Direct estimation of fibre orientations in partial volume contaminated regions using spherical deconvolution

J-D. Tournier[1], F. Calamante[1], D. G. Gadian[1], A. Connelly[1]

[1]Institute of Child Health, University College London, London, United Kingdom

Introduction: The diffusion tensor model is widely used to characterise the diffusion properties of brain tissue, but has been shown to be inadequate in partial volume contaminated regions [1-3]. Although a number of alternative models have been proposed, the direct computation of the orientations of the underlying fibres remains difficult. Fitting multiple tensors to the data has been proposed as a method to compute the fibre orientations [1]. However, this has been found to be unstable when more than two fibre populations are present, and relies on being able to determine the number of fibre populations present in each voxel prior to fitting. In this study, we propose a method that is able to estimate directly the distribution of fibre orientations within a voxel from high angular resolution data, without making assumptions regarding the number of fibre populations present.

Theory: In order to resolve different fibre populations with distinct orientations, we make the assumption that there is no exchange between these populations, such that the signals from each population can be assumed to add to the measured diffusion-weighted signal independently. We make the additional assumption that the diffusion characteristics of all fibre populations found in the brain are identical, and that variations in white matter fractional anisotropy (FA) can be attributed entirely to partial volume effects. The diffusion-weighted signal attenuation that would be measured if a single coherently oriented fibre population was present in the voxel can be represented by an axially symmetric *response function* $R(\theta)$. The problem can then be expressed as a convolution over the unit sphere of the response function $R(\theta)$ with a *fibre orientation density function* (fibre ODF) $F(\theta, \phi)$, as illustrated in figure 1 (see [4] for details on spherical convolution). In a high-angular resolution diffusion-weighted experiment [2], the diffusion-weighted signal attenuation $S(\theta, \phi)$ is sampled along a large number of directions. If the response function $R(\theta)$ is known, then the fibre ODF $F(\theta, \phi)$ can be obtained by deconvolving $R(\theta)$ from $S(\theta, \phi)$.

If $F(\theta, \phi)$ is decomposed into its spherical harmonic components and $R(\theta)$ into its rotational harmonic components, the spherical convolution operation becomes a simple set of matrix multiplications, much like the convolution of two functions in Cartesian space becomes a simple multiplication in Fourier space (see [4] for details). This operation can be trivially inverted to perform the spherical deconvolution operation.

Methods: Simulations were performed assuming the tensor model for the response function $R(\theta)$ (FA from 0.7 to 0.9), and a set of two delta functions for the fibre orientation density function $F(\theta, \phi)$ (separation angle from 0° to 90°). For each of 100 repetitions, the signal attenuation $S(\theta, \phi)$ was calculated for a set of N uniformly distributed diffusion-weighting gradient orientations [5] (N = 30, 60, or 100, b-value from 1000 to 10000 s/mm^2). Gaussian noise was added to these signals ($SNR_{b=0}$ from 20 to noiseless), which were then decomposed into their corresponding even spherical harmonic components by least-squares linear fitting [3], up to a maximum spherical harmonic order ranging from 6 to 10. The spherical harmonic decomposition of the estimated fibre ODF was finally obtained by spherical deconvolution as described above.

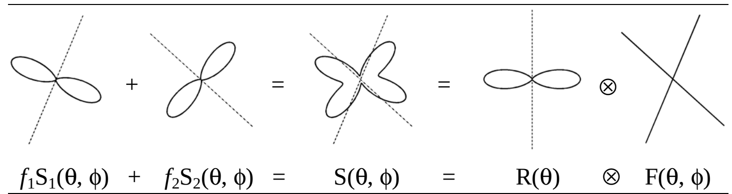

Figure 1: in this simple 2D illustration, the voxel contains two fibre populations with distinct orientations (θ_1, ϕ_1) and (θ_2, ϕ_2) and volume fractions f_1 and f_2. The diffusion-weighted signal attenuation, $S(\theta, \phi)$, is the sum of their contributions $S_1(\theta, \phi)$ and $S_2(\theta, \phi)$, weighted by their respective volume fractions. This can be expressed as a convolution over the unit sphere of an axially symmetric *response function* $R(\theta)$, describing the signal attenuation measured for a single fibre population, with a *fibre orientation density function* $F(\theta, \phi)$, describing the fibre orientations present in the voxel (in this case, $F(\theta, \phi) = \delta(\theta_1, \phi_1) + \delta(\theta_2, \phi_2)$).

Results & Discussion: As can be seen from the simulation results in figure 2, the spherical deconvolution method is able to reconstruct the original fibre ODF adequately. The angular resolution of the method is limited by the maximum harmonic order l used, since higher orders correspond to higher angular frequencies. This also limits the minimum angle that can be resolved (for l = 6, this angle is about 40°). However, the maximum harmonic order that can be reliably estimated is limited by the number of independent diffusion-weighting gradient orientations.

As expected, the precision of the reconstructed fibre ODF increases with SNR. The precision also increases as the FA and b-value are increased (results not shown). Using diffusion-weighting gradient schemes with a greater number of directions also increases the precision of the technique, and optionally allows the use of a higher maximum spherical harmonic order. However, the noise sensitivity of the method increases greatly as higher orders are used.

In practice, the response function should be measured in regions of the brain known to have minimal partial volume effects (such as the corpus callosum), in which case the diffusion tensor model would not need to be assumed. However, to demonstrate the efficacy of the method, the diffusion tensor model was used for the simulations in this study. If the assumption that all fibre populations have the same diffusion properties is invalid (in our case, if the FA of the response function $R(\theta)$ does not correspond to that of the fibres present), the simulations show that the directions in the reconstructed ODF are still preserved, while their volume fractions are not (results not shown).

An advantage of this method is that no prior assumptions need be made about the number of fibre populations present in each voxel. The method was also applied to the case of three orthogonal fibre populations, and was able to adequately reconstruct the corresponding fibre ODF (results not shown). In addition, the technique is fast, linear, and uses the set of measured signals directly, thus allowing the straightforward determination of its noise propagation properties.

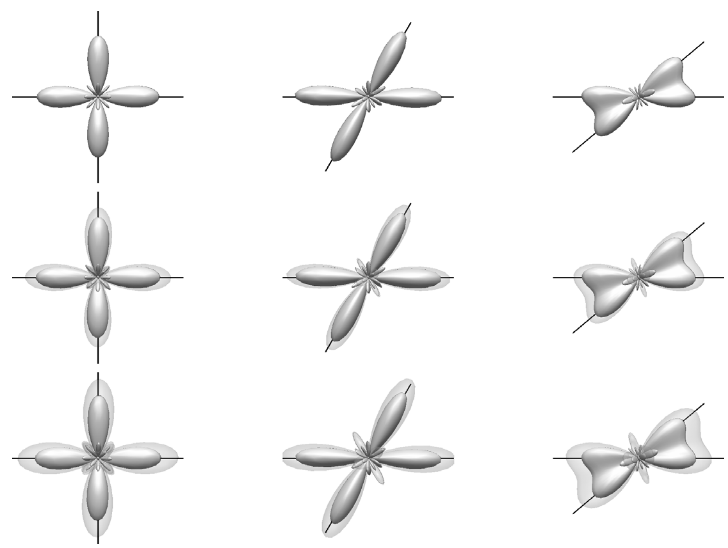

Figure 2: the fibre ODF $F(\theta, \phi)$ estimated using spherical deconvolution, for a system consisting of two fibre populations (FA=0.8) aligned along the solid lines. The mean fibre ODF is depicted by the opaque surface, and the mean+SD by the transparent surface. Top: no noise, middle: $SNR_{b=0}$=50, bottom: $SNR_{b=0}$=30. Left: angle=90º, centre: angle=60º, right: angle=40º. Other parameters were: max l=6, b=3000s/mm2, 60 directions, 100 repetitions.

Conclusion: A new method for *directly* estimating fibre orientations from diffusion-weighted data in partial-volume contaminated regions has been presented. This method can be applied to generate more meaningful measures of fibre orientation coherence (cf. anisotropy), or more robust tractography algorithms.

References: [1] Tuch DS *et al.* MRM 48:577 (2002). [2] Frank LR. MRM 47:1083 (2002). [3] Alexander DC *et al.* MRM 48:331 (2002). [4] Healy DM *et al.* J. Multivariate Anal. 67:1 (1998). [5] Jones DK *et al.* MRM 42:515 (1999).