Super-resolving patches in diffusion MRI using canonical fibre configurations

S. J. Prince[1], and D. C. Alexander[1]

[1]Computer Science, University College London, London, London, United Kingdom

Introduction: The utility of fibre-orientation estimates from diffusion MRI is limited by spatial resolution. This work aims to improve resolution by learning and exploiting prior knowledge about image structure. In particular, we hypothesize that fibre configurations are fractal in nature and exhibit self-similarity over scale. We learn a set of canonical configurations at the original image scale and fit the configurations to each voxel to increase spatial resolution, here by a factor of three in each direction.

Methods: The initial set of fibre configurations for learning come from a standard DTI dataset with 61 gradient directions at b=1000smm^{-2} and 7 measurements with b=100smm^{-2}. We fit Behrens' ball and stick model [1] in each voxel, which provides an estimate of the fibre orientation $\mathbf{n}$, the diffusivity d and the intra-cellular volume fraction f in each voxel.

We extract all 3x3x3 voxel patches of the fitted parameters from within the brain region of the image. Each patch is represent by parameters $f_{1..27}$, $d_{1..27}$, $\mathbf{n}_{1..27}$. We reject patches where the mean volume fraction μ_f is less than 0.1. We align the average fibre direction in the remaining voxel patches by rotating all $\mathbf{n}$ in each patch so that the first two eigenvectors of the mean dyadic $27^{-1}\Sigma\mathbf{n}\mathbf{n}^T$ align with the x and y axis respectively.

We then cluster the aligned data to find canonical patch configurations. We choose a mixture of Watson distributions for the clustering:

$$Pr(\mathbf{n}_i|\boldsymbol{\mu}_{1\ldots27,1\ldots J},\kappa) = \sum_{j=1}^{J} w_j \prod_{i=1}^{27} \mathrm{Wat}\left[\mathbf{n}_i;\boldsymbol{\mu}_{ij},\kappa\right]$$

where w_j is the weight of the j^{th} of J clusters, $\boldsymbol{\mu}_{ij}$ is the mean of the Watson distribution in the i^{th} spatial position for the j^{th} cluster and κ is a shared dispersion parameter. We fit this model using 40 iterations of the expectation maximization algorithm.

To determine the fibre configuration at a new voxel from the measurements, we need to identify the best cluster, its average fibre direction and the volume fraction and diffusivity (assumed constant for all 27 sub-voxels). We seek the combination that minimizes the sum of square differences between simulated and actual measurements. We search exhaustively over clusters and perform a Levenberg-Marquardt optimization over the remaining parameters. Simulated measurements are calculated by taking the mean over the patch of the measurements predicted by the forward model (ball and stick) at each sub-voxel.

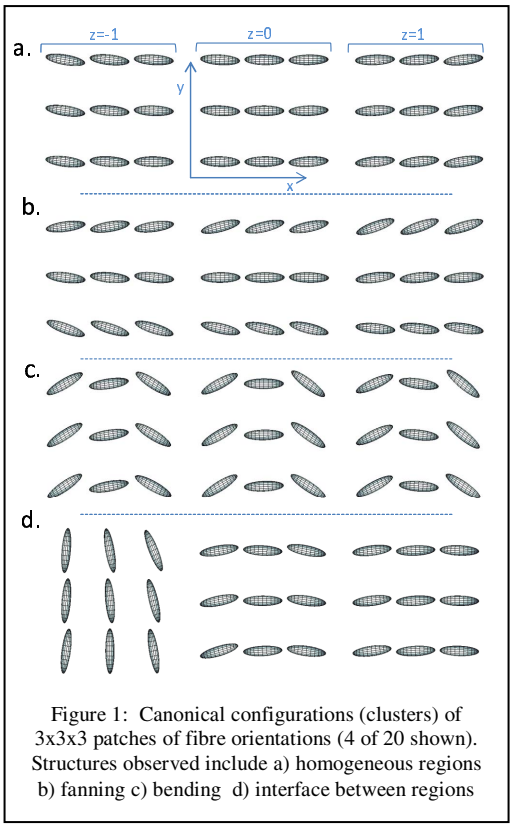

Figure 1: Canonical configurations (clusters) of 3x3x3 patches of fibre orientations (4 of 20 shown). Structures observed include a) homogeneous regions b) fanning c) bending d) interface between regions

Experiments: To test the approach, we try to recover the fibre configurations in a full resolution image after spatial downsampling by a factor of 3 in each direction. We fit the ball and stick model and extract 3x3x3 patches at this lower resolution, align and cluster the configurations. Examples of the resulting canonical fibre configurations are illustrated in Figure 1 for J=20 clusters.

We fit the clustered model to the downsampled measurements to generate an estimate of the higher resolution data. In Figure 2, we show a reconstructed slice of the brain, comparing high resolution ground truth (original data), our results and the fitted models at lower resolution (expanded to full resolution for comparison by nearest neighbour interpolation). There are some regions (shaded) in which our results clearly provide a better fit than the low resolution model. There are other areas where there is significant structure in the high resolution data that is not captured by either approximation.

Discussion: By building a model that captures the local structure of the direction field, it is possible to increase the spatial resolution of diffusion MRI. A limitation of our model is that it cannot distinguish between permutations of position within each fitted patch. This occasionally causes ambiguity leading to fitting errors. In future work, we will resolve this difficulty by modelling the spatial relationships between neighbouring canonical structures, which is a natural and simple extension of the framework we have presented.

Acknowledgement: This work was supported by EPSRC grant EP/E013309/1. Zoltan Nagy, UCL, provided the data. **References:** [1] Behrens MRM 2003.

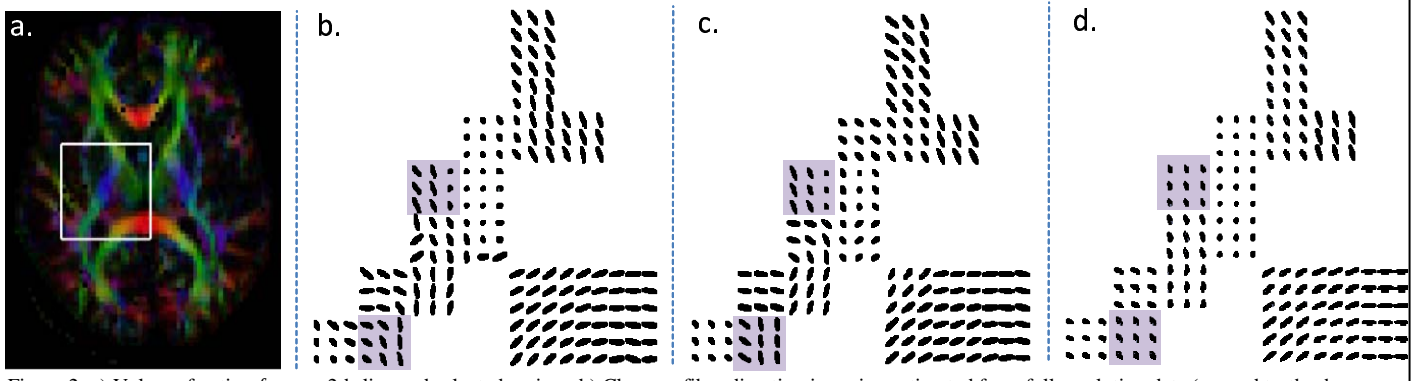

Figure 2: a) Volume fraction f over a 2d slice and selected region. b) Close up fibre direction in region estimated from full resolution data (ground truth, shown where volume fraction >0.3). c) Super-resolved from downsampled measurements and J=10 clusters (our approach). d) Downsampled data (repeated for comparison). Shaded patches indicate regions where our method c) successfully identifies sub-voxel structure in b) that is not represented in d).