

3D Curve Inference for Diffusion MRI Regularization

Peter Savadjiev¹, Jennifer S.W. Campbell²,
G. Bruce Pike², and Kaleem Siddiqi¹

McGill University, Montréal, QC, Canada

¹ School of Computer Science & Centre For Intelligent Machines

² McConnell Brain Imaging Centre, Montréal Neurological Institute

Abstract. We develop a differential geometric framework for regularizing diffusion MRI data. The key idea is to model white matter fibers as 3D space curves and to then extend Parent and Zucker’s 2D curve inference approach [8] by using a notion of *co-helicity* to indicate compatibility between fibre orientation estimates 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.

1 Introduction and Related Work

In this article, we consider the problem of regularizing orientation fields obtained using either Diffusion Tensor Imaging (DTI) [1] or High Angular Resolution Diffusion Imaging (HARDI) [12] measurements. We view both cases in a differential geometric setting where white matter fibers are modeled as 3D space curves. We extend Parent and Zucker’s 2D curve inference approach [8] by using a notion of *co-helicity* to model the compatibility between orientation estimates in 3D. In this context, curvature and torsion play key roles in the interpretation of a *tangent bundle*, where average local support is maximized using relaxation labeling techniques. We demonstrate that using 3D curve inference for the regularization of such data has advantages over earlier methods, including: 1) the possibility of representing multiple orientations at the same voxel, 2) applicability to both DTI and HARDI data, 3) numerical robustness in the vicinity of sparse measurements and 4) estimates of curvature and torsion at each voxel, which can be useful to guide fiber tracking algorithms.

We begin by briefly mentioning some of the key approaches to diffusion MRI regularization. Poupon *et al.* [9] use a “spaghetti plate” model for the regularization of DTI data. Westin *et al.* [7] regularize the individual elements of each DTI tensor using a six-dimensional, multivariate Gaussian Markov Random Field. Vemuri *et al.* [13] propose a constrained variational principle for the simultaneous estimation and regularization of DTI tensors directly from diffusion weighted images. Coulon *et al.* [5] combine variational and anisotropic diffusion methods to separately regularize orientation fields and diffusivities. Tschumperlé and

Deriche [11] also propose variational methods for regularizing orientation fields in DTI data while preserving discontinuities, with the constraint that orthonormality between the eigenvectors at each location is preserved.

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 fiber 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 [10, 3, 4]. 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 fiber orientations.

2 3D Curve Inference

We assume as input a 3D curve orientation distribution function sampled on a regular (typically rectangular) 3D lattice. Following Parent and Zucker’s 2D curve inference methodology [8], local curve orientation estimates can be interpreted as an initial *tangent bundle* at each location in the lattice. Our goal is to obtain estimates of the trace, tangent, curvature and torsion fields of curves in the 3D volume. This is done by using a notion of *co-helicity* between three tangents, which is the natural extension to 3D of Parent and Zucker’s *co-circularity* constraint between a pair of tangents in 2D [8]. In this framework, an *osculating helix* (which has constant curvature and torsion) is used to approximate a curve passing through a given point.

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, 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:

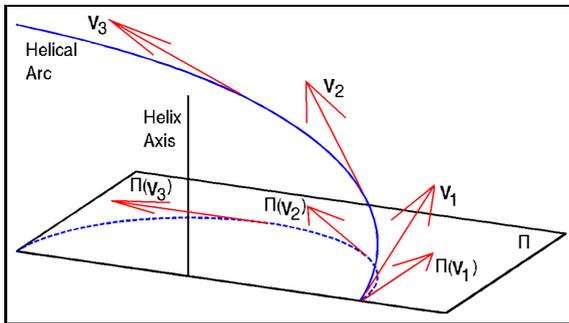


Fig. 1. An illustration of co-helicity between three vectors \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3

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

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

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

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). The orthographic projection of a helix onto the xy plane is a circle with radius r .

Definition 1. *Three orientations or vectors in 3D are co-helical if there is a helix to which all three are tangent.*

The concept of co-helicity is illustrated in Fig. 1, which shows a helical arc, three of its tangent vectors labeled \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 , along with the helix axis and its orthogonal plane Π . The orthographic projection of the helix onto Π is a circle (shown with a dashed curve), and the projections of \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 onto Π are labeled $\Pi(\mathbf{v}_1)$, $\Pi(\mathbf{v}_2)$ and $\Pi(\mathbf{v}_3)$. These three projections are co-circular [8], since they are all tangent to the same circle in plane Π .

Theorem 1. *Given three unit vectors \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 in \mathcal{R}^3 specified at three locations, it is possible to determine whether or not they are co-helical. Furthermore, if they are co-helical, it is possible to recover the parameters of the helix passing through these locations and having \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 as tangent vectors.*

Owing to space restrictions we omit the proof, but present an algorithm (Algorithm 1) for determining whether or not \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 in \mathcal{R}^3 are co-helical, and if so, for recovering the parameters of the helix.

2.2 Relaxation Labeling

In Parent and Zucker's 2D curve inference framework [8], tangent fields are estimated over a 2D image by maximizing the *average local support*, defined by $A(p) = \sum_{i=1}^n s_i(\lambda) p_i(\lambda)$. Expressions of this type can be maximized using the relaxation labeling algorithm of [6]. Here $p_i(\lambda)$ indicates the confidence in orientation λ at pixel i and its (local) neighborhood support given by $s_i(\lambda) = \sum_{j=1}^n \sum_{\lambda'=1}^m r_{ij}(\lambda, \lambda') p_j(\lambda')$. Each j denotes a node (i.e. pixel) in the neighborhood of i , λ' is an orientation at that node, $r_{ij}(\lambda, \lambda')$ represents the compatibility between orientation λ at node i and orientation λ' at node j , and there are a total of m orientations and n nodes. The compatibility coefficients are designed so that *co-circular* tangents lend one another support, while other configurations are suppressed. This imposes a constraint on the variation of curvature, while providing as a bi-product local estimates of curvature as well. The compatibility coefficients are further refined to include a partitioning into curvature classes [8].

The key to extending the above 2D curve inference framework to 3D is to replace the notion of *co-circularity* between two tangents with one of *co-helicity*

Algorithm 1: Determining co-helicity between three orientation estimates

Data : Three 3D vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$

Result : 1 if there is a helix to which $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are tangent, 0 otherwise.

1. (testing for a special case) If the three vectors are collinear, return 1.
2. Define a vector \mathbf{n} as the cross-product of the difference vectors $\mathbf{v}_3 - \mathbf{v}_2$ and $\mathbf{v}_2 - \mathbf{v}_1$. This is the axis of the putative helix.
3. Compute the projections $\Pi(\mathbf{v}_1), \Pi(\mathbf{v}_2)$ and $\Pi(\mathbf{v}_3)$ of $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ onto the plane that has \mathbf{n} as its normal vector. That fixes the value of the parameter r in the helix equation (1). That is, r is the radius of the circle, lying in the plane orthogonal to \mathbf{n} , to which $\Pi(\mathbf{v}_1), \Pi(\mathbf{v}_2)$ and $\Pi(\mathbf{v}_3)$ are all tangent.
4. Without loss of generality we can set the parameter t (see (1)) of \mathbf{v}_1 to $t_1 = 0$, the parameter t of \mathbf{v}_2 is then simply the angle $t_2 = \arccos\left(\left\langle \frac{\Pi(\mathbf{v}_1)}{\|\Pi(\mathbf{v}_1)\|}, \frac{\Pi(\mathbf{v}_2)}{\|\Pi(\mathbf{v}_2)\|} \right\rangle\right)$.
5. Using (1) and (2), it is easy to show that

$$\langle \mathbf{v}_1, \mathbf{v}_2 \rangle = \frac{r^2 \cos(t_2 - t_1) + c^2}{r^2 + c^2} \quad (4)$$

(assuming that \mathbf{v}_1 and \mathbf{v}_2 have been normalized to unit length). With the values for r, t_1 and t_2 obtained in steps 1 and 2, we can derive c using (4).

6. Given c , we can calculate (using (2)) the angle that the tangent vectors should make with the helix axis, as well as the displacement in the direction of the helix axis of the position of \mathbf{v}_1 and \mathbf{v}_2 ((1)). It is possible to verify whether \mathbf{v}_1 and \mathbf{v}_2 possess indeed these properties. If so, return 1. If not, return 0.
 7. If any of the above steps fail (e.g. there is no one circle to which $\Pi(\mathbf{v}_1), \Pi(\mathbf{v}_2)$ and $\Pi(\mathbf{v}_3)$ are all tangent in step 3) return 0.
-

between three tangents. The compatibility condition obviously has to be generalized to a notion of higher order compatibility. Following [6], this is done by replacing the support function by

$$s_i(\lambda) = \sum_{j, \lambda'} \sum_{k, \lambda''} r_{ijk}(\lambda, \lambda', \lambda'') p_j(\lambda') p_k(\lambda''). \quad (5)$$

Here $r_{ijk}(\lambda, \lambda', \lambda'')$ represents the compatibility (co-helicity) between orientation λ at node i , orientation λ' at node j and orientation λ'' at node k , as determined using Algorithm 1.

2.3 Implementation

The 3D curve inference algorithm can use an arbitrarily large label set. In our implementation, we use a label set of 90 unit direction vectors distributed isotropically over a hemisphere, obtained using an electrostatic charge repulsion algorithm. The relaxation labeling technique is implemented according to Algorithm 8.2 in [6], together with the radial projection method (Appendix A in [8]). The

support coefficients are calculated in a spherical neighborhood of each voxel. We obtain a significant increase in efficiency by precomputing the co-helical configurations of tangent triplets and storing these in a look-up table. The initial values for $p_i(\lambda)$ are obtained from the value of the orientation distribution function (ODF) along orientation λ at voxel i , and a discretization of the allowed range of variation of curvature and torsion (the 3D extension of the curvature classes in [8]) is implemented. Note that the output of our algorithm is not an ODF any longer, but a regularized tangent bundle, which is an indication of confidence in the underlying curve (i.e. white matter fibre) orientations at each voxel.

3 Validation

Quantitative Validation On a Biological Phantom: A biological phantom was created from two excised Sprague-Dawley rat spinal cords embedded in 2% agar. Two diffusion-weighted datasets were acquired using this phantom, with 90 diffusion encoding directions, with b values of 1300 s/mm^2 and 3000 s/mm^2 , respectively. The first was used for diffusion tensor reconstruction of the diffusion ODF and the second for high angular resolution reconstruction, using the q-ball technique [12]. A T1-weighted image of this phantom is shown in Fig. 2 (top left). The ground truth orientations were determined by extracting the centerlines of each cord using the technique of [2] and then smoothly extending the orientation estimates in the center to the boundary of the cord, for each cord (Fig. 2, top

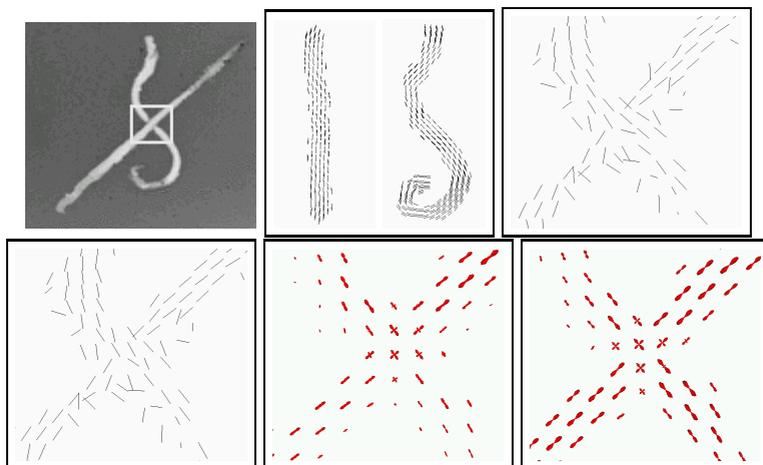


Fig. 2. Top Row: A biological phantom created by overlaying two rat cord spinal cords (left), ground truth orientation estimates shown separately for the two cords (middle), and the principal eigenvector orientations of the DTI image in the vicinity of the crossing (right). Bottom Row: The regularized orientation estimates obtained using the technique of [11] (left), those obtained using 3D curve inference on the DTI reconstruction (middle), and the regularized HARDI reconstruction using 3D curve inference (right).

middle). Fig. 2 (top right) shows the principal components of the diffusion tensors in the region of the crossing, indicated by a rectangle in Fig. 2 (top left). The bottom row of Fig. 2 shows the results obtained using Tschumperlé and Deriche’s orthonormal vector regularization [11] (OVR-DTI) (left), curve inference on the DTI dataset (CI-DTI) (middle) and curve inference on the HARDI dataset (CI-HARDI) (right). The average orientation errors in degrees (± 1 std. deviation) between the ground truth dataset and the unregularized as well as regularized datasets for each case are shown in Fig. 3 (left). Observe that the results obtained using curve inference give significantly lower errors. Furthermore, the lowest error is obtained by applying curve inference to HARDI data. The phantom data is challenging because of eddy current induced artefacts in many boundary voxels, where the measured principal direction of diffusion is perpendicular to the ground truth orientation. In the case of the curve inference experiments, the error is measured between the maximally supported orientation at each voxel in the regularized data, and the corresponding ground truth orientation.

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 intersect the sine wave, and background voxels were filled with isotropic (spherical) tensors. All tensors had a mean eigenvalue of 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 perturbed (independently) by adding Gaussian noise, with mean 0 and a standard deviation of 22.9° (0.4 rad). The original (noiseless) dataset was treated as the ground truth. One view of the noisy dataset is shown in Fig. 3 (right); 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. 3 (left). Observe that once again curve inference (CI-DTI) achieves a significant reduction in orientation error, compared to both the noisy unregularized data, as well as the result obtained with OVR-DTI [11].

Qualitative Validation On Human Brain Data: We conclude with regularization results using *in vivo* human brain data. Fig. 4 (top left) shows a slice through the data, which consists of a DTI reconstruction of the diffusion ODF. The data covers an area with multiple fibre directions, including voxels with partial volume averaging effects, as well as voxels with cerebro-spinal fluid and grey matter that do not exhibit high diffusion anisotropy. Fig. 4 (top right) shows one region of interest (ROI) of the regularization result using 3D curve inference, and Fig. 4 (bottom left) shows a slice through the corresponding 3D fractional anisotropy image of the brain, with a white rectangle indicating the ROI. A zoom-in on the bottom-left part of the regularization result is given in Fig. 4 (bottom right), showing the recovery of multiple fibre directions in voxels with partial volume averaging effects. Qualitatively, intra-voxel crossings and the apparent variation

	Phantom	Synthetic
Unreg. DTI	$27.1 \pm 29.0^\circ$	$24.9 \pm 14.5^\circ$
Unreg. HARDI	$21.0 \pm 21.9^\circ$	N/A
CI-DTI	$11.0 \pm 16.3^\circ$	$11.2 \pm 8.2^\circ$
CI-HARDI	$10.3 \pm 17.3^\circ$	N/A
OVR-DTI	$24.0 \pm 22.2^\circ$	$21.2 \pm 12.9^\circ$



Fig. 3. LEFT: Table of validation results showing average orientation errors in degrees for the biological phantom data set and the synthetic data set. RIGHT: A snapshot of the noisy synthetic data set, prior to regularization.

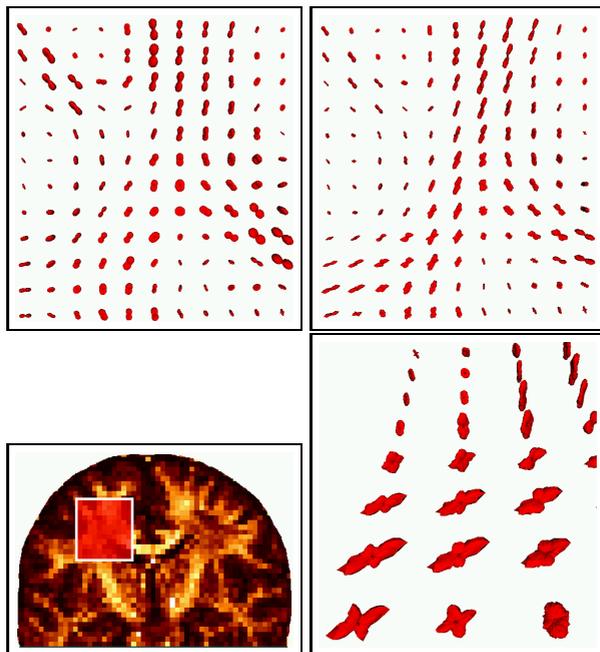


Fig. 4. A ROI through the brain DTI data (top left) with the regularization results using curve inference (top right). A slice through the associated fractional anisotropy image (bottom left) with a white rectangle enclosing the ROI. A zoom-in on the result from a different viewpoint in a region of partial volume averaging effects (bottom right).

of curvatures are recovered well from the DTI data. Quantitative evaluation of these results is difficult, due to the lack of ground truth.

4 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 orienta-

tions given neighbouring orientation estimates. Since multiple orientations can receive support in the same voxel, the algorithm is applicable to configurations of crossings or branchings, and it handles DTI and HARDI data in an identical way. It can also be applied to the regularization of any set of ODFs distributed over a discrete 3D lattice. As a bi-product of the algorithm, we obtain discrete estimates of the curvature and torsion of the likely curves at each voxel, which could be used to guide fiber tracking algorithms. We have validated the algorithm quantitatively on a biological phantom and on a synthetic data set, and qualitatively on *in vivo* human brain data.

Acknowledgments. We thank Sylvain Bouix for his centerline extraction algorithm and Vladimir V. Rymar for his help with preparing the biological phantom. This work was supported by grants from FQRNT, NSERC and CIHR.

References

1. P.-J. Basser, J. Mattiello, and D. Le Bihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
2. S. Bouix, K. Siddiqi, and A. Tannenbaum. Flux driven automatic centerline extraction. *Medical Image Analysis*, 9(3):209–221, 2005.
3. Y. Chen, W. Guo, Q. Zeng, X. Yan, F. Huang, H. Zhang, G. He, B.-C. Vemuri, and Y. Liu. Estimation, smoothing, and characterization of apparent diffusion coefficient profiles from high angular resolution DWI. In *Proc. IEEE CVPR*, pages 588–593, 2004.
4. Y. Cointepas, C. Poupon, D. Le Bihan, and J.-F. Mangin. A spin glass framework to untangle fiber crossing in MR diffusion based tracking. In *Proc. MICCAI 2002, LNCS 2488*, pages 475–482, 2002.
5. O. Coulon, D. Alexander, and S. Arridge. A regularization scheme for diffusion tensor magnetic resonance images. In *Proc. IPMI 2001, LNCS 2082*, pages 92–105, 2001.
6. R. Hummel and S. Zucker. On the foundations of relaxation labeling processes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 5:267–287, 1983.
7. M. Martin-Fernandez, C.-F. Westin, and C. Alberola-Lopez. 3D Bayesian regularization of diffusion tensor MRI using multivariate gaussian markov random fields. In *Proc. MICCAI 2004, LNCS 3216*, pages 351–359, 2004.
8. P. Parent and S. Zucker. Trace inference, curvature consistency, and curve detection. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 11:823–839, 1989.
9. C. Poupon, C. Clark, V. Frouin, J. Régis, I. Bloch, D. Le Bihan, and J.-F. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12:184–195, 2000.
10. A. Ramírez-Manzanares and M. Rivera. Brain nerve bundles estimation by restoring and filtering intra-voxel information in diffusion tensor MRI. In *Proc. IEEE VLISM*, 2003.
11. D. Tschumperlé and R. Deriche. Orthonormal vector sets regularization with PDE’s and applications. *Int. Journal of Computer Vision*, 50:237–252, 2002.
12. D. Tuch, T. Reese, M. Wiegell, and V. Wedeen. Diffusion MRI of complex neural architecture. *Neuron*, 40:885–895, 2003.
13. Z. Wang, B. Vemuri, Y. Chen, and T. Mareci. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Trans. on Medical Imaging*, 23:930–939, 2004.