Modelling uncertainty in two fibre-orientation estimates within a voxel

P. A. Cook[1], D. C. Alexander[1]

[1]Centre for Medical Image Computing, Department of Computer Science, University College London, London, United Kingdom

Introduction The Probabilistic Index of Connectivity (PICo) [1] framework uses Monte-Carlo streamline generation to create maps of connection probability. A probability density function (PDF) is associated with the fibre orientation in each voxel. For each Monte-Carlo iteration, the estimated fibre orientation is aligned with a sample from the PDF and streamline tractography provides a sample fibre tract. After N iterations, the probability of connection depends on the number of streamlines that cross through each voxel.

We calibrate the PICo PDFs by adding noise to synthetic data and extracting a large sample of independent fibre orientation estimates. We fit the parameters of the PDF to these estimates. Assuming a zero-mean Gaussian model for the diffusion displacement density p = G(D, t) in each voxel, where D is the diffusion tensor and t the diffusion time, we use a range of test functions with different diffusion tensors to produce a lookup table that maps tensor anisotropy to PDF concentration.

When there are two fibre bundles present in a voxel and there is sufficient data, it is possible to fit a Gaussian mixture model [2] with two tensors in each voxel. In this case, the calibration of N trials yields $2N$ fibre orientation estimates, one for each bundle at each trial. Parker and Alexander [3] separate the fibre orientation estimates into two groups and use a Gaussian model of uncertainty for each. Cook et al [4] fit two Watson PDFs [5] simultaneously, without making an assignment of fibres into groups. The methods of Parker and Cook fit PDFs that have circular contours on the sphere. However, certain variants of the Gaussian mixture model leads to PDFs with elliptical contours on the sphere. We propose a new PDF to explicitly model this case.

Methods We implement the PICo calibration algorithm by fitting three variants of the Gaussian mixture model. In each case, the test function is the same, $p = 0.5G(D_a, t) + 0.5G(D_b, t)$. The diffusion tensors in the test function are identically anisotropic, cylindrically symmetric and orthogonally oriented. We emulate a scanner sequence of 54 diffusion-weighted measurements distributed isotropically on a sphere in q-space (b = 1150 s mm^{-2}), with 6 unweighted (b = 0) images. The signal to noise ratio in the unweighted images is 14, which is similar to that found in white matter in brain scans using this scanner sequence. For each noisy set of measurements, we fit a Gaussian mixture model: $p = \alpha\, G(D_1, t) + (1-\alpha)G(D_2,t)$. We use three different inversion strategies. For the full inversion, we allow all 13 parameters (six parameters of D_1, six of D_2, and α) to vary. For the cylindrical inversion, we restrict the diffusion tensors to be cylindrically symmetric, so the second and third eigenvalues are the same. This reduces the number of free parameters from 13 to 9. For the restricted inversion, we enforce cylindrical symmetry and fix the mixing parameter $\alpha = 0.5$, so there are 8 free parameters.

In a similar way to Cook et al [4], we fit two spherical PDFs to the fibre orientation estimates $\mathbf{x}_1 \ldots \mathbf{x}_{2N}$, however we use the Bingham distribution $f(\pm\mathbf{x}) = B(\pm\mathbf{x};\ \kappa_1, \kappa_2, \mu_1, \mu_2) = M^{-1}\exp[\kappa_1(\mu_1^T\mathbf{x})^2 + \kappa_2(\mu_2^T\mathbf{x})^2]$, $\kappa_1 < \kappa_2 \le 0$, [5], where the axes μ describe the orientation of the distribution, the parameters κ describe the concentration, and M is a constant. The Bingham distribution does not assume circular contours on the sphere. Given the sample of fibre orientations we therefore maximise $l = \sum_{i=1}^{2N} 0.5B(\mathbf{x}_i;\kappa_1,\kappa_2,\mu_1,\mu_2) + 0.5B(\mathbf{x}_i;\kappa_3,\kappa_4\mu_3,\mu_4)$ with respect to the orientations and concentrations of the two Bingham PDFs.

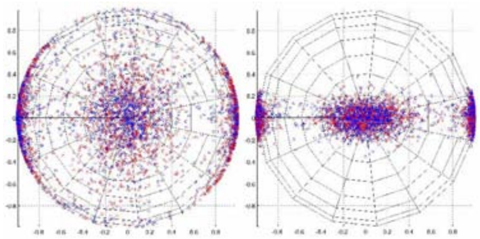

Fig 1: Estimated fibre orientations from the PICo calibration with the full (left) and restricted inversion (right). The raw data are the blue points; samples from the fitted Bingham PDFs are the red points.

Experiments Experiment 1 illustrates the elliptical contours of the fibre orientation PDF at intermediate anisotropy. We generate 1000 trials of the test function with the fractional anisotropy (FA) of D_a and D_b set to 0.6. Fig 1 shows the fibre orientation estimates for each of the inversions. The fibre orientations from the restricted inversion are more concentrated but distributed asymmetrically, with greater uncertainty towards the plane defined by the principal directions of D_a and D_b. The distribution of fibre orientations is similar for the restricted and cylindrical inversion. Fig 2 shows the histogram of the estimated anisotropy. The restricted model gives a much better estimate of the FA in this case. This is important for PICo because the calibration parameterises the PDF as a function of tensor anisotropy. We observe a similar result over a range of FA from 0.4 to 0.8. At FA = 0.9, the fibre orientations are clustered in circular patterns, and the FA estimates are similar from both inversions.

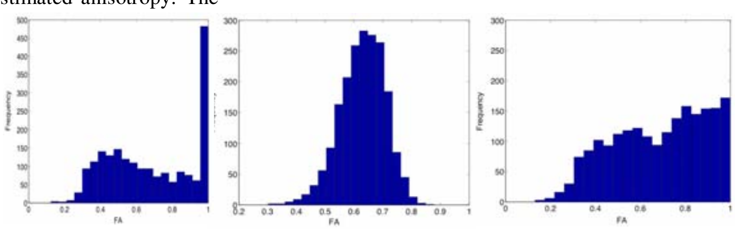

Fig 2: Histogram of FA of fitted tensors using the cylindrical (left) restricted (centre) and full (right) inversions.

Experiment 2 compares the two-Bingham and two-Watson model in a simulated fibre crossing. The crossing consists of two orthogonal fibre paths, which intersect through the centre of the image, forming a 20mm fibre-crossing region where both fibre bundles contribute equally to the signal. The imaging parameters are the same as in the first experiment, and the voxel dimensions are isotropic 2 mm. As a gold standard, we create a probability map using PICo without sampling from a PDF. We add noise to the test function at each of 5000 iterations and calculate the restricted inversion to obtain a fibre orientation estimate. We then run PICo in 50 different noisy images using the Watson and Bingham models, and compute the correlation between the 50 maps from each PDF and the gold standard. We carry out this experiment with two levels of noise, one with SNR=14 (as in experiment 1) and again with SNR=32. Fig. 3 plots the mean and standard error of $[C_B(i) - C_W(i)]$, where $C_B(i)$ is the correlation of the Bingham PICo map to the gold standard for map i and $C_W(i)$ is the same statistic for the Watson distribution. As expected, for very anisotropic fibres, both models perform similarly because the distribution of fibre orientation estimates in the fibre crossing is circular. However, at low anisotropy, the Bingham model correlates better to the gold standard. At SNR=14, this effect is not apparent at FA = 0.4 or 0.5, because the fibre crossing is not well resolved at this SNR and FA.

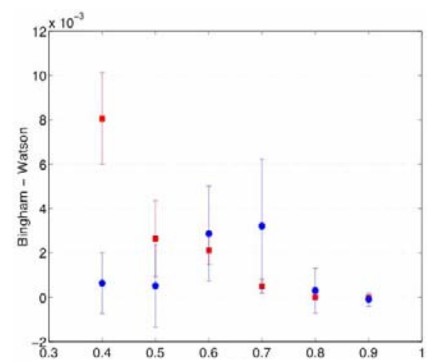

Fig 3: Correlation to the gold standard in a simulated fibre crossing. The results are plotted as squares for SNR = 32 and circles for SNR = 14.

Conclusion Our results suggest that the restricted inversion is more useful for PICo, but further work is required to understand the behaviour of this inversion when the mixing parameter in the test function is not 0.5, and when the diffusion tensors in the test function are not cylindrically symmetric. The two-Bingham model can describe a PDF with circular or elliptical contours on the sphere, as appropriate, making it more flexible than the previous two-Watson model proposed by Cook et al [4].

Acknowledgement We thank the UK Engineering and Physical Sciences Research Council for funding this work through GR/T22858/01.

References

1. G.J.M. Parker, et al, *JMRI* 18 (2003), 242–254
2. D.S. Tuch, *MRM* 48 (2002), 577–582
3. G.J.M. Parker, D. C. Alexander, *Proc, IPMI* (2003), 684-695
4. P.A. Cook, et al, *Proc. MICCAI* (2005), 164-171
5. K.V. Mardia, P.E, Jupp, *Directional Statistics*, Wiley, 2000