Exploiting peak anisotropy for tracking through complex structures

Kiran K Seunarine[1], P A Cook[2], M G Hall[1], K V Embleton[3], G J M Parker[3] and Daniel C Alexander[1]
[1]Centre for Medical Image Computing, University College London, k.seunarine@cs.ucl.ac.uk
[2]Department of Radiology, University of Pennsylvania
[3]Department of Imaging Science and Biomedical Engineering, University of Manchester

Abstract

This work shows that multi-fibre reconstruction techniques, such as Persistent Angular Structure (PAS) MRI or QBall Imaging, provide much more information than just discrete fibre orientations, which is all that previous tractography algorithms exploit from them. We show that the shapes of the peaks of the functions output by multiple-fibre reconstruction algorithms reflect the underlying distribution of fibres. Furthermore, we show how to exploit this extra information to improve Probabilistic Index of Connectivity (PICo) tractography. The method uses the Bingham distribution to model the uncertainty in fibre-orientation estimates obtained from peaks in the PAS or QBall Orientation Distribution Function (ODF). The Bingham model captures anisotropy in the uncertainty, allowing the method to track through fanning and bending structures, which previous methods do not recover reliably. We devise a new calibration procedure to construct a mapping from peak shape to Bingham parameters. We test the accuracy of the calibration using a bootstrap experiment. Finally, we show that exploiting the peak shape in this way can provide improved PICo tractography results.

1. Introduction

Diffusion MRI probes the microstructure of the brain by observing the diffusion of water molecules. In particular, it is used to examine the fibrous white-matter that connects the different regions of the brain. The technique has several applications; it can be used to detect white-matter diseases [28], generate atlases of connectivity of the brain [32] and in neurosurgical planning and evaluation [3, 20]. The standard technique in diffusion MRI is Diffusion Tensor Imaging (DTI), which fits the diffusion tensor to diffusion-weighted measurements in each image voxel. However, DTI fails when voxels contain complex structure, such as crossing fibres, because it assumes a single fibre direction in each voxel. Multi-tensor models [7, 11] overcome this limitation to some extent, but require knowledge of the number of fibre populations prior to fitting. More recently, various *multiple-fibre reconstruction algorithms* have been developed, such as Persistent Angular Structure (PAS) MRI [15], Spherical Deconvolution [14] and QBall Imaging [35], which resolve multiple fibre orientations in a single voxel without prior knowledge of the number of fibre directions.

Tractography uses fibre-orientation estimates to find the paths of white-matter tracts through the brain. The simplest method is streamline tractography [29, 30]. Starting from a seed point, a path (or streamline) is propagated from voxel to voxel using the local fibre orientation estimates to determine the direction of each step. A major limitation of streamline tractography is that noise in the diffusion-weighted measurements causes uncertainty in the estimates of the fibre orientations. The streamlines only reflect the maximum-likelihood path of the white-matter, which does not capture all of the connectivity. Probabilistic tractography [10, 17, 18, 31] uses estimates of uncertainty of each fibre-orientation estimate and Monte Carlo sampling of the streamline paths to create a map of an index of connectivity of every voxel in the brain to a specified seed voxel.

Most tractography algorithms still use the basic DTI single-fibre reconstruction and it is not clear how to generalize them to exploit the extra information that multiple fibre reconstructions provide. Parker and Alexander [27] introduced a generalization of PICo that uses the peaks of PASMRI functions as fibre-orientation estimates and the sharpness of the peaks as indicators of uncertainty. That method neglects much of the information stored in the PAS functions. Here, we hypothesize that the functions output by multiple-fibre reconstructions contain much more information about the true fibre orientations. We show that the shape of the peaks reflects the local spread of fibres and, in particular, the peak anisotropy (i.e. how elliptical the peak cross section is) reflects local fibre bending or fanning. We go on to show how to exploit this extra information to improve the ability of PICo to recover fanning or bending structures by constructing a new general framework for PICo with multiple fibre reconstructions, extending Parker and Alexander's method [27]. We use the Bingham distri-

bution, which can have elliptical isocontours, to model the uncertainty of each fibre-orientation estimate. We estimate the shape of the Bingham distribution from the peak shape. Fanning and bending structures are widespread in the human brain and are notoriously hard to reconstruct with tractography. We show results on one important structure, the descending motor pathways, which is important in a wide range of neuroscience studies.

2. Background

This section describes the main approaches to estimating uncertainty for probabilistic tractography. We then introduce a family of algorithms, which we refer to as multiple-fibre reconstruction algorithms. We finish by describing extensions to the calibration procedure used in the PICo algorithm.

2.1. Probabilistic Tractography

Probabilistic tractography uses models of uncertainty in fibre-orientation estimates, along with standard streamline tractography, to calculate the probability of voxels being connected to a seed voxel. The procedure runs multiple streamline tracking processes from each seed point. For each streamline, each fibre-orientation estimate is randomly sampled from a probability density function (PDF) that models the uncertainty of the fibre orientations. The index of connectivity of any other voxel to the seed voxel is the fraction of streamlines that pass through that voxel.

Several alternative methods for modelling the uncertainty of fibre-orientation estimates have been developed. Parker et al [10, 26, 27], Cook et al [23] and Lazar et al [18] use a calibration procedure. They construct a mapping from some rotationally invariant feature of the diffusion to the variance of the fibre-orientation estimate using simulations. In the simplest algorithm [10], for example, Parker et al create a population of deflection angles between a known fibre direction and an estimate reconstructed by fitting the diffusion tensor (DT) to noisy synthetic data. They model the deflection angles with a Gaussian distribution and repeat for several levels of anisotropy in the synthetic data. Finally, they fit a linear model of the relationship between the FA and the variance of the Gaussian model, which they use to predict the variance in each voxel during tractography.

Behrens et al [31], Friman et al [9] and Hosey et al [34] use a Bayesian approach combined with a multi-compartment model of diffusion and sample from the marginal posterior distribution of the fibre orientation directly. Jones et al [16, 17], Lazar [19] and Haroon et al [13] use statistical bootstrap techniques to obtain samples of the fibre-orientation estimate distribution.

The Bayesian and bootstrap methods require explicit and parsimonious models relating the fibre orientation to the measurements and are limited by the oversimplicity of these models. They do not extend naturally to exploit the general class of multiple fibre reconstructions we aim to use here. Here we use the calibration approach, which does extend naturally in that way.

2.2. Multiple Fibre Reconstruction Algorithms

Multiple-fibre reconstruction algorithms can model more than one fibre direction in each voxel. The algorithms all output a function of the sphere in each voxel that reflects the local distribution of fibre orientations. The PASMRI reconstruction algorithm calculates the Persistant Angular Structure of the particle displacements density p. The PAS is the function on a spherical shell embedded in three-dimensional space whose Fourier transform matches the observed diffusion MRI measurements [15]. Thus, the PAS $\tilde{p}$ is defined by

$$A(\mathbf{q}) = r^2 \int \tilde{p}(\hat{\mathbf{x}}) \cos(r\mathbf{q} \cdot \hat{\mathbf{x}}) \mathrm{d}\hat{\mathbf{x}}, \tag{1}$$

where r is the radius of the sphere on which $\tilde{p}$ is embedded in 3-space, $A(\mathbf{q})$ is the normalised diffusion-weighted measurement for wavevector $\mathbf{q}$.

QBall is an approximation of Diffusion Spectrum Imaging (DSI) [36]. DSI calculates the orientation distribution function (ODF) ϕ of p by projecting a discrete representation of p onto the unit sphere:

$$\phi(\hat{\mathbf{x}}) = \int_0^\infty p(\alpha\hat{\mathbf{x}}) \mathrm{d}\alpha. \tag{2}$$

QBall approximates the ODF by using the Funk-Radon of the diffusion signal on a sphere in Fourier space. The Funk Radon Transform of a spherical function at a point $\hat{\mathbf{x}}$ on the sphere is the integral of the function over the equator that lies in the plane perpendicular to $\hat{\mathbf{x}}$, through the origin. This approximation is the convolution of the true ODF with a Bessel function [35].

Both PASMRI and QBall compute a feature of the particle displacement density, which is assumed to reflect the underlying distribution of fibre-orientations, although the relationship is not clear. Spherical Deconvolution (SD) [14] aims to reconstruct the fibre orientation distribution (FOD) directly. This method makes the assumption that the measurement, $A(\mathbf{q})$, is equivalent to the response $R(\mathbf{q}; \hat{\mathbf{x}})$ from a single fibre with orientation $\hat{\mathbf{x}}$ convolved with the distribution of fibre-orientations present in the voxel. The methods use a model of R to deconvolve the signal to obtain the fibre orientation distribution. However, the relationship between the output and the true fibre orientation distribution is still complex and unclear, because the model of R is necessarily over simple.

All of these methods provide a function of the sphere with peaks in dominant fibre directions. The functions reflect the shape of the underlying fibre-orientation distribution to some extent and some authors, eg [21], use them directly as a PDF on the next step direction in tractography. However, none of the functions provide the true distribution of the fibre orientations and the mapping between these quantities may be very complex. In particular, the shapes of the spherical functions that the methods return are very different and depend on parameter settings in the reconstruction algorithms; see figure 1 and [1] for examples. One feature of the functions that does appear reliable, however, is the peak directions, which usually provide good estimates of the dominant fibre orientations.

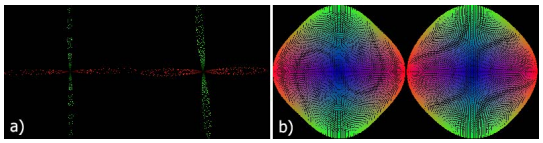

Figure 1. Two voxels of synthetic data reconstructed using a) PASMRI b) QBall. Each algorithm computes a different mathematical feature of the true PDF. As a result, the reconstructions differ in shape depending upon the method chosen to reconstruct.

2.3. Extensions to PICo

Cook et al [23] use calibrated PICo with more standard spherical distributions such as the Watson and Bingham distributions [22] instead of Parker et al's Gaussian to model fibre-orientation uncertainty. They evaluate their method using synthetic data and within PICo tractography tasks on brain data. They show that both the Watson and Bingham distributions are better for modelling the uncertainty of fibre-orientation estimates than Parker et al's Gaussian model. In particular, the Bingham model captures anisotropy in the uncertainty.

Both Parker and Alexander [26] and Cook et al [24, 6] generalize calibrated PICo to use multi-tensor models, which can improve tracking through fibre-crossings. However, this framework suffers from the limitations of the multi-tensor models such as fitting problems and the need to prespecify the number of fibres per voxel. Hosey et al [34] and Behrens et al [33] extend the Bayesian approach from [31] to model up to two fibre-orientation estimates within each voxel.

Parker and Alexander [27] extend calibrated PICo to exploit multiple fibre reconstruction algorithms. They use peak directions as fibre-orientation estimates, of which there may be several in each voxel. In place of the FA of the diffusion tensor, they use the sharpness of the fibre orientation distribution peak to predict the uncertainty of the estimate. Broad peaks have a high uncertainty, sharp peaks have a low uncertainty. Specifically, they compute the Hessian, or second derivative matrix, of the PAS function at each peak and use its trace as a measure of mean curvature. The uncertainty mapping is indexed using the log of the trace of the Hessian of each peak to find the variance in a Gaussian model of uncertainty.

A weakness of the approach in [27] is that the Gaussian model does not account for anisotropy in the uncertainty of the fibre-orientation estimate. Anisotropy in the uncertainty occurs in regions of fanning and bending where the fibres spread out more in one direction than another. Isotropic uncertainty models fail to capture this structure, which compromises reconstruction of these kinds of architectures. Here, we show that the shape of the peaks of functions from multi-fibre reconstructions reflects anisotropy in the spreading of fibre directions and we can use the peak anisotropy to predict uncertainty anisotropy.

3. Methods

This section describes how to calculate the shape of the cross-section of peaks from multiple fibre reconstructions and introduces the peak anisotropy, which is a measure of the eccentricity of the cross section of a peak. We then give details of the bi-polar Bingham distribution, which can model anisotropy in a spherical distribution. Finally, we describe a calibration procedure that exploits peak shape to improve estimates of uncertainty required for PICo tractography.

3.1. Peak Anisotropy

Peak anisotropy describes how elliptical the cross-sections of the peaks of multiple-fibre reconstructions are. We hypothesize here that the peak anisotropy reflects the underlying distribution of fibre-orientations. In particular, where there is a fanning or bending structure present in a voxel, the peaks of the function will have a highly elliptical cross-section that corresponds to the distribution of fibre-orientations present. We use the Hessian to describe each peak. Figure 2a shows an illustration of the peak shape for each peak of a multiple-fibre reconstruction. We compute a separate peak Hessian for each peak in the reconstruction.

To calculate the peak anisotropy, we generalize the standard definition of the fractional anisotropy [2] to

$$FA = \sqrt{\frac{N \sum (\lambda_i - \bar{\lambda})^2}{(N-1) \sum \lambda_i^2}}, \tag{3}$$

where λ_i, $i = 1, ..., N$ are eigenvalues of some N-dimensional quadratic form and $\bar{\lambda}$ is their mean. For $N = 3$, we get the familiar FA for three-dimensional DTs proposed in [2]. Here we use $N = 2$ to compute the anisotropy of the Hessian at PAS or ODF peaks. We refer to this quantity as the *peak anisotropy*.

We compute the FA of the two minor eigenvalues of the diffusion tensor, which provides an analogue to peak anisotropy for DT-MRI. The FA of the minor eigenvalues quantifies diffusion anisotropy orthogonal to the principal direction of the diffusion tensor.

3.2. Spherical Distributions

The Watson distribution [22] is

$$p(\mathbf{x}) = W(\mathbf{x}; \mathbf{z}, \kappa) = M\left(\frac{1}{2}, \frac{3}{2}, \kappa\right)^{-1} \exp[\kappa(\mathbf{z}\cdot\mathbf{x})^2], \quad (4)$$

where M denotes the confluent hypergeometric function of the first kind [22], the vector $\mathbf{z}$ is the mean orientation and κ controls the concentration of the distribution. The distribution is isotropic about $\mathbf{z}$. When $\kappa > 0$, the distribution is "bipolar" and is most concentrated at $\pm\mathbf{z}$. At $\kappa = 0$ the estimates are uniformly distributed over the sphere and when $\kappa < 0$ the distribution forms a girdle. The Bingham distribution is a generalisation of the Watson distribution with elliptical contours,

$$p(\mathbf{x}) = \frac{1}{M_2(\frac{1}{2}, \frac{3}{2}, \mathbf{A})} \exp[\kappa_1(\mathbf{z_1}\cdot\mathbf{x})^2 + \kappa_2(\mathbf{z_2}\cdot\mathbf{x})^2], \quad (5)$$

where

$$\mathbf{A} = (\mathbf{z_3}, \mathbf{z_2}, \mathbf{z_1}) \begin{bmatrix} \kappa_3 & 0 & 0 \\ 0 & \kappa_2 & 0 \\ 0 & 0 & \kappa_1 \end{bmatrix} (\mathbf{z_3}, \mathbf{z_2}, \mathbf{z_1})^T \quad (6)$$

and M_2 is the confluent hypergeometric function of the first kind for matrix argument [22]. There are two parameters, $\kappa_1 \leq \kappa_2 \leq 0$, that define the concentration of the distribution, since $\mathbf{A}$ and $\mathbf{A} + d\mathbf{I}$ give equivalent distributions for any real d, where $\mathbf{I}$ is the identity. Therefore it is customary to set $\kappa_3 = 0$. When $\frac{\kappa_2}{\kappa_1} \approx 1$ the distribution is circular. As $\frac{\kappa_2}{\kappa_1} \to 0$ the distribution becomes increasingly elliptical. The distribution is a girdle when $\kappa_2 = 0$, and when $\kappa_1 = \kappa_2 = 0$ the distribution is uniform. The mean of the distribution is $\mathbf{z}_1 \times \mathbf{z}_2$ and $\mathbf{z}_1$ and $\mathbf{z}_2$ are the principal axes of the elliptical contours of the PDF.

3.3. Calibration

For calibration, we construct a mapping from two Hessian eigenvalues λ_1 and λ_2 to the Bingham parameters κ_1 and κ_2 using simulations on two-tensor mixture models with known peak directions. Specifically, for a large number of noisy trials we synthesize data from test functions and reconstruct fibre-orientation estimates and associated Hessian matrices using the multiple-fibre reconstruction of choice. In each trial, we rotate the true direction to a common reference frame and apply the same rotation to the reconstructed direction. We collect the rotated reconstruction directions into bins with similar Hessian eigenvalues. The bin size is set to 0.45 for the PASMRI calibration and 0.2 for the QBall calibration, and the bins are indexed using the log of the Hessian eigenvalues. We fit the parameters of the Bingham distribution in each bin containing 50 or more samples. Finally, we fit linear surfaces to the log of each Bingham parameter as a function of the log of the Hessian eigenvalues. We create two calibration mappings; one for the voxels containing a single fibre population and one for voxels containing two fibre populations. In voxels where more than two fibre populations are reconstructed we use the two-fibre calibration mapping to estimate uncertainty.

The synthetic data come from test functions representative of one or two fibre orientations per voxel. The datasets containing two fibre orientations per voxel were generated using variations of the test function

$$p(\mathbf{x}) = aG(\mathbf{x}; \mathbf{D}_1, t) + (1-a)G(\mathbf{x}, \mathbf{R}_\theta^T \mathbf{D}_2 \mathbf{R}_\theta, t), \quad (7)$$

where a is a mixing parameter, $G(\mathbf{x}; \mathbf{D}, t)$ is a zero mean 3D Gaussian with covariance $2t\mathbf{D}$, $\mathbf{D}_1 = diag(\lambda_1, \lambda_2, \lambda_2)$, $\mathbf{D}_2 = diag(\lambda_2, \lambda_1, \lambda_2)$ and $\mathbf{R}_\theta$ is a rotation by θ about the z-axis. We use combinations of $\lambda_2 \in [1, 5] \times 10^{-10}\,\mathrm{m}^2\,\mathrm{s}^{-1}$, $a \in [0.1, 0.5]$ and $\theta \in [0, 45^o]$. We set $Tr(\mathbf{D}) = \lambda_1 + 2\lambda_2 = 2.1 \times 10^{-9}\,\mathrm{m}^2\,\mathrm{s}^{-1}$, which is approximately the value expected in brain data. For datasets containing a single fibre orientation per voxel, the test function is $p(\mathbf{x}) = G(\mathbf{x}; \mathbf{D}_1, t)$, where $\lambda_2 \in [1, 5] \times 10^{-10}\,\mathrm{m}^2\,\mathrm{s}^{-1}$. The parameters are varied between their minimum and maximum values to create all possible variations of the test function. A random rotation is then applied to the test function to remove any directional bias due to the chosen acquisition scheme. The data is synthesized by sampling the Fourier transform of p at each wavenumber in a spherical acquisition scheme according to the standard Fourier relationship between the particle displacement density p and the diffusion weighted measurement [5]. We add Rician noise to make S, the SNR at $b = 0$ equal to 20, which is approximately the same value as the average SNR of the white matter in the $b = 0$ image of each brain volume.

The synthetic data used for calibrating the original diffusion tensor implementation of PICo is constructed using the method described above to generate data from a single Gaussian model; see Cook et al [23] for details.

Figure 2b shows an example calibration mapping. Each of the surfaces maps the Hessian eigenvalues of a peak to one of the Bingham parameters. In this example, the green surface maps λ_1 and λ_2 to κ_1 and the blue surface maps λ_1 and λ_2 to κ_2. Only the highlighted half of the surfaces is ever used, since $\lambda_1 \geq \lambda_2$. Thus, to estimate the Bingham parameters of the uncertainty model for a peak with Hessian eigenvalues λ_1 and λ_2 from the calibration mapping, we set

$$\kappa_1 = -\exp(c_1^{(1)} + c_2^{(1)}\log(\lambda_1) + c_3^{(1)}\log(\lambda_2)) \quad (8)$$

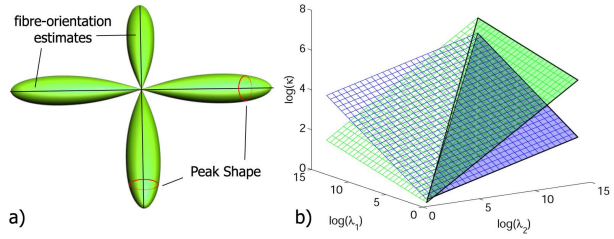

Figure 2. a) Features of the multiple-fibre reconstructions. For each peak there is a fibre-orientation estimate (black line) and a peak Hessian, which describes the shape of the peak cross-section (red circle). b) Example of calibration mapping surfaces. The calibration outputs two linear mappings from the Hessian eigenvalues, $\log(\lambda_1)$ and $\log(\lambda_2)$, and the Bingham parameters $\log(-\kappa_1)$ (green surface) and $\log(-\kappa_2)$ (blue surface).

$$\kappa_2 = -\exp(c_1^{(2)} + c_2^{(2)}\log(\lambda_1) + c_3^{(2)}\log(\lambda_2)), \qquad (9)$$

where $c_i^{(1)}$ and $c_i^{(2)}$, $i = 1, 2, 3$, are the coefficients for the two linear surfaces. During tractography we draw vectors from the Bingham distribution with shape parameters κ_1 and κ_2, using the fibre-orientation estimate as the mean of the distribution.

3.4. Human Brain Data

High angular resolution diffusion-weighted imaging (HARDI) data were acquired on a 3 T Philips Achieva scanner using an 8-element SENSE head coil. A pulsed gradient spin echo (PGSE) echo planar imaging (EPI) sequence was implemented with TE=54ms, TR=6000ms, 112×112 matrix reconstructed to 128×128 using zero filling, reconstructed resolution $1.836, \text{mm} \times 1.836, \text{mm}$, slice thickness $2.1, \text{mm}$, 34 contiguous slices, 61 gradients at $b = 1200\,\text{s}\,\text{mm}^{-2}$, 1 at b = 0, SENSE acceleration factor = 2.5. The total imaging time for each HARDI acquisition was approximately 7 minutes. This acquisition was repeated 8 times provide a conventional bootstrapping dataset. All diffusion-sensitised images were registered to the corresponding b = 0 image within each slice location to remove eddy current-induced distortions and for all scanning repetitions to the first scan, using the 6-degrees-of-freedom 2D registration schedule file available in FSL's FLIRT. The average SNR in the white-matter regions of the $b = 0$ image is 20. For tractography, we acquired a separate dataset with 60 contiguous axial slices, which provides complete brain coverage. The imaging parameters for the tractography data are as above except that TR=11884ms.

4. Experiments and Results

In this section, we show that the peaks of multiple-fibre reconstructions do provide useful information that can be used to improve tractography results. We validate our calibration procedure by comparing the Bingham parameter estimates from our calibration procedure to estimates generated using the bootstrap method for several regions of interest in human brain data. PICo tractography results are shown for human brain data seeded at the base of the corticospinal/corticopontine tracts. For the QBall reconstruction we use spherical harmonic basis functions [4], since the spherical integral can be computed analytically and therefore does not incur the numerical inaccuracies of the radial basis function implementation described in [35]. We use the PASMRI standard settings listed in [15]. The peaks of the PAS or ODF in each voxel is found using the peak-finding algorithm described in [15]. The Hessians of the PAS and ODF functions at each peak are calculated numerically.

4.1. Correlation between Hessian anisotropy and fanning in the human brain

We start by testing the hypothesis that the peaks of the PAS or ODF contain more useful information than simply the peak sharpness and direction by generating images of the peak anisotropy of the dominant peaks of the PAS and the ODF in each voxel, as well as the anisotropy of the DT perpendicular to its principal axis. The peak anisoptropy images (figure 3) generated are colour-coded to show the direction of the anisotropy using the principal eigenvector of the peak Hessian of the spherical functions and the second eigenvector of the DT respectively. Both point in the direction of the largest anisotropy of the peak cross-section, which we hypothesize is the direction of the fibre-spread/uncertainty.

Figure 3 shows the magnitude and direction of the peak anisotropy of the dominant peak from the DT (b), ODF (c) and PAS (d) in each voxel. In figures 3b and 3c, light pixels indicate that the cross-section of the principal peak of the reconstruction is elliptical; dark regions indicate that the peak-cross section is circular. Where there are multiple peaks, figures 3b and 3c show only the anisotopy of the largest peak. The peak anisotropy at the centre of the corpus callosum (highlighted with the upper box) is low, since the fibres are approximately parallel in this region and any spread is isotropic. Note that this is in stark contrast to the diffusion tensor anisotropy (Fractional Anisotropy) in the corpus callosum, which is usually very high. However, the PAS peaks in the descending motor pathways (highlighted with the left box) have high anisotropy with the largest axis in the posterior-anterior direction (green), which, according to anatomical knowledge [12], is the direction of the fanning of this structure. The ODF peak anisotropy image highlights the fanning structure less clearly than PAS. This is a result of the fourth-order spherical-harmonic representation of the ODF, which cannot model anisotropy in the peak sufficiently well. However, both the PASMRI and QBall peak anisotropy maps broadly agree. In the case of the DT, the

FA of the two-dimensional DT orthogonal to the dominant fibre direction generally agrees with the other two images.

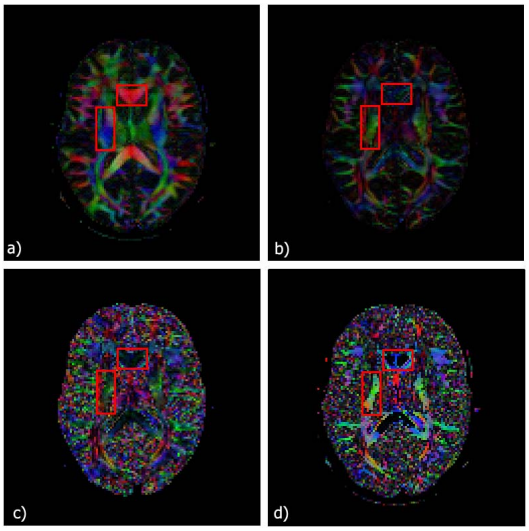

Figure 3. images of a) the standard colour-coded principal direction map from the DT weighted by FA of the full three dimensional DT. Red indicates left-right; green indicates front to back; blue top to bottom; intensity is FA. b) Colour-coded map of the second eigenvector weighted by the FA of the two-dimensional DT orthogonal to the dominant fibre direction. c) Colour-coded peak major axis orientation weighted by peak anisotropy for the dominant ODF peak. d) Colour-coded peak major axis orientation weighted by peak anisotropy for the dominant PAS peak. In cases where the PAS or ODF has several peaks, the peak anisotropy of the peak with the greater magnitude is shown. Note that the peak major axis describes the orientation of the elliptical peak cross-section rather than the fibre orientation.

4.2. Calibration validation

We validate the estimates of the Bingham distribution parameters provided by the calibration mapping by comparison with bootstrap samples from the 8 repeats of the human brain data. We identify 4 regions of interest (ROIs), one containing voxels with highly coherent fibre bundles, one with fibre-crossings, one in a fanning region and one in an isotropic region. Each ROI contains 25 voxels. The ROIs are shown in figure 4. Each ROI was extracted from all 8 datasets and we create 100 bootstrap sets of measurements for each voxel by sampling with replacement from the 8 options for each measurement [8]. We reconstruct using PASMRI and QBall and estimate the fibre-orientation estimates and associated Hessians for each bootstrap. To estimate the true uncertainty anisotropy, we fit a Bingham distribution to the set of all peak directions from all 100 bootstraps and compute the 2D FA of the Bingham parameters $c_b = FA(diag(\kappa_1, \kappa_2))$. Then, for each bootstrap, we use the calibration to estimate the Bingham parameters. We draw 100 samples from all of the Bingham PDFs in each bootstrap, fit a single Bingham PDF to the whole colleection of these samples and compute $\bar{c}_c$, the average 2D FA. In cases where there are several peaks, it is likely that samples will be drawn from both Bingham distributions during tractography. We fit a single Bingham to both populations when computing c_b to capture this effect. For effective calibration, we should find that $c_b \simeq \bar{c}_c$. Figures 5a and 5b plot c_b against $\bar{c}_c$ for each voxel in each ROI for PASMRI and QBall, respectively.

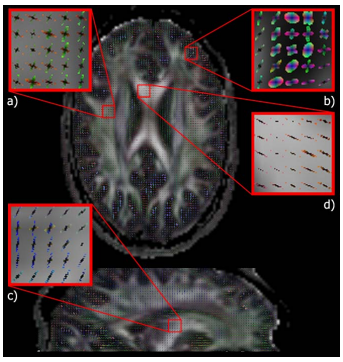

Figure 4. The regions of interest used for validating the performance of the fitted mapping (the PAS is shown in each voxel of the ROIs). The ROIs were chosen to include a) crossing fibre, b) grey matter, c) fanning and d) single fibre regions of the brain.

For QBall, the FA of the Bingham parameters from the calibration (figure 5b) are generally very low $(0.1 - 0.2)$ in comparison to those from the bootstrapped estimate, although the calibration FA improves for grey matter and fibre-crossing regions. The fourth-order spherical-harmonic basis function representation of the ODF cannot adequately capture the anisotropy in the peaks of the QBall ODF. The PASMRI calibration gives better results (figure 5a), with c_b and $\bar{c}_c$ generally agreeing. Some disagreement occurs for the high FA samples, where bootstrapping gives much higher anisotropy. Occasionally, when the uncertainty is anisotropic, PASMRI produces small peaks in the perpendicular direction rather than a single elliptical peak. The small spurious peak causes unexpectedly low $\bar{c}_c$ in the main peak. Future work may correct for this by including the peak height in the calibration mapping.

4.3. Tractography

We run PICo tractography on the human brain data using the new multiple-fibre PICo algorithm in conjunction with PASMRI and QBall and compare the results to those of the algorithm described in [26]. The tractography process was started from a manually defined region of interest at the base of the corticospinal/corticopontine tracts (see Figure 6, inset). The dominant pathway from the seed region runs

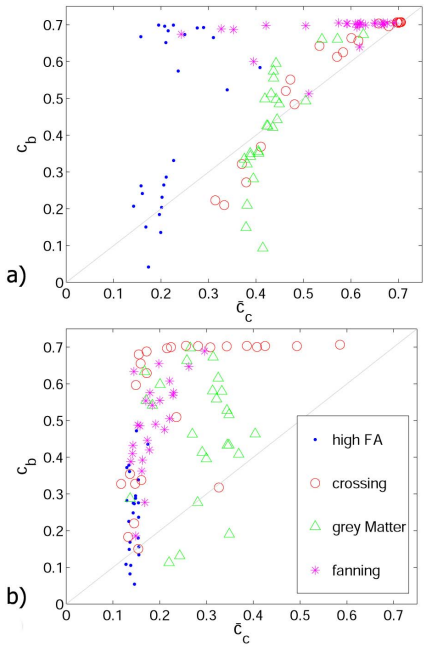

Figure 5. Validation results for calibration mapping estimates of the Bingham parameters. a) and c) are plots of the mean FA, $\bar{c}_c$, of the Bingham parameters estimated by our mapping against the FA, c_b of the Bingham parameters estimated using bootstrapped data for both PAS and QBall respectively.

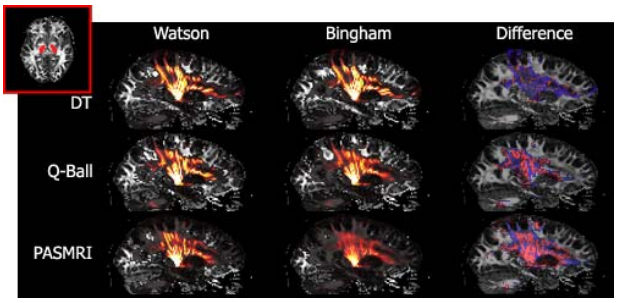

Figure 6. PICo tractography connectivity maps using a DT reconstruction (top row), QBall reconstruction (middle row) and PASMRI reconstruction (bottom row) made semi-transparent and overlaid onto FA images. The difference map shows where the probability of connection is higher when using the Bingham distribution (blue) or the Watson distribution (red). The seed ROIs are shown in the axial view in the top-left corner (inset).

inferior-superior into the corona radiata, where the decending motor pathway fibres cross lateral fibres projecting from the corpus callosum. Figure 6 shows the results of tractography using the multi-fibre algorithm with both PASMRI and QBall as well as results from the DT algorithm introduced by [23] using both the Watson and Bingham distributions to model uncertainty. As expected, the original DT-PICo algorithm fails at fibre-crossings, which results in large holes in the descending motor pathway reconstructions. In the DT case, the connection probabilities in the descending motor pathways vary widely, as opposed to the connection probabilities from PAS-PICo and QBall-PICo, which are more consistent over the tract. For QBall, the PICo results have fewer holes than the DT-PICo results, since the QBall ODFs are able to model multiple peaks. However, the reconstruction of the descending motor pathways generally fails to reach the cortical layer. The results of the PAS-PICo tractography experiment are more favorable. Not only does the algorithm correctly map the fanning structure of the descending motor pathways, the reconstructed tracts extend to the cortical layer. The PAS-PICo results show fewer holes in the descending motor pathways than those from QBall PICo. These results reflect the greater ability of PASMRI to resolve fibre-crossings with a small number of measurements over QBall [1]. The difference map from the PAS-PICo experiment shows that using the information about the shape of the peak results in a more even probability of connection over the descending motor pathways. Although the overall magnitude of the connection probabilities appears lower, this is merely because the total connectivity is the same but it is now more evenly spread across the structure.

5. Discussion and Conclusions

We have introduced a generalization of the PICo algorithm that allows us to use more of the information contained in the fibre-orientation distributions to improve tractography through complex white-matter structures. We have demonstrated that the algorithm described here is completely general and works for both PASMRI and QBall. It extends easily to similar algorithms such as Spherical Deconvolution. The bootstrap validation shows that the FA of the Bingham parameters from the PASMRI calibration and the bootstrap estimates approximately correlate. For QBall-PICo, the choice of PDF used to model uncertainty makes little difference to the resulting PICo connectivity map. This results from using a fourth-order spherical-harmonic basis function representation of the ODF. Using a spherical radial basis function representation of the ODF or higher-order SH may allow peak anisotropy to be captured better. PAS-PICo gives the most compelling results. The difference maps in figure 6 demonstrate clear changes by modelling anisotropy. The PAS-PICo difference map clearly shows that using both the shape and sharpness of the PAS peaks results in a higher connection probability in regions of fibre crossings. Utilizing other information in the peaks of multi-fibre reconstructions, such as the magnitude of the peaks of functions, in PICo may improve results further. The bootstrap validation highlights one weakness of the PASMRI reconstruction algorithm, which is that it sometimes produces a spurious perpendicular peak in fanning structures. Future work will investigate and compare the ability of other multiple

fibre reconstructions to capture the shape of fanning and bending structures. We also aim to compare the uncertainty estimates using our calibration approach to those from the Bayesian and bootstrap approaches. These implementations of PICo are freely available in Camino [25], an open-source diffusion MRI toolkit.

References

[1] D. C. Alexander. Multiple-fibre reconstruction algorithms for diffusion MRI. *Annals of the NYAS*, 1046:113–133, 2005. 3, 7

[2] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI. *J. Magn. Reson. B*, 111:209–219, 1996. 3

[3] C. A. Clark et al. White matter fiber tracking in patients with space-occupying lesions of the brain: a new technique for neurosurgical planning? *NeuroImage*, 20:1601–1608, 2003. 1

[4] C. R. Hess et al. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic. *JMRI*, 56:104–117, 2006. 5

[5] P. T. Callaghen. *Principles of Magnetic Resonance Imaging*. Oxford: Oxford Science Publications, 1991. 4

[6] P. A. Cook and D. C. Alexander. Modelling uncertainty in two fibre-orientation estimates within a voxel. In *Proc ISMRM*, page 1629, Seattle, 2006. 3

[7] D. S. Tuch et al. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *MRM*, 48:577–582, 2002. 1

[8] B. Efron. Bootstrap methods: another look at the jackknife. *Annuals of Statistics*, 1046:113–133, 2005. 6

[9] O. Friman and C.-F. Westin. Uncertainty in fiber tractography. Volume 3749 of Lecture Notes in Computer Science, pages 107–114. Edited by James S. Duncan and Guido Gerig. Springer, 2005. 2

[10] G. J. M. Parker et al. A framework for a streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *JMRI*, 18:242–254, 2003. 1, 2

[11] G. J. M. Parker et al. Lateralization of ventral and dorsal auditory-language pathways in the human brain. *NeuroImage*, 24:656–666, 2005. 1

[12] H. Grey. Greys anatomy (figure 684). http://www.bartleby.com/107/187.html. 5

[13] H. A. Haroon and G. J. Parker. Using the wild bootstrap to quantify uncertainty in fibre orientations from q-ball analysis. In *Proc ISMRM*, Berlin, 2007. 2

[14] J. D. Tournier et al. Direct estimation of the fibre orientation density function from diffusion-weighted MRI data using spherical deconvolution. *NeuroImage*, 23:1176–1185, 2004. 1, 2

[15] K. M. Jansons and D. C. Alexander. Persistent angular structure: new insights into diffusion MRI. *Inverse Problems*, 19:1031–1046, 2003. 1, 2, 5

[16] D. K. Jones. Tractography gone wild: Probabilistic tracking using the wild bootstrap. In *Proc ISMRM*, page 435, Seattle, 2006. 2

[17] D. K. Jones and C. Pierpaoli. Confidence mapping in diffusion tensor magnetic resonance imaging tractography using a bootstrap approach. *MRM*, 53:1143–1149, 2005. 1, 2

[18] M. Lazar and A. L. Alexander. White matter tractography using random vector (rave) perturbation. In *Proc ISMRM*, page 523, Hawai'i, 2002. 1, 2

[19] M. Lazar and A. L. Alexander. Bootstrap white matter tractography (boot-trac). *NeuroImage*, 24:524–532, 2005. 2

[20] M. Lazar et al. A white matter tractography study of white matter reorganization after surgical resection of brain neoplasms. In *Proc ISMRM*, page 1259, Kyoto, 2004. 1

[21] M. Perrin et al. Fiber tracking in q-ball fields using regularized particle trajectories. In *Proc IPMI*, Spring Valley, 2005. 3

[22] K. V. Mardia and P. E. Jupp. *Directional Statistics*. Wiley, 2000. 3, 4

[23] P. A. Cook et al. Modelling noise-induced fibre-orientation error in diffusion tensor MRI. In *Proceedings of the IEEE ISBE*, pages 332–336, Arlington, 2004. 2, 3, 4, 7

[24] P. A. Cook et al. An automated approach to connectivity-based partitioning of brain structures. In *Proc MICCAI*, pages 164–171, Palm Springs, 2005. 3

[25] P. A. Cook et al. Camino: open source diffusion-MRI reconstruction and processing. In *Proc ISMRM*, page 2759, Seattle, 2006. 8

[26] G. J. M. Parker and D. C. Alexander. Probabilistic monte-carlo based mapping of cerebral connections utilizing whole-brain crossing fibre information. In *Proc IPMI*, pages 684–695, 2003. 2, 3, 6

[27] G. J. M. Parker and D. C. Alexander. Probabilistic anatomical connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Phil Trans Roy Soc*, 360:893–902, 2005. 1, 2, 3

[28] Q. Dong et al. Clinical applications of diffusion tensor imaging. *JMRI*, 66:259–267, 1994. 1

[29] S. Mori et al. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45:265–269, 1999. 1

[30] T. E. Conturo et al. Tracking neuronal fiber pathways in the living human brain. In *Proc NAS*, volume 96, pages 10422–10427, USA, 1999. 1

[31] T. E. J. Behrens et al. Characterization and propagation of uncertainty in diffusion-weighted mr imaging. *MRM*, 50:1077–1088, 2003. 1, 2, 3

[32] T. E. J. Behrens et al. Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging. *Nature Neuroscience*, 6:750–757, 2003. 1

[33] T. E. J. Behrens et al. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34:144–155, 2007. 3

[34] T. Hosey et al. Inference of multiple fiber orientations in high angular resolution diffusion imaging. *MRM*, 54:1480–1489, 2005. 2, 3

[35] D. S. Tuch. Q-Ball Imaging. *MRM*, 52:1358–1372, 2004. 1, 2, 5

[36] V. J. Weeden et al. Mapping fibre orientation spectra in cerebral white matter with fourier-transform diffusion MRI. In *Proc ISMRM*, pages 259–267, Philadelphia, 1999. 2