

Beyond crossing fibers: Tractography exploiting sub-voxel fibre dispersion and neighbourhood structure

Matthew Rowe, Hui Zhang, Neil Oxtoby, Daniel C Alexander

Centre for Medical Image Computing, Department of Computer Science, University
College London, UK

`matthew.rowe.09@ucl.ac.uk`

Abstract. In this paper we propose a novel algorithm which leverages models of white matter fibre dispersion to improve tractography. Tractography methods exploit directional information from diffusion weighted magnetic resonance (DW-MR) imaging to infer connectivity between different brain regions. Most tractography methods use a single direction (e.g. the principal eigenvector of the diffusion tensor) or a small set of discrete directions (e.g. from the peaks of an orientation distribution function) to guide streamline propagation. This strategy ignores the effects of within-bundle orientation dispersion, which arises from fanning or bending at the sub-voxel scale, and can lead to missing connections. Various recent DW-MR imaging techniques estimate the fibre dispersion in each bundle directly and model it as a continuous distribution. Here we introduce an algorithm to exploit this information to improve tractography. The algorithm further uses a particle filter to probe local neighbourhood structure during streamline propagation. Using information gathered from neighbourhood structure enables the algorithm to resolve ambiguities between converging and diverging fanning structures, which cannot be distinguished from isolated orientation distribution functions. We demonstrate the advantages of the new approach in synthetic experiments and *in vivo* data. Synthetic experiments demonstrate the effectiveness of the particle filter in gathering and exploiting neighbourhood information in recovering various canonical fibre configurations and experiments with *in vivo* brain data demonstrate the advantages of utilising dispersion in tractography, providing benefits in practical situations.

1 Introduction

Tractography is a powerful tool to probe the geometric structure of white matter *in vivo* from non-invasive DW-MR imaging, allowing us to infer the anatomical connectivity of the separate functional regions of the brain. Information gained from tractography has great potential to advance our understanding of neurological function and disease.

Tractography algorithms infer connectivity by propagating streamlines between locations in the brain using directional information derived from DW-MR

images. Early techniques were deterministic and followed a single direction per voxel, taken from the principal eigenvector of the diffusion tensor. The simplest method involves nearest neighbour interpolation FACT (Fiber Assignment by Continuous Tracking) [1], others involve trilinear interpolation [2, 3]. As the local fibre orientation estimates are prone to errors and uncertainty, probabilistic techniques were developed which assume a variance on the dominant fibre orientation and propagate a large collection of streamlines from each seed location using monte-carlo sampling of each fibre orientation estimate, thus inferring an index of connection probability between separate locations related to the number of streamlines which connect them [4, 5]. Multi-fibre techniques addressed the potential existence of multiple fibre populations traversing a voxel [6–9]. The majority of these techniques make use of a discrete set of fibre orientation estimates in each voxel (e.g. from the peaks of a fibre orientation distribution function). They assume that the uncertainty in the peak directions captures orientation dispersion due to noise and underlying fibre dispersion [10], while some [11] sample a fibre orientation distribution function (fODF) directly via a rejection sampling scheme for streamline propagation. While these methods provide good solutions to the crossing fibre problem, they do not account for other sub-voxel fibre configurations such as fanning and bending. This can lead to false negative connections in regions where such fibre architecture exists such as the corona radiata.

Global tractography methods [12–14] pose a potential solution to these limitations of local tractography. These methods search for the set of all streamlines which best explain the entire DW-MR data set. In theory, this approach handles complex sub-voxel fibre architectures such as fanning and crossing since the distribution of orientations formed by the candidate streamlines passing through each voxel is required to support the DW-MR signal in each voxel. Also local ambiguities such as fanning vs. crossing and fanning polarity could also be resolved in theory, since the streamlines are continuous and must reflect the diffusion weighted data globally, hence data from multiple voxels supports or opposes the existence of any particular candidate streamline. Sherbondy [15] demonstrates that by combining such a global technique with microstructure modelling can resolve classic confounding fibre architectures such as kissing vs. crossing. However, the major drawback of global tractography is the computational cost. The search space is very high-dimensional so globally optimal solutions are impossible to find in practical timescales with current technology. Recent work [16] has questioned whether these suboptimal solutions offer any real advantage over local tractography in real brain data.

Recent DW-MR modelling and imaging techniques [17–20] use parametric models of the fODF to recover estimates of within-voxel dispersion. These techniques avoid the instability with non-parametric fODFs such as those from standard spherical deconvolution and have been shown to match well fODFs measured from histology [21]. Preliminary investigations have demonstrated that tractography techniques sampling these fODFs in full for streamline propagation can reduce false negatives in regions of high anatomical fibre dispersion such as

the the corona radiata [22]. However, Jeurissen *et al.* [10] show that making use of the full fODF for sampling propagation directions renders tractography more susceptible to false positives due to ambiguities about the underlying anatomical basis of the orientation dispersion.

To resolve ambiguities in the local voxel model of fibre orientation, we can make use of information from local neighbourhood voxels in the vicinity of the region which a streamline is traversing in addition to the local model. Savadjiev *et al.* [23] demonstrate that information from a voxels neighbourhood can be used to disambiguate sub-voxel fibre architectures such as curving and fanning. Using helical curves projected into the neighbourhood of a voxel, Savadjiev derives markers distinguishing and quantifying fanning and crossing fibres and fanning polarity within each voxel and parameterises the set of streamline selections enabling evaluation of those which are most consistent with forthcoming local structure. These methods demonstrate nicely the potential of leveraging voxel neighbourhood information for disambiguating fibre architecture at the subvoxel level. These methods however do not make use of a parameterised local model of dispersion and carry a computational expense due to the complexity of the helical model of streamlines, which would be costly to apply in a probabilistic tractography framework.

In this paper we present a new tractography method which combines local models of sub-voxel fibre dispersion with voxel neighbourhood exploration via a particle filter. Particle filter methods have been used previously in tractography [24–26] to estimate global connectivity in DW-MR images. In contrast to these methods we use the particle filter over a short range at each step of the streamline propagation to sample the set of possible future trajectories suggested by the local dispersion model and probe their compatibility with the neighbouring image structure. The particle filter informs the next single step selected from the local model, after which we repeat the whole procedure. Each step is stochastic and so, therefore, is the final resulting streamline. Thus the full algorithm repeats the whole process, as in traditional probabilistic tractography, to establish a collection of candidate streamlines from which we can derive probabilistic indices of connectivity in the usual way [4, 5]. We demonstrate the behaviour of the particle filter in gathering information on neighbourhood geometry in common fibre configurations on synthetic data and demonstrate the advantages of using dispersion in tractography on *in vivo* brain data tracking through the corona radiata.

2 Methods

In this section we describe our tracking algorithm that exploits local fibre dispersion and neighbourhood structure via a particle filter. Section 2.1 explains the details of the local model and Section 2.2 presents the technique to create a directional distribution informed by both local and neighbourhood information.

2.1 Locally-estimated dispersion ODF

To capture dispersion at each voxel locally, we model the fODF as a Bingham distribution. The Bingham distribution is a function on the sphere with a mean direction $\boldsymbol{\mu}$ and two concentration parameters κ_1 and κ_2 which control the degree of dispersion along the two axes orthogonal to the mean direction $\boldsymbol{\mu}_1$ and $\boldsymbol{\mu}_2$ respectively:

$$f(\mathbf{n}) = F_1\left(\frac{1}{2}, \frac{3}{2}, \kappa_1, \kappa_2\right)^{-1} \exp[\kappa_1(\boldsymbol{\mu}_1 \cdot \mathbf{n})^2 + \kappa_2(\boldsymbol{\mu}_2 \cdot \mathbf{n})^2], \quad (1)$$

where F_1 is the hypergeometric function (note: $F_1(1/2, 3/2, \kappa_1, \kappa_2)$ is a number, not a function).

The Bingham distribution has been used to model fibre dispersion in [17, 19]. For this work, we make use of the NODDI technique for fitting fibre dispersion models. The NODDI technique fits the cylindrically symmetric Watson distribution and introduces an imaging protocol better able to support the estimation of fibre dispersion. Here we use an extension of the NODDI technique to fit a Bingham distribution to the DW data. The Bingham distribution captures cylindrically asymmetric dispersion, which can better represent planar dispersion, representative of fanning white matter fibre structure [19].

Although such a local model of dispersion can improve exploration of sub-voxel fibre trajectories, it exhibits a number of ambiguities which cannot be resolved by examining each voxel in isolation. Due to the symmetrical nature of DW-MR measurements, the resulting distributions are symmetric, meaning there is an ambiguity between sub-voxel curving and fanning, and also fanning polarity. To address these ambiguities, it is necessary to gather information from voxels in the neighbourhood of the subject voxel to inform on the treatment of the local model for streamline propagation. For this purpose we use a particle filter framework explained in Section 2.2.

2.2 Voxel neighbourhood-informed dispersion ODF

In this section, we describe a technique that creates a neighbourhood-informed fODF from the local dispersion estimates by fusing the information from the neighbourhood structure of the dominant fibre orientation. We draw candidate directions from the local dispersion fODF and propagate these directions into the local neighbourhood. By examining the coherence of each projected streamline with neighbourhood structure we can then weight each of these candidate directions according to the coherence of the respective projected streamline with the neighbourhood structure. This process is illustrated in Figure 1 for various canonical neighbourhood structures. The streamlines which are misaligned with neighbourhood structure are downweighted (coloured in blue) and the streamlines which align with the neighbourhood structure have their weights increased (coloured in red). This shows that fanning polarity and curvature can be distinguished. In the case of a diverging neighbourhood structure, the streamlines propagated from the dispersed candidate directions drawn from the fODF in

the subject voxel find good alignment with the neighbourhood structure and are therefore evenly weighted. In contrast, in the presence of a convergent neighbourhood structure, the streamlines propagated from the peripheral candidate directions drawn from the local fODF misalign with neighbourhood structure and are downweighted. In the case of a curving structure, the candidate directions propagating against the curve are downweighted.




Fig. 1. Illustration of neighbourhood exploration in the case of of diverging (a), converging (b), coherent (c) and curving (d) neighbourhood structure in the tracking direction. Red streamlines are highly weighted, blue are low weighted.

This neighbourhood exploration scheme falls naturally into a particle filter framework. The particle filter is a simulation based model estimation technique used in non-linear, non-Gaussian dynamical systems. Through a process of prediction and update, the particle filter provides a discrete approximation of a posterior distribution $p(x_k|y_{0:k})$ on a time-varying parameter x_k at timestep k given the observations $y_{0:k}$ for timesteps $0, 1, 2, \dots, k$ and the initial state distribution $p(x_0)$. At each timestep k , N particles are propagated by sampling from an importance density $\pi(x_k^{(i)}|x_{0:k-1}^{(i)}, y_{0:k})$, then assigned importance weights $w_k^{*(i)}$ which depend on a likelihood model $p(y_k|x_k^{(i)})$. Subsequently the discrete approximation to the posterior distribution $p(x_k|y_{0:k})$, denoted by $\tilde{w}_k^{(i)}$, is computed by

normalising $w_k^{*(i)}$. This operation of a particle filter is summarised in Algorithm 1. Further details on particle filtering can be found in [27].

```

State initialization, sample  $x_0$  from  $p(x_0)$ ;
Initialise importance weights
for  $i = 1, \dots, N$ , do
    |  $w_0^{*(i)} = \frac{1}{N}$ 
end
for times  $k = 1, 2, \dots, K$  do
    | for  $i = 1, \dots, N$ , do
    | | sample  $x_k^{(i)}$  from  $\pi(x_k | x_{0:k-1}^{(i)}, y_{0:k})$ 
    | end
    | calculate weight up to normalisation factor:
    | for  $i = 1, \dots, N$ , do
    | |  $w_k^{*(i)} = w_{k-1}^{*(i)} p(y_k | x_k^{(i)})$ 
    | end
    | normalise the importance weights:
    | for  $i = 1, \dots, N$ , do
    | |  $\tilde{w}_k^{(i)} = \frac{w_k^{*(i)}}{\sum_{j=1}^N w_k^{*(j)}}$ 
    | end
end

```

Algorithm 1: Sequential importance sampling

In this implementation, the importance density $\pi(x_k^{(i)} | x_{0:k-1}^{(i)}, y_{0:k})$ is chosen as a Watson distribution and the initial state distribution $p(x_0)$ is the Bingham distribution (described in Section 2.1) from the current voxel. A cloud of N particles defines a set of N streamlines defined by a string of vectors of fixed length connected end to end. At each timestep k each streamline is propagated one step from its previous location $u_{k-1}^{(i)}$ with a direction vector $v_k^{(i)}$ sampled from the importance density by a step length d such that $u_k^{(i)} = u_{k-1}^{(i)} + dv_k^{(i)}$. The state of a particle at timestep k $x_k^{(i)}$ is defined by its location $u_k^{(i)}$ and direction vector $v_k^{(i)}$. At each step the particle weights $w_k^{*(i)} = w_{k-1}^{*(i)} p(y_k | x_k^{(i)})$ are calculated to reflect their alignment with neighbourhood structure and the process is repeated for K steps. The likelihood $p(y_k | x_k^{(i)}) = (v_k \cdot D(u_k))^\gamma$ where $D(u_k)$ is the interpolated direction of the vector field D , defined by the mean directions of the Bingham distributions in each voxel, at the point location u_k . The stages of the particle filter scheme from streamline propagation to the selection of tract propagation direction is illustrated in Figure 2.

3 Experiments and results

3.1 Synthetic experiments

We use simulated data to examine the behaviour of the neighbourhood exploration by the particle filter. We simulate vector fields to mimic neighbourhood




Fig. 2. Illustration of the stages of neighbourhood exploration: initialisation from the local Bingham distribution (a), particle update and weighting (b), particle selection (c) and tract propagation direction selection (d).

structures which can arise in conjunction with dispersive local fODFs which are not distinguishable from the information from an isolated voxel. Figure 3 shows the behaviour of the particle filter in regions of diverging 3(a), converging 3(b), coherent 3(c), and curving 3(d) neighbourhood structure. The figure shows the particle weights at their final iteration step K .

Figure 4 shows a demonstration of the proposed tracking algorithm over 100 repetitions in diverging and converging local structure.

3.2 *in vivo* experiments

We apply the algorithm on *in vivo* brain data. DW-MR images of a healthy male were acquired on a clinical 3T Philips system with isotropic voxels of 2mm, TE=78ms, TR=12.5ms, with one 30 direction shell and one 60 direction shell with b -values of 1000 s/mm² and 2000 s/mm², respectively. This dataset is the same as that used in [20]. Tracking experiments were performed from a single seed in the midbody of the corpus callosum (CC) at the mid-sagittal position, using the proposed algorithm and diffusion tensor (DT) PICo [4]. 5000 repetitions were used and results are shown in Figure 5. Tractography is also performed on 4 major white matter pathways: the inferior longitudinal fasciculus, the superior longitudinal fasciculus, the cingulum and the occipito-frontal fasciculus,




Fig. 3. Demonstration of particle filter behaviour in regions of diverging (a), converging (b), coherent (c) and curving (d) neighbourhood structure in the tracking direction. Colour denotes particle weight, blue is low through to red, which is high, yellow is intermediate.

for validation of expected performance in standard tracts. All *in vivo* results are displayed as maximum connection probability maps with a threshold of 1%.

4 Discussion and conclusions

We have presented a tractography algorithm that combines information from both a local model of fibre orientation, which captures sub-voxel fibre dispersion with information drawn from the structural information available in the voxels in the immediate neighbourhood of the current tracking location. The simulations in Figure 3 demonstrate that the neighbourhood exploration particle filter behaves in practice like the conceptual illustration in Figure 1. In the case of divergent structure in the tracking direction (Figure 3(a)), the particles have an evenly spread distribution of weights, allowing full exploration of the orientations from the local model. In the case of convergent neighbourhood structure in the tracking direction (Figure 3(b)), the high particle weights are concentrated in the middle, limiting the selection of directions to align with forthcoming structure. Figure 4 shows the result of multiple iterations of the tractography algorithm in converging and diverging structure. In Figures 4(a) and 4(c) neighbourhood exploration is used, while in 4(b) and 4(d) neighbourhood exploration is not used, a curvature prior is used instead. The dispersion of streamline trajectories is exploited in the diverging case while the dispersion is restrained in the




Fig. 4. 100 repetitions of the tracking algorithm in regions of diverging (a), (b) and converging (c) and (d) neighbourhood structure in the tracking direction. In (a) and (c), particle filter neighbourhood exploration is used, while in (b) and (d) only a curvature prior is used.




Fig. 5. Tracking from a single seed in the corpus callosum.




Fig. 6. Tracking of major white matter structures using proposed algorithm. a) inferior longitudinal fasciculus, b) superior longitudinal fasciculus, c) cingulum, d) occipito-frontal fasciculus

converging case over multiple repetitions. These simulations show that combining information from the voxel neighbourhood with information from the local fODF, we can resolve some of the remaining structural ambiguities of the local model such as fanning polarity, which cannot be resolved from the voxel information in isolation.

Figure 5(a) shows the advantages of using dispersion in tractography in *in vivo* data. Tracking from a single seed in the corpus callosum, we can recover a wide range of connectivity to areas of the peripheral cortex. Both lateral and vertical connections are covered evenly with the oblique connections between. These streamlines traverse the corona radiata, which is a brain region known from anatomy to exhibit dispersing white matter structure. Figure 5(b) shows tracking from the same seed with PICO tractography based on the diffusion tensor (DT), which only uses a single discrete direction per voxel. DT PICO only recovers the vertical connections from the CC to the cortex as it does not exploit the dispersion in the corona radiata.

Figure 6 shows validation of the algorithm in tracking major white matter pathways, for which standard tractography algorithms work well, showing that the algorithm performs as expected in these structures.

Using a hybrid method utilising local fODFs modelling dispersion and neighbourhood exploration, we can capture complex dispersing subvoxel fibre architectures and include information from multiple voxels to resolve ambiguities of the local model and improve connectivity estimation. The synthetic experiments demonstrate the ability of the neighbourhood exploration to resolve ambiguities of local dispersive fODF such as fanning polarity and curvature, while the *in vivo* experiments show the advantages of exploiting fibre dispersion in tractography.

This technique represents a middle ground between local and global tractography, making use of the global tractography paradigm of pooling information from multiple voxels to overcome the limitations of considering each voxel individually, while not inflating the problem to computationally intractable proportions. In future work we can explore other areas of this middle ground with larger neighbourhood exploration regions and using alternative fODFs such as those derived from spherical deconvolution.

References

1. Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.M.: Three dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol* **45** (1999) 265–269
2. Conturo, T.E., Lori, N.F., Cull, T.S., Akbudak, E., Snyder, A.Z., Shimony, J.S., McKinstry, R.C., Burton, H., Raichle, M.E.: Tracking neuronal fiber pathways in the living human brain. *Proc. Natl. Acad. Sci. U.S.A.* **96** (1999) 10422–10427
3. Basser, P.J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A.: In vitro fiber tractography using DT-MRI data. *Magn. Reson. Med.* **44** (2000) 625–632
4. Parker, G.J., Alexander, D.C.: Probabilistic monte carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. *Lecture Notes in Computer Science* **2732** (2003) 684–695
5. Behrens, T.E.J., Johansen-Berg, H., Woolrich, M.W., Smith, S.M., Wheeler-Kingshott, C.A.M., Boulby, P.A., Barker, G.J., Sillery, E.L., Sheehan, K., Cicarelli, O., Thompson, A.J., Brady, J.M., Matthews, P.M.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.* **50** (2003) 1077–1088
6. Parker, G.J., Alexander, D.C.: Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions of the Royal Society B: Biological Sciences* **360** (2005) 893–902
7. Behrens, T.E.J., Johansen-Berg, H., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage* **34** (2007) 144–155
8. Tournier, J., Calamante, F., Connelly, A.: Robust determination of the fibre orientation distribution in diffusion mri: non-negativity constrained super-resolved spherical deconvolution. *NeuroImage* **35** (2007) 1459–1472
9. Jeurissen, B., Leemans, A., Tournier, J.D., Sijbers, J.: Estimation of uncertainty in constrained spherical deconvolution fiber orientations. In: *IEEE International Symposium on Biomedical Imaging: Macro to Nano.* (2008) 907–910
10. Jeurissen, B., Leemans, A., Jones, D.K., Tournier, J.D., Sijbers, J.: Probabilistic fiber tracking using the residual bootstrap with constrained spherical deconvolution. *Human Brain Mapping* **32** (2011) 461–479
11. Tournier, J.D., Calamante, F., Gadian, D.G., Connelly, A.: Probabilistic fibre tracking through regions containing crossing fibres. In: *Proc. Intl. Soc. Mag. Reson. Med.* Volume 13. (2005) 1343
12. Sherbondy, A.J., Dougherty, R.F., Ananthanarayanan, R., Modha, D.S., Wandell, B.A.: Think Global, Act Local; Projectome Estimation with Bluematter. *MICCAI* (2009) 861–868
13. Fillard, P., Poupon, C., Mangin, J.F.: A Novel Global Tractography Algorithm Based on an Adaptive Spin Glass Model. *MICCAI* **60** (2009) 927–934

14. Kreher, B.W., Mader, I., Kiselev, V.G.: Gibbs Tracking: A Novel Approach for the Reconstruction of Neuronal Pathways. *Magnetic Resonance in Medicine* **60** (2008) 953–963
15. Sherbondy, A.J., Rowe, M., Alexander, D.C.: MicroTrack: an algorithm for concurrent projectome and microstructure estimation. *MICCAI* **13** (2010) 183–90
16. Li, L., Rilling, J.K., Preuss, T.M., Glasser, M.F., Damen, F.W., Hu, X.: Quantitative assessment of a framework for creating anatomical brain networks via global tractography. *NeuroImage* **61** (2012) 1017–30
17. Kaden, E., Knosche, T.R., Anwender, A.: Parametric spherical deconvolution: Inferring anatomical connectivity using diffusion MR imaging. *NeuroImage* **37** (2007) 474–488
18. Zhang, H., Hubbard, P.L., Parker, G.J.M., Alexander, D.C.: Axon diameter mapping in the presence of orientation dispersion with diffusion MRI. *NeuroImage* **56** (2011) 1301–1315
19. Sotiropoulos, S.N., Behrens, T.E.J., Jbabdi, S.: Ball and rackets: Inferring fiber fanning from diffusion-weighted MRI. *NeuroImage* **60** (2012) 1412–1425
20. Zhang, H., Shneider, T., Wheeler-Kingshott, C., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61** (2012) 1000–1016
21. Jespersen, S.N., Leigland, L.A., Cornea, A., Kroenke, C.D.: Determination of axonal and dendritic orientation distributions within the developing cerebral cortex by diffusion tensor imaging. *IEEE Trans. Med. Imaging* **31** (2012) 16–32
22. Rowe, M., Zhang, H., Alexander, D.C.: Utilising measures of fiber dispersion in white matter tractography. *MICCAI CDMRI workshop* (2012)
23. Savadjiev, P., Campbell, J.S.W., Descoteaux, M., Deriche, R., Pike, G.B., Siddiqi, K.: Labeling of ambiguous subvoxel fibre bundle configurations in high angular resolution diffusion MRI. *NeuroImage* **41** (2008) 58–68
24. Zhang, F., Hancock, E.R., Goodlett, C., Gerig, G.: White matter tractography using sequential importance sampling. *Proc. ISMRM Annual meeting* **10** (2002)
25. Zhang, F., Hancock, E.R., Goodlett, C., Gerig, G.: Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. *Med. Image Anal.* **13** (2009) 5–18
26. Pontabry, J., Rousseau, F.: Probabilistic tractography using Q-ball modeling and particle filtering. *MICCAI* **14** (2011) 209–16
27. Doucet, A., Godsill, S., Andrieu, C.: On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. and Comput.* **10** (2000) 197–208