
An introduction to computational diffusion MRI: the diffusion tensor and beyond

Daniel C Alexander

Department of Computer Science, University College London, Gower Street,
London, WC1E 6BT, UK.
D.Alexander@cs.ucl.ac.uk

Summary. This chapter gives an introduction to the principles of diffusion magnetic resonance imaging (MRI) with emphasis on the computational aspects. It introduces the philosophies underlying the technique and shows how to sensitize MRI measurements to the motion of particles within a sample material. The main body of the chapter is a technical review of diffusion MRI reconstruction algorithms, which determine features of the material microstructure from diffusion MRI measurements. The focus is on techniques developed for biomedical diffusion MRI, but most of the methods discussed are applicable beyond this domain. The review begins by showing how the standard reconstruction algorithms in biomedical diffusion MRI, diffusion-tensor MRI and diffusion spectrum imaging, arise from the principles of the measurement process. The discussion highlights the weaknesses of the standard approaches to motivate the development of a new generation of reconstruction algorithms and reviews the current state-of-the-art. The chapter concludes with a brief discussion of diffusion MRI applications, in particular fibre tracking, followed by a summary and a glimpse into the future of diffusion MRI acquisition and reconstruction.

1 Introduction

Diffusion magnetic resonance imaging (MRI) provides a unique probe into the microstructure of materials. The method observes the displacements of particles that are subject to Brownian motion within a sample material. Specifically, it measures the probability density function p of particle displacements \mathbf{x} over a fixed time t . The microstructure of the material determines the mobility of the particles within and thus determines p . Conversely, features of p provide information about the material microstructure.

In biomedical diffusion MRI, the particles of interest are usually water molecules. Water is a major constituent of biological tissue. Water molecules within tissue undergo random motion due to thermal fluctuations. Currently, brain imaging is the most common application of biomedical diffusion MRI.

The brain has a complex architecture of grey-matter areas connected by white-matter fibres. Diffusion MRI allows non-invasive mapping of the connectivity of the brain.

Figure 1 shows schematic diagrams of four different microstructures that appear in brain tissue together with contours of the p that we expect to observe within each kind of tissue. The diagrams do not aim to reproduce true brain-tissue microstructure, but merely to show how different shapes of p can arise from different configurations of barriers to water mobility. Some regions of the brain, such as the ventricles, contain mostly cerebro-spinal fluid (CSF) and figure 1(a) depicts such a fluid-filled region. Microstructural barriers to water mobility are sparse in these regions, although a few membranes may be present. The function p is isotropic, since displacements are equally likely in all directions. Figure 1(b) depicts grey-matter microstructure. Grey matter is dense tissue containing many barriers to water mobility, such as cell walls and membranes. However, the barriers in grey matter often, as in the picture, have no preferred orientation and so hinder the water movement equally in all directions. The function p thus remains isotropic, but is less spread out than in the CSF region, since the average length of displacements is smaller. Figure 1(c) depicts the microstructure in a white-matter fibre bundle. White matter contains bundles of parallel axon fibres that connect different regions of the brain. The orientations of the cell walls that form barriers to water mobility have much greater consistency in white matter than in grey matter. The microstructure hinders movement more in directions perpendicular to the fibre than along the fibre axis. Displacements along the fibre are larger on average than displacements across it and p is anisotropic with a ridge in the direction of the fibre. More complex microstructure also appears in white matter. Figure 1(d) depicts the microstructure at an orthogonal fibre crossing. Displacements are largest on average in the fibre directions and p has ridges in the directions of each fibre. Other configurations of white matter fibres also occur in the brain.

If we can determine the orientations of the ridges of p , we can infer the dominant orientations of the microstructural fibres. With fibre-orientation estimates in each voxel of a three-dimensional MR image volume, we can follow fibres through the image, using so-called “tractography” algorithms, see Chapter 7 by Villanova et al, and construct a connectivity map of the imaged sample.

The following length scales of brain tissue and the measuring process help appreciation of the discussion in the rest of the chapter:

- The voxel volume in biomedical MRI is of order 10^{-9} m^3 .
- The diffusion time, t , in biomedical diffusion MRI is of order 10^{-2} s and, over this time, the root-mean-squared displacement of water molecules is in the micrometer range.
- The diameters of axon fibres in human white matter can reach $2.50 \times 10^{-5} \text{ m}$, but most axon-fibre diameters are less than 10^{-6} m [1, 2, 3].

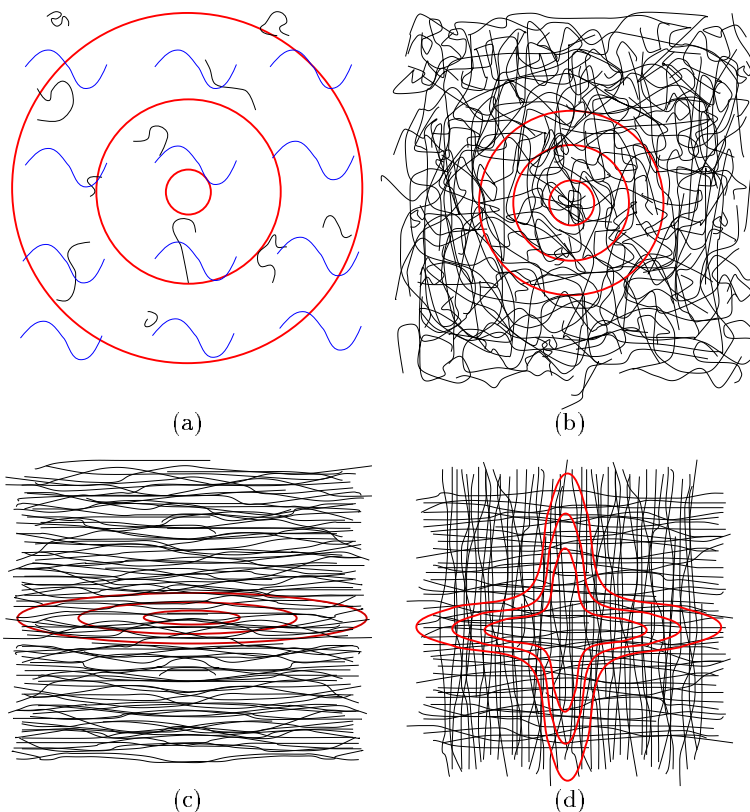


Fig. 1. See colour plates. Schematic diagrams four microstructures found in the brain. The black lines are barriers to the movement of water molecules. The red contours show the expected shape of p in each tissue. Panel (a) shows a fluid-filled region. Panel (b) shows isotropic grey matter. Panels (c) and (d) show white matter with one and two dominant fibre orientations, respectively.

- Coherent white-matter fibre-bundles vary widely in size from several centimetres across down to a few axons.
- The packing density of axon fibres in white matter is of order 10^{11} m^{-2} [1, 2, 3].

The next section introduces the basic diffusion MRI measurement and its relationship to the function p . Section 3 reviews diffusion MRI reconstruction algorithms, which determine features of the microstructure from diffusion MRI measurements. Section 4 gives a brief review of diffusion MRI applications that concentrates on fibre-tracking and connectivity-mapping methods. We conclude in section 5 with a summary of the field and some pointers for future research in diffusion MRI methods.

2 Diffusion-weighted MRI

Diffusion-weighted MRI acquires measurements that are sensitive to the motion of nuclei possessing a net spin (“spins”), most commonly hydrogen nuclei. We can sensitize the MRI measurement to spin displacements by introducing magnetic-gradient pulses to the standard spin-echo sequence (or other standard sequences, such as the stimulated-echo sequence). Figure 2 shows the pulsed-gradient spin-echo (PGSE) sequence [4], which is the most common pulse sequence for diffusion-weighted MRI. The scanner maintains a constant and approximately homogeneous magnetic field \mathbf{H}_0 over the sample. The spins align with \mathbf{H}_0 and have a slightly higher probability of having spin up state than spin down, which causes a non-zero net magnetization of the material. The 90° radio-frequency (RF) pulse P90, centred at time $\tau = 0$, tips the spins into the “transverse” plane perpendicular to \mathbf{H}_0 . The spins then precess about \mathbf{H}_0 at the Larmor frequency, which is proportional to $|\mathbf{H}_0|$. Immediately after P90, the spins precess in phase so that the net magnetization rotates about \mathbf{H}_0 . Inhomogeneities in \mathbf{H}_0 cause the spin precessions to dephase gradually so that the net magnetization decays. The 180° RF pulse P180, centred at time $\tau = \text{TE}/2$ where TE is the “echo time”, negates the phase of each spin. In the absence of the gradient pulses Γ_1 and Γ_2 , the rate of dephasing is the same before and after P180 so the spins come back into phase at time TE. The “spin echo” occurs when the spins come back into phase and recover their net magnetization, which is the MR signal.

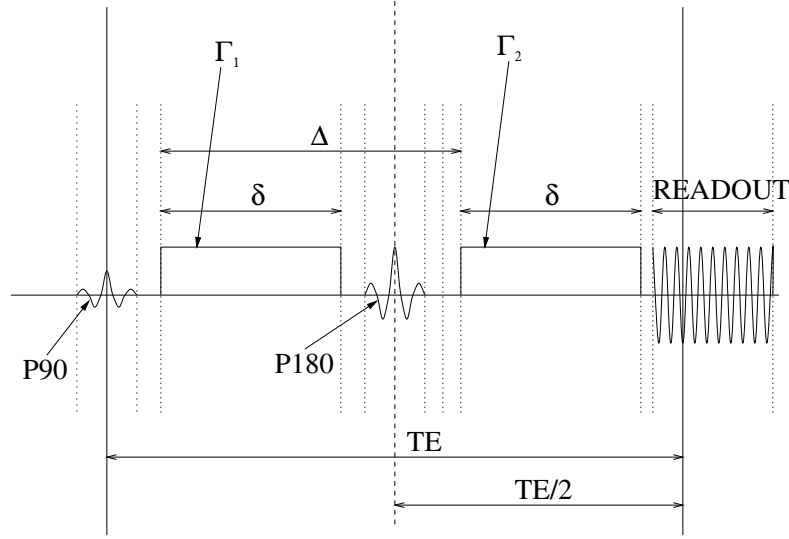


Fig. 2. Shows the pulsed-gradient spin-echo sequence.

The spatially homogeneous diffusion-weighting gradient offsets the phase of each spin by a linear function of the spin position. A gradient pulse $\mathbf{\Gamma}$ offsets the phase of a spin at position \mathbf{r} by $\mathbf{r} \cdot \mathbf{q}$, where

$$\mathbf{q} = \gamma \int_0^\infty \mathbf{\Gamma}(\tau) d\tau,$$

$\mathbf{\Gamma}(\tau)$ is the component of the magnetic-field gradient parallel to \mathbf{H}_0 (i.e. $(\nabla \mathbf{H}_0) \mathbf{H}_0$) at time τ and γ is the gyromagnetic ratio of the spins. The gyromagnetic ratio for protons in water, which are usually the spins in biomedical diffusion MRI, is $2.675 \times 10^8 \text{ s}^{-1} \text{ T}^{-1}$. In practice, the pulses are usually approximately rectangular (with brief rise and fall times) so that $\mathbf{\Gamma}_1$ and $\mathbf{\Gamma}_2$ have constant value \mathbf{g} over the pulse duration δ and $\mathbf{q} = \gamma \delta \mathbf{g}$. Since P180 negates the phase of each spin, $\mathbf{\Gamma}_2$ cancels the phase offset from $\mathbf{\Gamma}_1$ for a stationary spin. However, if a spin moves from position \mathbf{r}_1 to \mathbf{r}_2 between the two pulses, it retains a residual phase offset of $\mathbf{q} \cdot (\mathbf{r}_2 - \mathbf{r}_1) = \mathbf{q} \cdot \mathbf{x}$, where \mathbf{x} is the spin displacement. The magnetic moment of the spin at the spin echo is thus $\mathcal{M} \exp(i\mathbf{q} \cdot \mathbf{x})$, where \mathcal{M} is the magnitude of the magnetic moment. The MRI signal $A^*(\mathbf{q})$ is the magnetization of all contributing spins. If we sum over all possible spin displacements, we see that

$$A^*(\mathbf{q}) = A^*(\mathbf{0}) \int_{\mathbb{R}^3} p(\mathbf{x}) \exp(i\mathbf{q} \cdot \mathbf{x}) d\mathbf{x},$$

where $A^*(\mathbf{0})$ is the signal with no diffusion-weighting gradients. Diffusion MRI usually assumes that the local advection velocity is zero (no net motion of the spin population), so that $p(\mathbf{x}) = p(-\mathbf{x})$ and $A^*(\mathbf{q})$ is real valued in the absence of noise. Moreover, we use the normalized signal

$$A(\mathbf{q}) = (A^*(\mathbf{0}))^{-1} A^*(\mathbf{q}) = \int_{\mathbb{R}^3} p(\mathbf{x}) \cos(\mathbf{q} \cdot \mathbf{x}) d\mathbf{x}, \quad (1)$$

which is the Fourier transform of the function p at wavenumber \mathbf{q} . A measurement $A(\mathbf{q})$ thus provides the *apparent diffusion coefficient* (ADC) $d = -b^{-1} \log(A(\mathbf{q}))$, where $b = t|\mathbf{q}|^2$ is the *diffusion-weighting factor*, on the assumption that p is an isotropic zero-mean Gaussian function [4, 5]. Researchers in the 1980s [6, 7] combined the basic diffusion-weighted NMR measurement described above with MRI to obtain image maps of the ADC.

The derivation of equation (1) assumes that the movement of particles during the gradient pulses is negligible. This assumption is justified if δ is small compared with the pulse separation Δ , which is then the diffusion time $t = \Delta$. In practice however, δ and Δ usually have similar magnitude, as in figure 2. When δ is non-negligible, the phase offset of a spin depends on its trajectory during $\mathbf{\Gamma}_1$ and $\mathbf{\Gamma}_2$ rather than just its displacement, which complicates the model relating the measurements to p ; see discussions in [8, 9, 5]. With some assumptions, we can model the effects of non-negligible δ analytically. For

example, if p is Gaussian and the gradient pulses are rectangular, then non-negligible δ reduces t to an effective diffusion time of $\Delta - \delta/3$, see [4, 5]. The effective diffusion time reduces still further for higher moments of p [10]. Mitra and Halperin [9] show that, if p is the displacement density of the centres of mass (COM) of particle trajectories over time δ , rather than the displacements of particles themselves, equation (1) holds with $t = \Delta$ even with non-negligible δ . The COM displacement density has similar shape to the particle displacement density (though somewhat blurred) and, in particular, indicates fibre directions in the same way. Thus, although non-negligible δ confounds absolute measurements of the particle displacement density, the features of p of interest in brain imaging are relatively unaffected. Lori et al [8] and Brihuega-Moreno et al [11] provide some further analysis of the non-negligible Δ problem.

The MRI measurement is complex-valued, since the magnetization has magnitude and phase. Often the phase of the measurements is discarded, since inhomogeneities in \mathbf{H}_0 and movement of the sample make it unstable. In practice, it is common to take the modulus of $A^*(\mathbf{q})$ as the real-valued MR signal. An additive Gaussian noise model is common in MRI. With this model, the real and imaginary parts of the signal are independent and identically distributed with distribution $N(0, \sigma^2)$. Noise on the modulus of the signal thus follows a Rician distribution [12], which tends to a Gaussian distribution as the signal-to-noise ratio increases. A common measurement of quality of diffusion MRI data sets is the signal-to-noise ratio $S = \mathcal{A}^*(\mathbf{0})/\sigma$ of the measurement with $\mathbf{q} = \mathbf{0}$, where \mathcal{A}^* is the noise-free signal.

3 Diffusion MRI reconstruction algorithms

This section reviews diffusion MRI reconstruction algorithms. We focus here on reconstructing fibre orientations, but note that some diffusion MRI techniques aim to estimate other features of the microstructure, such as the ratio of intracellular to extracellular water [13], by targeting other features of p . For this discussion, a diffusion MRI reconstruction algorithm inputs a set of diffusion-weighted MRI measurements from one voxel and outputs, at least, i) the number n of dominant fibre directions and ii) an estimate of each dominant fibre direction. Most of these algorithms determine a feature of p that highlights fibre orientations. In addition to fibre-orientation estimates, these features of p usually provide scalar indices of shape that discriminate different kinds of material and can indicate the reliability of the fibre-orientation estimates.

3.1 Diffusion Tensor MRI

Diffusion-tensor (DT) MRI [14] computes the *apparent diffusion tensor* on the assumption that p is a zero-mean trivariate Gaussian distribution:

$$p(\mathbf{x}) = G(\mathbf{x}; \mathbf{D}, t), \quad (2)$$

where

$$G(\mathbf{x}; \mathbf{D}, t) = ((4\pi t)^3 \det(\mathbf{D}))^{-\frac{1}{2}} \exp\left(-\frac{\mathbf{x}^T \mathbf{D}^{-1} \mathbf{x}}{4t}\right),$$

\mathbf{D} is the diffusion tensor and t is the diffusion time. Since the Gaussian function has a single ridge, DT-MRI assumes $n = 1$. Substitution of (2) into (1) gives

$$A(\mathbf{q}) = \exp(-t \mathbf{q}^T \mathbf{D} \mathbf{q}). \quad (3)$$

If we take the logarithm of (3), we see that each $A(\mathbf{q})$ provides a linear constraint on the elements of \mathbf{D} . The Gaussian model has six free parameters, which are the elements of the symmetric three-by-three matrix \mathbf{D} . To fit the six free parameters, we need a minimum of six $A(\mathbf{q})$ with independent \mathbf{q} , although many more are often acquired. Note that six $A(\mathbf{q})$ requires a minimum of seven $A^*(\mathbf{q})$ including one for normalization. Practitioners most often use the linear least-squares fit of \mathbf{D} to the log measurements. However, fitting directly using equation (3), as in [15], can improve results, since the error distribution is closer to normal on $A(\mathbf{q})$ than on $\log(A(\mathbf{q}))$. When fitting directly to $A(\mathbf{q})$, we can include constraints on the diffusion tensor, such as positive definiteness, using the Cholesky decomposition as in [15], or cylindrical symmetry, by writing

$$\mathbf{D} = \alpha \mathbf{n} \mathbf{n}^T + \beta \mathbf{I}, \quad (4)$$

where \mathbf{n} is the principal direction of the diffusion tensor, \mathbf{I} is the identity tensor and \mathbf{D} has eigenvalues $\alpha + \beta$, β and β .

Diffusion-tensor MRI generalizes the ADC calculation from simple diffusion-weighted MRI to three dimensions. It provides two extra insights into the material microstructure over simple diffusion-weighted MRI. First, it provides rotationally invariant statistics of the anisotropy of p , which reflect the anisotropy of the microstructure. Second, it provides an estimate of the dominant orientation of microstructural fibres. The eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ of \mathbf{D} determine the shape of p . The Gaussian function has ellipsoidal contours and the relative lengths of the major axes of the ellipsoids have the same proportions as the $(\lambda_i)^{\frac{1}{2}}$. Statistics of anisotropy come from the distribution of eigenvalues. A common statistic is the *fractional anisotropy* [16]

$$\nu = \left(\frac{3}{2} \sum_{i=1}^3 \left(\lambda_i - \frac{1}{3} \text{Tr}(\mathbf{D}) \right)^2 \right)^{\frac{1}{2}} \left(\sum_{i=1}^3 \lambda_i^2 \right)^{-\frac{1}{2}}, \quad (5)$$

which is the normalized standard deviation of the eigenvalues. Figure 3(a) shows ν over a coronal slice through a healthy human brain. The highest values of ν are in regions of dense white matter, such as the corpus callosum, where the fibres are packed most densely and have consistent orientation. Other common scalar statistics derived from the diffusion tensor are $\text{Tr}(\mathbf{D})$

and the skewness μ . The trace of the diffusion tensor $\text{Tr}(\mathbf{D}) = \sum_{i=1}^3 \lambda_i$ is proportional to the mean squared displacement of water molecules and thus indicates the mobility of water molecules within each voxel, which reflects the density of microstructural barriers. The skewness

$$\mu = \left(\frac{9}{2} \sum_{i=1}^3 \left(\lambda_i - \frac{1}{3} \text{Tr}(\mathbf{D}) \right)^3 \right)^{\frac{1}{3}} \left(\sum_{i=1}^3 \lambda_i^3 \right)^{-\frac{1}{3}}, \quad (6)$$

is close to zero for isotropic diffusion tensors with near spherical contours ($\lambda_1 \approx \lambda_2 \approx \lambda_3$), positive for prolate diffusion tensors with cigar-shaped contours ($\lambda_1 \gg \lambda_2 \approx \lambda_3$) and negative for oblate diffusion tensors with pancake-shaped contours ($\lambda_1 \approx \lambda_2 \gg \lambda_3$). Figures 3(b) and 3(c) show $\text{Tr}(\mathbf{D})$ and μ , respectively, over the coronal slice in figure 3(a). The highest values of $\text{Tr}(\mathbf{D})$ are in the ventricles and other regions of cerebro-spinal fluid, where the density of barriers to water mobility is low. The skewness is positive in most white-matter regions. In some white-matter regions, such as the pons, the skewness is negative showing oblate diffusion tensors. Oblate diffusion tensors arise in regions contain orthogonally crossing fibres, as depicted in figure 1(d), where the best-fit Gaussian model has oblate contours. Many other configurations of fibres or microstructures can also give rise to oblate diffusion tensors.

The eigenvectors \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 of \mathbf{D} determine the orientation of p . In regions of prolate diffusion tensors, the principal eigenvector \mathbf{e}_1 (that with eigenvalue λ_1) provides an estimate of the single fibre orientation. At fibre crossings where the diffusion tensor is oblate, \mathbf{e}_1 and \mathbf{e}_2 span the plane of the crossing fibres. Pajevic and Pierpaoli [17] use colour for a compact visualization of fibre orientations. A popular choice is to use RGB vectors proportional to $\nu^{1/2} \mathbf{e}_1$. Figure 3(d) uses this colour orientation-encoding for the coronal slice in figure 3(a); red indicates left-right orientation, green indicates anterior-posterior (front to back of the head) and blue indicates inferior-superior (top to bottom of the head).

Diffusion-tensor MRI requires a minimum of seven MRI measurements. Most diffusion-tensor MRI sequences acquire more than the minimum seven measurements to reduce the effects of noise. The standard approach [18] is to acquire M measurements with $\mathbf{q} = \mathbf{0}$ and N measurements with non-zero wavenumbers \mathbf{q}_i , $i = 1, \dots, N$. The $|\mathbf{q}_i|$ are all equal and the diffusion time, t , and hence b , is fixed for all the $A(\mathbf{q}_i)$. The directions $\hat{\mathbf{q}}_i$ are unique and distributed uniformly over the sphere. This kind of scheme gives less rotational dependence of the fibre-orientation estimates and shape statistics than schemes that acquire repeated measurements at a smaller number of \mathbf{q}_i [19]. The images in figure 3 come from a data set with $M = 6$, $N = 54$, $\delta = 0.034$ s, $\Delta = 0.040$ s and $|\mathbf{g}| = 0.022$ Tm $^{-1}$. Thus $|\mathbf{q}_i| = 2.0 \times 10^5$ m $^{-1}$ and $b = 1.15 \times 10^9$ s m 2 using $t = \Delta - \delta/3$. This scheme is typical for whole-brain clinical DT-MRI and requires around 20 minutes scan time on standard hardware. In white-matter regions, the signal to noise ratio at $\mathbf{q} = 0$, S , is around 16 on average.

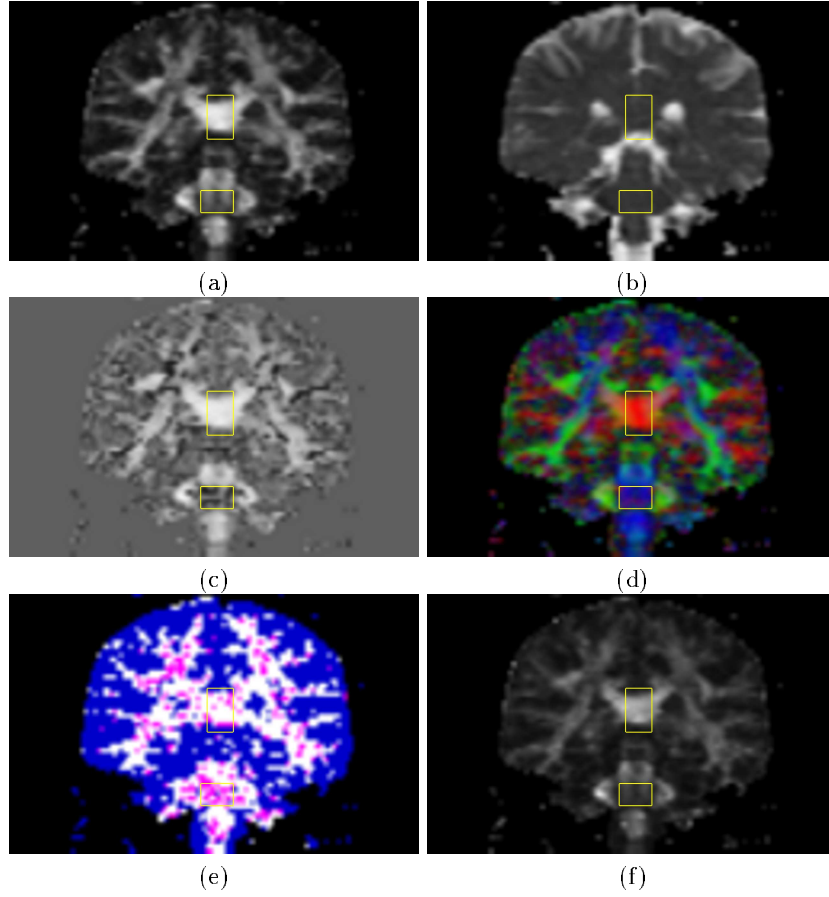


Fig. 3. See colour plates. Shows various features of p plotted over a coronal slice through a healthy human brain. Panel (a) shows the fractional anisotropy, ν . Panel (b) shows $\text{Tr}(\mathbf{D})$. Panel (c) shows the skewness, μ . Panel (d) shows the colour coded principal direction, \mathbf{e}_1 . Panel (e) shows the output of Alexander's voxel classification algorithm (section 3.2); black is background, blue is order 0, white is order 2 and pink is order 4. Panel (f) shows the spherical-harmonic anisotropy (section 3.2). In each panel, the upper region of interest contains some grey matter (top), part of the corpus callosum (middle) and some CSF (bottom). The lower region contains the fibre crossing in the pons.

Diffusion-tensor MRI is the most popular diffusion MRI reconstruction algorithm by far. It is the simplest technique that provides anisotropy statistics and fibre-orientation estimates. The computational and data requirements of the technique are modest. Modern scanners come with built-in acquisition sequences for DT-MRI and the post-processing is simple to implement and fast to run on modern desktop computers. However, a major drawback of DT-MRI is that the Gaussian model is often a poor fit to the data. Diffusion-tensor MRI provides only one fibre-orientation estimate in each voxel. In regions where fibres cross within one voxel, p has multiple ridges, as figure 1(d) depicts. The Gaussian model of p in figure 1(d) has oblate ellipsoidal contours. When the Gaussian model is poor the two major selling-points of DT-MRI fail. First, indices of anisotropy derived from the diffusion tensor, such as ν , underestimate the true directional variability of p . The Gaussian model for p in figure 1(d) smoothes out the ridges in the plane of the page. Second, fibre-orientation estimates are incorrect. For a perfectly oblate Gaussian distribution, \mathbf{D} has no unique principal eigenvector. In regions with the microstructure depicted in figure 1(d), measurement noise will ensure that \mathbf{e}_1 is randomly oriented in the plane of the crossing fibres. Non-orthogonally intersecting fibres are potentially more dangerous. The single fibre-orientation estimate from DT-MRI then lies consistently between the two true fibre directions; see [20] for an example.

3.2 Modelling the “ADC profile”

Equation (3) shows that, with no noise, $\log(A)$ is quadratic in \mathbf{q} when p is Gaussian. Several authors model $\log(A)$ with higher-order polynomials both to detect departures from the Gaussian model and to obtain more reliable indices of anisotropy.

Frank [21] and Alexander et al [22] both fit the spherical-harmonic series to $\log(A)$ at a fixed $|\mathbf{q}|$. In the literature, the term “ADC profile” refers to $-b^{-1}\log(A)$ as a function of $\hat{\mathbf{x}}$ with fixed $|\mathbf{q}|$. The spherical harmonics Y_{lm} , $l = 0, \dots, \infty$, $m = -l, \dots, l$, form a basis for complex-valued functions on the unit sphere in three dimensions \mathcal{S}^2 . Thus we can write any complex-valued function f of the sphere as

$$f = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{lm} Y_{lm}.$$

Each spherical-harmonic series containing only terms up to order $l = L$ is the restriction to \mathcal{S}^2 of an order L polynomial, and vice versa. Series with only even-order terms are symmetric, so that $f(\hat{\mathbf{x}}) = f(-\hat{\mathbf{x}})$, and constraining both $\text{Im}(a_{l0}) = 0$ and $a_{lm} = (-1)^m a_{l(-m)}^*$ for all l and m ensures that f is real-valued [22]. Reference [22] shows how to compute the least-squares-fit symmetric real-valued spherical-harmonic series to $\log(A(\mathbf{q}_i))$, $i = 1, \dots, N$, robustly via a single matrix multiplication.

If $\log(A)$ is quadratic, its spherical-harmonic series contains only terms up to order 2. If the fitted spherical-harmonic series contains significant higher-order terms, the Gaussian model for p is poor. Frank [21] observes significant fourth-order terms in the spherical-harmonic series in various white-matter regions in the human brain. Alexander [22] uses the analysis of variance (ANOVA) test for deletion of variables, the “ F -test” [23], to choose the lowest-order series that fits the data. This simple voxel-classification algorithm classifies each voxel as isotropic (order 0), anisotropic Gaussian (order 2), or non-Gaussian (order 4 or above). Results show clusters of order 4 voxels in several fibre-crossing regions in human-brain data similar to that used for figure 3. Figure 3(e) shows Alexander’s voxel classification over the coronal slice in figure 3(a). Order 4 voxels appear consistently in the pons and other fibre crossings showing failure of the Gaussian model.

Spherical-harmonic models of $\log(A)$ provide anisotropy indices that are robust to departures from the Gaussian model. The moments of a spherical function f are

$$\omega_n[f] = (4\pi)^{n/2} \int f^n(\hat{\mathbf{x}}) d\hat{\mathbf{x}}.$$

A general index of anisotropy of f is $(\omega_1[f])^{-1}(\omega_2[f] - (\omega_1[f])^2)^{1/2}$. For a real-valued symmetric spherical-harmonic series, $\omega_1[f] = 4\pi a_{00}$ and $\omega_2[f] = 4\pi \sum_{l=0}^{\infty} \sum_{m=-2l}^{2l} |a_{(2l)m}|^2$. Figure 3(f) shows the spherical-harmonic anisotropy from series including terms up to order 4. Differences between the spherical-harmonic anisotropy and ν are more noticeable at higher b . Other moments may also provide useful shape indices.

Ozarslan et al [24] use a higher-order tensor model of $\log(A)$ so that

$$\log(A(\mathbf{q})) = -t \mathbf{q}^{(j)} \mathbf{D}^{(2j)} \mathbf{q}^j, \quad (7)$$

where the term on the right contains the contraction of the order $2j$ tensor $\mathbf{D}^{(2j)}$ by $\mathbf{q}^{(j)}$, which is the outer product of $\mathbf{q}^{(1)} = \mathbf{q}$ and $\mathbf{q}^{(j-1)}$. The tensors $\mathbf{D}^{(2j)}$ are real valued and have symmetry ensuring that $\log(A(\mathbf{q})) = \log(A(-\mathbf{q}))$. It is straightforward to demonstrate that the unique elements of the tensor model in (7) with order $2j$ are a linear transformation of the real and imaginary parts of the coefficients of the real symmetric spherical-harmonic model including terms up to order $2j$. In this sense, the two methods equivalent, both theoretically and computationally. Liu et al [10] model $\log(A(\mathbf{q}))$ by a sequence of higher-order tensors, which includes both odd and even-order tensors. The inclusion of odd-order tensors allows the model to capture non-symmetric spin displacements.

Neither the spherical-harmonic nor the higher-order-tensor models provide fibre-orientation estimates. Both model $\log(A)$ rather than p and the peaks of $\log(A)$ at a fixed radius are not in the directions of the ridges of p in general. The scalar anisotropy of $\log(A)$ correlates with that of p , so we can compute anisotropy indices from $\log(A)$. Also when p is Gaussian, A is Gaussian, so we can infer departures from the Gaussian model of p from departures of the

$A(\mathbf{q}_i)$ from the best-fit Gaussian. To estimate fibre orientations, however, we must invert the Fourier transform in (1) and reconstruct directional features of p .

3.3 Multi-compartment models

A simple generalization of DT-MRI replaces the Gaussian model for p with a mixture of Gaussian densities:

$$p(\mathbf{x}) = \sum_{i=1}^n a_i G(\mathbf{x}; \mathbf{D}_i, t), \quad (8)$$

where each $a_i \in [0, 1]$ and $\sum_i a_i = 1$. Particle displacements in media containing n distinct compartments, between which no exchange of particles occurs, follow the distribution in equation (8) if the displacement density in the i -th compartment, which has volume fraction a_i , is $G(\mathbf{x}; \mathbf{D}_i, t)$.

We take the Fourier transform of (8) and substitute into (1) to relate the measurement values to the model parameters (\mathbf{D}_i and a_i , $i = 1, \dots, n$):

$$A(\mathbf{q}) = \sum_{i=1}^n a_i \exp(-t \mathbf{q}^T \mathbf{D}_i \mathbf{q}).$$

The constraint on the model parameters from each measurement is non-linear so we must fit the model to the data by non-linear optimization using, for example, a Levenberg–Marquardt algorithm [25]. The principal eigenvector of each \mathbf{D}_i provides a fibre-orientation estimate. The multi-compartment model assumes n is fixed. Practical considerations, such as the number of MRI measurements and the measurement noise level, limit the number of orientations the method can resolve reliably. Most work to date uses a maximum n of 2.

Two problems accompany the use of multi-compartment models. First, the choice of n presents a model-selection problem. Second, the non-linear fitting procedure is unstable and starting-point dependent, because of local minima in the objective function. Parker and Alexander [26] and Blyth et al [27] use Alexander’s voxel classification algorithm [22] to solve the model-selection problem. They use $n = 2$ in order 4 voxels, where DT-MRI fails, and $n = 1$ elsewhere. This method does not extend naturally above $n = 2$, however. Although a fourth-order polynomial is a good approximation to $\log(A)$ from a mixture of two Gaussian densities [21], a mixture of three Gaussians does not necessarily require a sixth-order polynomial. Tuch [28] thresholds the correlation of the measurements with their predictions from a single-component model to decide whether to use one or two components. Constraints on the diffusion tensors in the multi-compartment model can help stabilize the fitting procedure. For example, we can enforce positive definiteness on the \mathbf{D}_i , using the Cholesky decomposition [29], or cylindrical symmetry on \mathbf{D}_i using equation (4), or specific eigenvalues as in [28]. Spatial regularization techniques also help overcome the fitting problem by ensuring voxel to voxel coherence, see [29] and Chapter 9 by Pasternak et al.

3.4 Fibre models

A similar model-based approach [30] assumes that particles belong to one of two populations: a restricted population within or around microstructural fibres and a free population that are unaffected by microstructural barriers. With negligible exchange between the populations, $p = ap_f + (1 - a)p_r$, where p_f is the spin-displacement density for the free population, p_r that for the restricted population, and a is the fraction of particles in the free population. Behrens et al [30] use an isotropic Gaussian model for p_f . They use a Gaussian model for p_r in which the diffusion tensor has only one non-zero eigenvalue so that particle displacement is restricted to a line. Assaf et al [31] describe a similar approach. They model p_r with Neuman’s model for restricted diffusion in a cylinder [32]. The fitted p_r provides the fibre-orientation estimate. For p_f , which they call the “hindered compartment”, they use an anisotropic Gaussian model.

Both approaches extend naturally to the multiple-fibre case by including multiple restricted populations in the model, which gives a more physically-based mixture model than the multi-compartment models in section 3.3. In the multiple-fibre case, fibre-model approaches have the same model-selection and fitting problems as multi-compartment models. Assaf et al [31] show promising results in the two-fibre case in simulation.

3.5 Diffusion spectrum imaging

Diffusion spectrum imaging [33], unlike the approaches discussed earlier in this section, does not use a parametric model for p . Instead, DSI reconstructs a discrete representation of p directly from measurements on a regular grid of wavenumbers via a fast Fourier transform. The reconstruction gives values of p on a grid of displacements.

The orientation distribution function (ODF)

$$\phi(\hat{\mathbf{x}}) = \int_0^\infty p(\alpha\hat{\mathbf{x}})d\alpha, \quad (9)$$

where $\hat{\mathbf{x}}$ is a unit vector in the direction of \mathbf{x} , is the radial projection of p onto the unit sphere. The ODF has peaks in the directions in which p has most mass and thus has peaks in the directions of the ridges of p . In DSI, therefore, the peaks of ϕ provide the fibre-orientation estimates. The function ϕ can have multiple pairs of equal and opposite peaks. Each pair provides a separate fibre-orientation estimate, which enables DSI to resolve the orientations of fibres that cross within a single voxel. The ODF also provides anisotropy indices. For example, we can use the standard deviation $(\omega_2[\phi] - 4\pi)^{1/2}$ of ϕ as an analogue of the fractional anisotropy, ν .

Qualitative results from DSI in [33, 34] and subsequent publications show ODF peaks in the expected fibre directions at known crossings in human and animal brain data. However, the results also show ODFs with multiple peaks

in grey-matter regions and it is unclear whether these peaks show genuine anatomic structure or simply arise from measurement noise. Diffusion spectrum imaging has clear advantages over DT-MRI and multi-compartment modelling, since it can resolve multiple fibre orientations, it does not require non-linear fitting and it does not involve a model-selection problem. Despite its advantages, DSI is not used as widely as DT-MRI. The main drawback of the technique is that acquisition times are long, since it requires an order of magnitude more measurements than DT-MRI to get sufficient detail in the reconstructed p . Wedeen and Tuch and coworkers [33, 34] use around 500 measurements for DSI. They acquire images with $64 \times 10^{-9} \text{ m}^3$ voxels, compared with $10 \times 10^{-9} \text{ m}^3$ voxels typical in DT-MRI, to keep the acquisition time manageable. Furthermore, DSI ignores the effects of non-negligible δ discussed at the end of section 2.

3.6 A new generation of multiple-fibre reconstructions

The main drawback of DSI is the long acquisition time. However, in many applications, DSI wastes much of the information in the measurements. The projection of p onto the sphere to obtain the ODF, ϕ , discards the radial component of p to which much of the information in the measurements contributes. In some applications, the radial component of p may be useful. However, the primary interest is often in the angular structure, which provides fibre-orientation estimates and anisotropy indices. An emerging new generation of diffusion MRI reconstruction algorithm reconstructs the angular structure of p directly from the measurements. Rather than acquiring measurements on a grid of wavenumbers, as in DSI, the new methods use sets of wavenumbers chosen to contribute mostly to the angular structure of p . Specifically, methods to date use the spherical acquisition schemes popular in DT-MRI (see section 3.1).

Approximations to the ODF

Several methods approximate the ODF from measurements acquired using a spherical acquisition scheme. Tuch’s \mathbf{q} -ball imaging method [34, 35] approximates the ODF by the Funk transform [36] of the measurements. (For brevity in the remainder of the chapter, we shall refer to Tuch’s method simply as “ \mathbf{q} -ball”.) The Funk transform is a mapping between functions of the sphere. The value of the Funk transform of a function f at a point $\hat{\mathbf{x}}$ is the integral of f over the great circle perpendicular to $\hat{\mathbf{x}}$. In [34], Tuch shows that, in the absence of noise, the approximation becomes closer as $|\mathbf{q}|$ increases. Qualitative results in [34, 35] show good agreement between \mathbf{q} -ball and full DSI in a fibre-crossing region in the human brain. Tuch uses high-quality test data with $N = 492$ and with $|\mathbf{q}| = 3.6 \times 10^5 \text{ m}^{-1}$ ($b = 4.0 \times 10^9 \text{ s m}^{-2}$) and $|\mathbf{q}| = 5.4 \times 10^5 \text{ m}^{-1}$ ($b = 12.0 \times 10^9 \text{ s m}^{-2}$). Lin et al [37] propose a similar algorithm independently. They test their algorithm on data acquired from a

phantom containing water-filled capillaries in two orientations, which simulates crossing white-matter fibres. The algorithm recovers the orientation of the capillaries consistently.

In Chapter 10 and [38], Ozarslan et al fit higher-order tensor models (see section 3.2) to measurements from a spherical acquisition scheme. They assume that $A(\mathbf{q})$ decays exponentially with increasing $|\mathbf{q}|$ and fixed $\hat{\mathbf{q}}$. This assumption allows them to estimate the measurements on a regular grid of wavenumbers, which they use as input to DSI. The method finds the ridge directions of simple test functions and qualitative results on rat-brain data, with $N = 81$ and $b = 1.5 \times 10^9 \text{ s m}^2$, are promising.

Deconvolution techniques

Deconvolution methods generalize the fibre-model methods by assuming a distribution of fibre orientations. The diffusion MRI signal is the convolution of the *fibre orientation distribution* (FOD) with the response from a single fibre [30, 39, 40]. Any fibre model can provide the response for a single fibre. References [30, 39] use Gaussian fibre models. Tournier [40] derives a fibre model directly from the data by taking an average signal from the most anisotropic voxels. Deconvolution is linear using a linear set of basis functions, such as the spherical harmonics, for the FOD [39, 40]. The peaks of the FOD provide fibre-orientation estimates. Like the ODF, the FOD can have any number of pairs of equal and opposite peaks and each pair provides a separate fibre-orientation estimate. Thus, deconvolution methods avoid the model-selection problems associated with multi-compartment and fibre models.

Other methods

Jansons and Alexander's PASMRI algorithm [41] computes another feature of p called the persistent angular structure (PAS). The PAS is the function \tilde{p} of the sphere that, when embedded in three-dimensional space on a sphere of radius r , has Fourier transform that best fits the measurements. Thus

$$\tilde{p} = \arg \min_{\tilde{p}} \left[\sum_{i=1}^N \left(A(\mathbf{q}_i) - \hat{A}(\mathbf{q}_i; \tilde{p}) \right)^2 \right],$$

where

$$\hat{A}(\mathbf{q}_i; \tilde{p}) = \int \tilde{p}(\hat{\mathbf{x}}) \cos(r\mathbf{q}_i \cdot \hat{\mathbf{x}}) d\hat{\mathbf{x}}. \quad (10)$$

Jansons and Alexander use a maximum-entropy parametrization of \tilde{p} . They fit the $N + 1$ parameters of \tilde{p} using a Levenberg–Marquardt algorithm and numerical approximations of the integrals in (10). The function \tilde{p} can have any number of pairs of equal and opposite peaks and each pair provides a fibre-orientation estimate. The parameter r controls the smoothness of \tilde{p} .

The iterative optimization required to compute \tilde{p} makes PASMRI a much slower algorithm than the other algorithms discussed in this section. However, Alexander [42] shows that PASMRI reconstructs fibre directions more consistently than **q**-ball. On the human-brain data used for figure 3, the **q**-ball algorithm fails to resolve the orientations at known fibre-crossings, where PASMRI succeeds. Simulations show that PASMRI is more sensitive than **q**-ball to anisotropy in test functions and recovers ridge directions more reliably, particularly at low b and S . Figures 4 and 5 show the PAS, computed by PASMRI, and ODF, approximated by **q**-ball, respectively, in the coronal brain slice in figure 3. The PAS has sharper peaks than the ODF and resolves the crossing fibres in the pons more consistently.

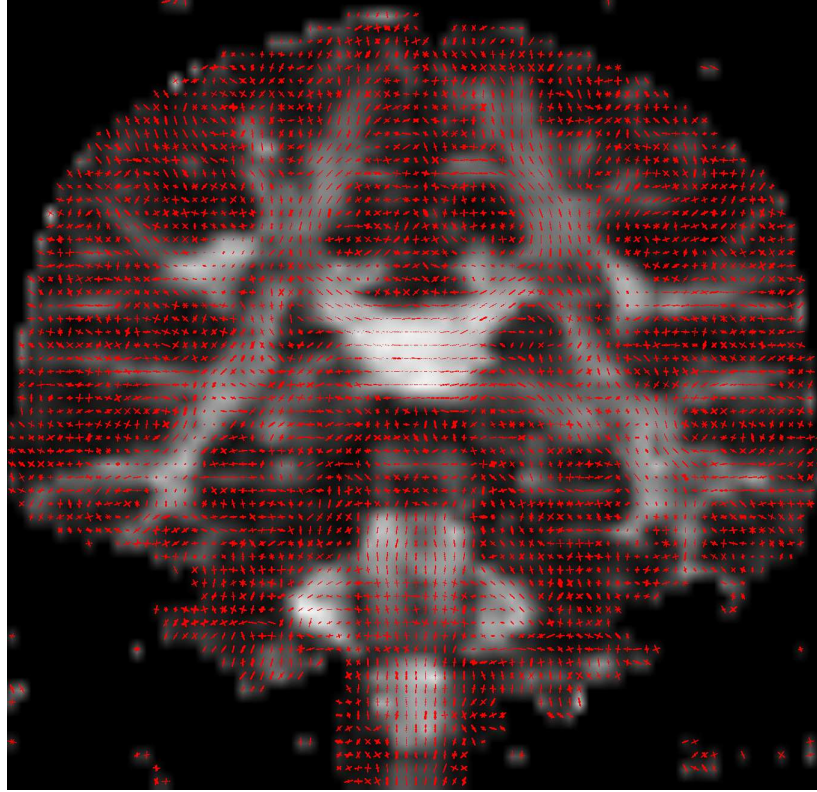


Fig. 4. See colour plates. Shows the PAS (in red) in brain voxels of the coronal slice in figure 3 superimposed on the fractional anisotropy map.

Liu et al [10] outline a general inversion of their higher-order-tensor series model (see section 3.2) of $\log(A)$ to obtain p . They simulate random walks of molecules through restricted media to obtain synthetic MRI measurements. In

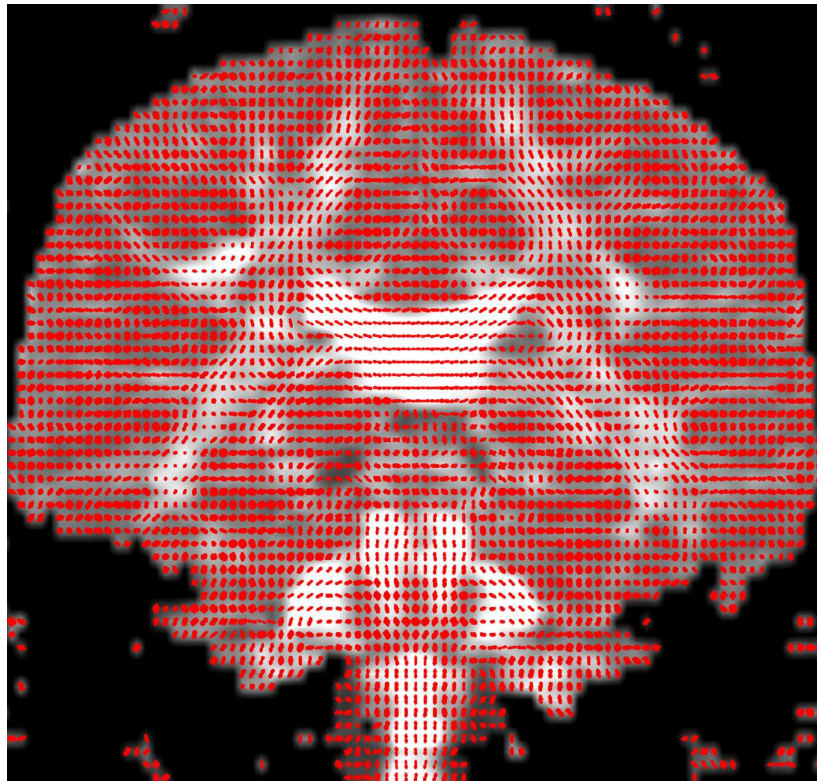


Fig. 5. See colour plates. Shows the ODF (in red) approximated using q -ball in brain voxels of the coronal slice in figure 3 superimposed on the fractional anisotropy map.

simulation, the reconstructed p reflects the geometry of several simple media. Liu et al use the phase of the MRI measurements and include odd-order tensors in their model, which allows them to determine net motion of particles (advection) as well as symmetric motion.

4 Applications

Diffusion-tensor MRI is now a routine clinical technique. Scalar statistics derived from the apparent diffusion tensor, such as $\text{Tr}(\mathbf{D})$ and the fractional anisotropy [16], are used to study a broad range of conditions including stroke, epilepsy, multiple sclerosis, dementia and many other white-matter diseases; see [43] for a recent review. Diffusion MRI is also used to probe the microstructure of a variety of other materials including muscle tissue, e.g. in the heart [44], cartilage [45], plant tissue [46] and porous rock [47].

A major application area of diffusion MRI is fibre-tracking or “tractography”. References [48, 49] contain reviews of tractography techniques with qualitative and quantitative performance comparisons. Chapter 7 by Villanova et al also discusses tractography techniques. Simple “streamline” tractography algorithms trace fibre trajectories by following fibre-orientation estimates from point to point through an image volume. Probabilistic tractography algorithms use a probability density function to model the uncertainty in the fibre-orientation estimate in each voxel. The algorithms run repeated streamline tractography processes with fibre-orientation estimates drawn from the model in each voxel. The fraction of streamlines that pass through a voxel provides an index of the connectivity of that voxel to the starting point. The distributions on the fibre-orientation estimates generally come from modelling the distribution of estimates from repeated trials of adding synthetic noise to the measurements. Some implementations [20, 50] use shape indices, such as the fractional anisotropy, to predict the parameters of the distribution.

Most tractography in the literature uses DT-MRI for fibre-orientation estimates. Several authors [20, 26, 27, 34, 51] use multiple-fibre reconstructions in tractography applications. Tractography algorithms based on multiple-fibre reconstruction use fibre-orientation estimates from multi-compartment models [20, 26, 27] or the peaks of the ODF [34] or PAS [51] together with uncertainty models obtained from simulations. Blyth et al [27] provide direct evidence that multiple-fibre reconstruction improves tractography results over DT-MRI. One might be tempted to use the PAS, the ODF or the FOD as a direct estimate of the distribution of fibre orientations for probabilistic tractography. However, the physical basis of these functions is a gross simplification of the complex distribution of barriers to diffusion within material such as brain tissue. Any supposed relationship between these features of p and the true distribution of fibre orientations would require a great deal of validation and verification.

Tractography algorithms have undergone intensive development since the introduction of DT-MRI and exciting applications are now beginning to emerge. By mapping fibre pathways in abnormal brains [52, 53], we can monitor disease progression and assist neurosurgical planning. Probabilistic tractography has lead to profound insights in human neuroanatomy [26, 54] and highlights region-connectivity differences between normal and patient groups [55]. Behrens et al [54] use probabilistic tractography to segment the human thalamus into regions that connect to different cortical regions. The segmentation they produce is consistent among individuals and similar to a connectivity-based segmentation of the monkey thalamus performed by histology. Barrick et al [56] use a similar idea to segment the whole human brain into connected regions. Behrens et al [57] also propose a method for automatic segmentation based on connectivity information. For the future, tractography and connectivity-mapping hold great promise for studies of brain development [58, 59].

5 Discussion

We have reviewed the principles of diffusion MRI measurements and reconstruction. We have seen how the standard diffusion MRI reconstruction algorithms, in particular DT-MRI and DSI, arise from a simple model relating the measurements to the spin-displacement density function p . We have highlighted the drawbacks of these basic approaches: DT-MRI provides only a single fibre-orientation estimate in each voxel and fails at fibre crossings; DSI requires too many measurements for routine use on current hardware. We have reviewed a new generation of multiple-fibre reconstruction algorithms, including multi-compartment and fibre models and all the methods in section 3.6, that can resolve the orientations of crossing fibres from sparse sets of measurements, similar to those acquired routinely in DT-MRI.

The new generation of reconstruction algorithms is still in its infancy and requires refinement and validation before routine application in clinical studies. As yet, no single algorithm has emerged as a comprehensive replacement for DT-MRI. The most tested new algorithms, which are multi-compartment models, PASMRI and \mathbf{q} -ball, all have problems. Multi-compartment models have problems with model selection and model fitting. The \mathbf{q} -ball algorithm does not resolve crossing fibres reliably on current standard data sets, but does produce acceptable results with a moderate increase in data quality [42]. The PASMRI algorithm works well on current data, but is too computationally heavy in practice. The great interest in fibre tractography continues to expand. Such are the problems caused by fibre crossings that development of the new algorithms will be rapid and we can expect to see them in routine use within the next few years.

We shall conclude with some specific questions for the further development of diffusion MRI reconstruction algorithms:

What new shape indices can we derive from multiple-fibre reconstructions?

Multiple-fibre reconstructions produce a range of new features of p that can provide new scalar indices of shape beyond the common anisotropy and skewness statistics. We can compute higher-order moments of these functions, which may highlight previously unseen tissue-type boundaries. The number of peaks and relative peak strength of the PAS, ODF or FOD may also provide useful stains for analysis and diagnosis.

How reliable are fibre-orientation estimates and shape indices?

We can determine the accuracy of fibre-orientation estimates in simulation from test functions, as in [41], and also with bootstrap methods from repeated scanner acquisitions, as in [60]. Such experiments on multiple-fibre reconstructions will provide performance comparisons for selecting the best algorithms and uncertainty estimates for probabilistic tractography. We can assess the reliability of shape indices in the same ways. The reliability determines the diagnostic power of shape indices as well as their potential as indicators of fibre-orientation-estimate accuracy in probabilistic tractography.

These performance estimates will depend on the imaging parameters, such as $|\mathbf{q}|$, t , N , M and S . We must optimize the trade-off between imaging time and data quality to maximize performance.

Can we detect and reject spurious structure in isotropic areas?

A major concern with multiple-fibre reconstructions is that they often show spurious structure in isotropic regions. Jansons and Alexander [41] and Alexander [61] illustrate this problem on synthetic data for PASMRI and \mathbf{q} -ball. On scanner data, these methods invariably produce functions with strong peaks in grey-matter and CSF regions. However, the fibre orientation estimates have little or no voxel-to-voxel coherence, which suggests that the peaks come from measurement noise rather than genuine anatomy. The success of the new generation of algorithms will require methods for distinguishing spurious from genuine structure. Voxel classification algorithms, such as Alexander's [22], may help solve the problem, as may methods that analyze voxel-to-voxel consistency.

Can we do better than spherical acquisition schemes?

In the literature, multi-compartment models and fibre models, ODF approximations, PASMRI and deconvolution methods mostly use data acquired with a spherical acquisition scheme. However, these methods can work with data acquired with any set of \mathbf{q}_i . Other distributions of sample points surely exist that will improve the methods, but the diffusion MRI community is yet to investigate methods for choosing and optimizing these distributions.

What is the best model for p in brain tissue?

The literature contains a variety of parametric models for p in white-matter fibres; see for example [30, 31, 39, 40]. Multi-compartment, fibre-model and deconvolution methods can use any such model. Quantitative comparisons of these models, again using simulations and bootstrap techniques, will determine which models best fit the data and produce the most reliable fibre-orientation estimates and shape statistics.

Can we estimate the fibre-orientation distribution reliably?

Diffusion-weighted MRI [4, 5, 6, 7] and one-dimensional \mathbf{q} -space imaging [5] were the first generation of diffusion MRI algorithms. The second generation, diffusion-tensor MRI and DSI, generalizes the first to three-dimensions. The third generation consists of the multiple-fibre reconstructions from data designed to emphasize the angular structure of p . Despite the names of the features of p that these algorithms compute ("orientation-distribution function" and "fibre-orientation distribution"), only the peaks of these functions are generally considered reliable as fibre-orientation estimates. Perhaps generation four will provide reliable estimates of the true distribution of fibre orientations within each voxel of an image and help distinguish crossing, kissing and diverging patterns of fibres.

The questions above are the tip of an iceberg. The remaining chapters of this book will reveal many other questions that demand answers in this exciting, expanding and fast-moving area of research.

Acknowledgements

The author would like to thank Dr. Claudia Wheeler-Kingshott from the Institute of Neurology, University College London, who provided the MRI data used to generate the images in this chapter. Dr. Dave Tuch, from the Athinoula A. Martinos Imaging Center for Biomedical Imaging, Massachusetts General Hospital, provided the code for the **q**-ball algorithm. Dr. Alexander's work is supported by the MIAS IRC, EPSRC GR/N14248/01 and MRC D2025/31, and EPSRC GR/T22858/01.

References

1. Pierpaoli C, Jezzard P, Basser P J, Barnett A and Di Chiro G 1996 Diffusion tensor MR imaging of the human brain *Radiology* **201** 637–48
2. Blinkov S and Glezer I 1968 *The human brain in figures and tables* (New York, USA: Plenum Press)
3. Highley J R, Esiri M M, McDonald B, Cortina-Borja M, Herron B M and Crow T J 1999 The size and fibre composition of the corpus callosum with respect to gender and schizophrenia: a post-mortem study *Brain* **122** 99–110
4. Stejskal E O and Tanner T E 1965 Spin diffusion measurements: spin echoes in the presence of a time-dependent field gradient *The Journal of Chemical Physics* **42** 288–92
5. Callaghan P T 1991 *Principles of Magnetic Resonance Microscopy* (Oxford, UK: Oxford Science Publications)
6. Merboldt K D, Hanicke W and Frahm J 1985 Self-diffusion NMR imaging using stimulated echoes *Journal of Magnetic Resonance* **64** 479–486
7. LeBihan D and Breton E 1985 Imagerie de diffusion in-vivo par resonance magnetique nucleaire *C. R. Acad. Sci. (Paris)* **301** 1109–1112
8. Lori N F, Conturo T E and Le Bihan D 2003 Definition of displacement probability and diffusion time in *q*-space magnetic resonance measurements that use finite-duration diffusion-encoding gradients *Journal of Magnetic Resonance* **165** 185–195
9. Mitra P P and Halperin B I 1995 Effects of finite gradient-pulse widths in pulsed-field-gradient diffusion measurements *Journal of Magnetic Resonance* **113** 94–101
10. Liu C, Bammer R, Acar B and Moseley M E 2004 Characterizing Non-Gaussian Diffusion by Using Generalized Diffusion Tensors *Magnetic Resonance in Medicine* **51** 924–937
11. Brihuega-Moreno O, Heese F P, Tejos C and Hall L D 2004 Effects of, and corrections for, cross-term interactions in *q*-space MRI *Magnetic Resonance in Medicine* **51** 1048–1054
12. Sijbers J, den Dekker A J, Van Audekerke J, Verhoye M, Van Dyck D 1998 Estimation of the noise in magnitude MR images *Magnetic Resonance Imaging* **16** 87–90.
13. Inglis B A, Bossart E L, Buckley D L, Wirth E D, Mareci T H 2001 Visualization of neural tissue water compartments using biexponential diffusion tensor MRI *Magnetic Resonance in Medicine* **45** 580–587.

14. Basser P J, Mattiello J and Le Bihan D 1994 MR diffusion tensor spectroscopy and imaging *Biophysical Journal* **66** 259–67
15. Wang Z, Vemuri B C, Chen Y and Mareci T 2003 A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from DWI *Proc. of 18th International Conference on Information Processing in Medical Imaging (Ambleside)* (Springer: LNCS 2732) 660–671
16. Basser P J and Pierpaoli C 1996 Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI *Journal of Magnetic Resonance Series B* **111** 209–19
17. Pajevic S and Pierpaoli C 1999 Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fibre tract mapping in the human brain *Magnetic Resonance in Medicine* **42** 526–540
18. Jones D K, Horsfield M A and Simmons A 1999 Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging *Magnetic Resonance in Medicine* **42** 515–525
19. Jones D K 2004 The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study *Magnetic Resonance in Medicine* **51** 807–815
20. Parker G J M and Alexander D C 2003 Probabilistic Monte Carlo Based Mapping of Cerebral Connections Utilising Whole-Brain Crossing Fibre Information *Proc. of 18th International Conference on Information Processing in Medical Imaging (Ambleside)* (Springer: LNCS 2732) 684–695
21. Frank L R 2002 Characterization of anisotropy in high angular resolution diffusion-weighted MRI *Magnetic Resonance in Medicine* **47** 1083–99
22. Alexander D C, Barker G J and Arridge S R 2002 Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data *Magnetic Resonance in Medicine* **48** 331–40
23. Armitage P and Berry G 1971 *Statistical methods in medical research* (Oxford, UK: Blackwell Scientific Publications)
24. Ozarslan E and Mareci T H 2003 Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging *Magnetic Resonance in Medicine* **50** 955–965
25. Press W H, Teukolsky S A, Vetterling W T and Flannery B P 1988 *Numerical Recipes in C* (New York, USA: Press Syndicate of the University of Cambridge)
26. Parker G J M, Luzzi S, Alexander D C, Wheeler-Kingshott C A M, Ciccarelli O and Lambon-Ralph M A 2004 Non-invasive structural mapping of two auditory-language pathways in the human brain *Neuroimage* In press
27. Blyth R, Cook P A and Alexander D C 2003 Tractography with multiple fibre directions *Proc. 11th Annual Meeting of the ISMRM (Toronto)* (Berkeley, USA: ISMRM) 240
28. Tuch D S, Reese T G, Wiegell M R, Makris N, Belliveau J W and Wedeen V J 2002 High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity *Magnetic Resonance in Medicine* **48** 577–582
29. Chen Y, Guo W, Zeng Q, He G, Vemuri B and Liu Y 2004 Recovery of intravoxel structure from HARD DWI *Proc. IEEE International Symposium on Biomedical Imaging (Arlington)* (IEEE)
30. Behrens T E J, Woolrich M W, Jenkinson M, Johansen-Berg H, Nunes R G, Clare S, Matthews P M, Brady J M and Smith S M 2003 Characterization and propagation of uncertainty in diffusion-weighted MR imaging *Magnetic Resonance in Medicine* **50** 1077–1088

31. Assaf Y, Freidlin R Z, Rohde G K and Basser P J 2004 New modelling and experimental framework to characterize hindered and restricted water diffusion in brain white matter *Magnetic Resonance in Medicine* **52** 965–978
32. Neuman C H 1974 Spin echo of spins diffusing in a bounded medium *Journal of Chemical Physics* **60** 4508–4511
33. Wedeen V J, Reese T G, Tuch D S, Dou J-G, Weiskoff R M and Chessler D 1999 Mapping fiber orientation spectra in cerebral white matter with Fourier-transform diffusion MRI *Proc. 7th Annual Meeting of the ISMRM (Philadelphia)* (Berkeley, USA: ISMRM) 321
34. Tuch D S 2002 *Diffusion MRI of Complex Tissue Structure* (Doctor of Philosophy in Biomedical Imaging at the Massachusetts Institute of Technology)
35. Tuch D S, Reese T G, Wiegell M R and Wedeen V J 2003 Diffusion MRI of complex neural architecture *Neuron* **40** 885–895
36. S Helgason 1999 *The Radon Transform* (Birkhäuser)
37. Lin C P, Tseng W Y I, Kuo L, Wedeen V J and Chen J H 2003 Mapping orientation distribution function with spherical encoding *Proc. 11th Annual Meeting of the ISMRM (Toronto)* (Berkeley, USA: ISMRM) 2120
38. Ozarslan E, Vemuri B C and Mareci T H 2004 Fiber orientation mapping using generalized diffusion tensor imaging *Proc. IEEE International Symposium on Biomedical Imaging (Arlington)* (IEEE)
39. Anderson A and Ding Z 2002 Sub-voxel measurement of fiber orientation using high angular resolution diffusion tensor imaging *Proc. 10th Annual Meeting of the ISMRM (Honolulu)* (Berkeley, USA: ISMRM) 440
40. Tournier J-D, Calamante F, Gadian D G and Connelly A 2004 Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution *NeuroImage* **23** 1176–1185
41. Jansons K M and Alexander D C 2003 Persistent Angular Structure: new insights from diffusion MRI data *Inverse Problems* **19** 1031–1046
42. Alexander D C 2004 A comparison of q -ball and PASMRI on sparse diffusion MRI data *Proc. 12th Annual Meeting of the ISMRM (Kyoto)* (Berkeley, USA: ISMRM) 90
43. Dong Q, Welsh R C, Chenevert T L, Carlos R C, Maly-Sundgren P, Gomez-Hasan D M and Mukherji S K 2004 Clinical applications of diffusion tensor imaging *Journal of Magnetic Resonance Imaging* **19** 6–18
44. Dou J, Reese T G, Tseng W-Y I and Wedeen V J 2002 Cardiac diffusion MRI without motion effects *Magnetic Resonance in Medicine* **48** 105–114
45. Hsu E W and Setton L A 1999 Diffusion tensor microscopy of the intervertebral disc annulus fibrosus *Magnetic Resonance in Medicine* **41** 992–999
46. Li T-Q 1997 Porous structure of cellulose fibers studied by Q-Space NMR Imaging *Langmuir* **13** 3570–3574
47. Mansfield P and Issa B 1994 Studies of fluid transport in porous rocks by echo-planar MRI *Magnetic Resonance Imaging* **12** 275–278
48. Mori S and van Zijl P C M 2002 Fiber tracking: principles and strategies — a technical review *NMR in Biomedicine* **15** 468–480
49. Lori N F, Akbudak E, Shimony J S, Cull T S, Snyder A Z, Guillory R K and Conturo T E 2002 Diffusion tensor fiber tracking of human brain connectivity: acquisition methods, reliability analysis and biological results *NMR in Biomedicine* **15** 494–515

50. Cook P A, Alexander D C and Parker G J M 2004 Modelling noise-induced fibre-orientation error in diffusion-tensor MRI *Proc. IEEE International Symposium on Biomedical Imaging (Arlington)* (IEEE)
51. Parker G J M and Alexander D C 2005 Probabilistic anatomic connectivity derived from the microscopic persistent angular structure of cerebral tissue. *Philosophical Transactions of the Royal Society* **In press**
52. Clark C A, Barrick T R, Murphy M M, Bell B A 2003 White matter fiber tracking in patients with space-occupying lesions of the brain: a new technique for neurosurgical planning? *NeuroImage* **20** 1601–1608
53. Mori S, Frederiksen K, van Zijl P C M, Steiltjes B, Kraut M A, Solaiyappan M and Pomper M G 2002 Brain white matter anatomy of tumor patients evaluated with diffusion tensor imaging *Annals of Neurology* **51** 377–380
54. 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, Ciccarelli O, Thompson A J, Brady J M and Matthews P M 2003 Non-invasive mapping of connections between human thalamus and cortex using diffusion imaging *Nature Neuroscience* **7** 750–757
55. Ciccarelli O, Toosey A T, Parker G J M, Wheeler-Kingshott C A M, Barker G J, Miller D H and Thompson A J 2003 Diffusion tractography based group mapping of major white-matter pathways in the human brain *NeuroImage* **19** 1545–1555
56. Barrick T R, Lawes I N and Clark C A 2004 Automatic segmentation of white matter pathways by application of a region growing algorithm *Proc. 12th Annual Meeting of the ISMRM (Kyoto)* (Berkeley, USA: ISMRM) 619
57. Behrens T E J, Johansen-Berg H, Drobnjak I, Brady J M, Matthews P M, Smith S M and Higham D J 2004 Delineation of functional subunits in the human cortex from diffusion based connectivity matrices *Proc. 12th Annual Meeting of the ISMRM (Kyoto)* (Berkeley, USA: ISMRM) 621
58. Mori S, Itoh R, Zhang J Y, Kaufmann W E, van Zijl P C M, Solaiyappan M and Yarowsky P 2001 Diffusion tensor imaging of the developing mouse brain *Magnetic Resonance in Medicine* **46** 18–23
59. Ulug A M 2002 Monitoring brain development with quantitative diffusion tensor imaging *Developmental Science* **5** 286–292
60. Jones D K 2003 Determining and visualizing uncertainty in estimates of fiber orientation from diffusion tensor MRI *Magnetic Resonance in Medicine* **49** 7–12
61. Alexander D C 2005 Multiple-fibre reconstruction algorithms in diffusion MRI *Annals of the New York Academy of Sciences* **In Press**