the original unregularized DTI dataset in the case of the biological phantom, where ground truth is available. Furthermore, experiments in the phantom and in the brain datasets show a significant improvement of the performance of streamline fibre tracking in regularized DTI, as opposed to unregularized DTI. Experimental results also demonstrate that regularization can improve tracking in the phantom even in the case of QBI ODF reconstructions.

## Acknowledgments

We are grateful to Sylvain Bouix and Svetlana Stolpner for help with the centerline extraction algorithm, and to Vladimir V. Rymer for preparing the biological phantom. This work was supported by grants from NSERC, FQRNT, CIHR and CFI.

## References

- Alibhai, S., Zucker, S., 2000. Contour-based correspondence for stereo. In: Proceedings of the 6th European Conference on Computer Vision (ECCV), vol. 1, LNCS 1842, pp. 314–330.
- Basser, P.-J., Mattiello, J., Le Bihan, D., 1994. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal* 66, 259–267.

- Basser, P., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A., 2000. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine* 44, 625–632.
- 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. *Magnetic Resonance in Medicine* 50, 1077–1088.
- Ben-Shahar, O., Zucker, S., 2003. The perceptual organization of texture flow: a contextual inference approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25, 401–417.
- Bouix, S., Siddiqi, K., Tannenbaum, A., 2005. Flux driven automatic centerline extraction. *Medical Image Analysis* 9 (3), 209–221.
- 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.
- Chan, T., Shen, J., 1999. Variational restoration of non-flat image features: model and algorithm, Tech. Rep. CAM-TR 99-20, UCLA.
- Chefd'Hotel, C., Tschumperlé, D., Deriche, R., Faugeras, O., 2004. Regularizing flows for constrained matrix-valued images. *Journal of Mathematical Imaging and Vision* 20, 147–162.
- Chen, Y., Guo, W., Zeng, Q., Yan, Q., Huang, F., Zhang, H., He, G., Vemuri, B.-C., Liu, Y., 2004. Estimation, smoothing, and characterization of apparent diffusion coefficient profiles from high angular resolution DWI. In: *Proceedings of the IEEE CVPR*, pp. 588–593.
- 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: *Proceedings of the MICCAI 2002, LNCS 2488*, pp. 475–482.
- Coulon, O., Alexander, D., Arridge, S., 2001. A regularization scheme for diffusion tensor magnetic resonance images. In: *Proceedings of the IPMI 2001, LNCS 2082*, pp. 92–105.
- Descoteaux, M., Angelino, E., Fitzgibbons, S., Deriche, R., 2005. A linear and regularized odF estimation algorithm to recover multiple fibers in q-ball imaging, Tech. Rep. 5768, INRIA (November 2005).
- Hummel, R., Zucker, S., 1983. On the foundations of relaxation labeling processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 5, 267–287.
- Jackowski, M., Kao, C.Y., Qiu, M., Constable, R.T., Staib, L.H., 2005. White matter tractography by anisotropic wavefront evolution and diffusion tensor imaging. *Medical Image Analysis* 9, 427–440.
- Jones, D., Horsfield, M., Simmons, A., 1999. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magnetic Resonance in Medicine* 42, 515–525.
- Lazar, M., Alexander, A.L., 2003. An error analysis of white matter tractography methods: synthetic diffusion tensor field simulations. *NeuroImage* 20, 1140–1153.
- Le Bihan, D., Mangin, J.-F., Poupon, C., Clark, C., Pappata, S., Molko, N., Chabriat, H., 2001. Diffusion tensor imaging: concepts and applications. *Journal of Magnetic Resonance Imaging* 13, 534–546.
- Li, G., Zucker, S., 2003. A differential geometrical model for contour-based stereo correspondence. In: *Proceedings of the IEEE Workshop on Variational, Geometric, and Level Set Methods in Computer Vision*.
- 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.
- Martin-Fernandez, M., Westin, C.-F., Alberola-Lopez, C., 2004. 3D Bayesian regularization of diffusion tensor MRI using multivariate Gaussian Markov random fields. In: *Proceedings of the MICCAI 2004, LNCS 3216*, pp. 351–359.
- Medioni, G., Lee, M.S., Tang, C.K., 2000. *A Computational Framework for Segmentation and Grouping*. Elsevier, Amsterdam.
- Mori, S., Crain, B., Chacko, V., van Zijl, P., 1999. Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology* 45, 265–269.
- 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: *Proceedings of the MICCAI 2002, LNCS 2488*, pp. 459–466.
- Özarslan, E., Shepherd, T.M., Vemuri, B.C., Blackband, S.J., Mareci, T.H., 2005. Fast orientation mapping from HARDI. In: *Proceedings of the MICCAI 2005, LNCS 3749*, pp. 156–163.
- Parent, P., Zucker, S., 1989. Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11, 823–839.
- Parker, G., Alexander, D., 2005. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions of the Royal Society of London Series B* 360, 893–902.
- Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 629–639.
- Poupon, C., Clark, C., Frouin, V., Régis, J., Bloch, I., Le Bihan, D., Mangin, J.-F., 2000. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage* 12, 184–195.
- Ramírez-Manzanares, A., Rivera, M., 2003. Brain nerve bundles estimation by restoring and filtering intra-voxel information in diffusion tensor MRI. In: *Proceedings of the IEEE VLSM*.
- Rosenfeld, A., Hummel, R., Zucker, S., 1976. Scene labeling by relaxation operations. *IEEE Transactions on Systems, Man and Cybernetics SMC-6*, 420–433.
- Tang, B., Sapiro, G., Caselles, V., 2000. Diffusion of general data on non-flat manifolds via harmonic maps theory: the direction diffusion case. *International Journal of Computer Vision* 36, 149–161.
- 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.
- Tschumperlé, D., Deriche, R., 2002. Orthonormal vector sets regularization with PDE's and applications. *International Journal of Computer Vision* 50, 237–252.
- Tuch, D.-S., Weisskoff, R.-M., Belliveau, J.-W., Wedeen, V.-J., 1999. High angular resolution diffusion imaging of the human brain. In: *Proceedings of the International Society for Magnetic Resonance in Medicine: 7th Scientific Meeting and Exhibition*, p. 321.
- Tuch, D., Reese, T., Wiegell, M., Wedeen, V., 2003. Diffusion MRI of complex neural architecture. *Neuron* 40, 885–895.
- Wang, Z., Vemuri, B., Chen, Y., Mareci, T., 2004. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging* 23, 930–939.