

# Diffusion Tensor Imaging

Derek K. Jones and Alexander Leemans

## Abstract

Diffusion tensor MRI (DT-MRI) is the only non-invasive method for characterising the microstructural organization of tissue in vivo. Generating parametric maps that help to visualise different aspects of the tissue microstructure (mean diffusivity, tissue anisotropy and dominant fibre orientation) involves a number of steps from deciding on the optimal acquisition parameters on the scanner, collecting the data, pre-processing the data and fitting the model to generating final parametric maps for entry into statistical data analysis. Here, we describe an entire protocol that we have used on over 400 subjects with great success in our laboratory. In the ‘Notes’ section, we justify our choice of the various parameters/choices along the way so that the reader may adapt/modify the protocol to their own time/hardware constraints.

**Key words:** Diffusion tensor, MRI, sampling schemes, pulse sequence, optimal, data quality.

---

## 1. Introduction

Diffusion tensor MRI (DT-MRI), developed in the early- to mid-1990s (1, 2), provides a means for non-invasively characterising the properties of soft tissue on a microstructural scale. It works by sensitising the MRI signal to the random molecular motion of water molecules (diffusion) by addition of ‘diffusion-encoding gradients’ to a standard MR pulse sequence (3). At the typical resolution of a DT-MRI experiment (2–3 mm voxel sizes), in the grey matter and cerebro-spinal fluid, the diffusion-weighted signal is independent of the direction in which the gradients are applied, and the diffusion appears to be *isotropic* (2). In white matter, water molecules diffuse more freely along the dominant fibre orientation than across them (4). This *anisotropy* of diffusion provides insights into the microstructural organisation

of the white matter. The simplest model that encapsulates this anisotropic diffusion is the diffusion tensor (1). By applying diffusion-encoding gradients in at least six non-collinear and non-coplanar orientations, one can estimate the six unknown elements of the diffusion tensor (1) and thus characterise the anisotropy. Further, the direction in which the diffusion-weighted signal has the greatest attenuation gives an indication of the dominant fibre orientation – which can be used to create voxel-wise maps of fibre orientation (e.g. 5) or pieced together to reconstruct continuous trajectories throughout the white matter (i.e. ‘tractography’) (e.g. 6–10). Such information has previously been unavailable in vivo – and so it is understandable that the technique has attracted huge interest from, and has enjoyed rapid uptake by, clinical and neuroscientific communities.

Here we present a protocol that begins by bringing the subject to the scanner room and preparing them for the diffusion tensor imaging data acquisition, through the acquisition of the data and pre-processing, through to estimation of the diffusion tensor in each voxel and subsequent computation of quantitative parametric maps. We stress the importance of checking the data at each stage of the pipeline to ensure that the data are robust.

It should be noted that what is ‘optimal’ for one laboratory – which can only allocate 10 min to DT-MRI on a 1.5 T system made by manufacturer X, for example, will be of little interest to the group that can allocate 30 min on a 3 T system made by manufacturer Y. Therefore, in the Notes section, details will be provided on how to proceed in choosing what is optimal for a given set of circumstances.

---

## 2. Materials

MRI Scanner: General Electric HDx 3.0 T system (*see Note 1*).

- a. Gradients: Twin-speed gradient system with gradient strength = 40 mT/m and maximum slew rate = 150 T/m/s (*see Note 2*).
- b. RF Coils: Whole-body birdcage coil used for RF transmit; eight-channel head coil (made by MRI Devices Corp.) used for RF receive (*see Note 3*).
- c. Scanner Software Capability: Software to provide diffusion tensor imaging capability (*see Note 4*).
- d. Peripherals: Adequate padding for the head (wedge cushions, etc.); hearing protection (ear plugs); a peripheral pulse-oximeter; a squeeze-bulb (for the participant to communicate to the operator).

---

### 3. Methods

1. *Preparing the Participant:* After successful completion of standard MR screening and appropriate informed consent forms, the participant is led to the MR magnet room. They are given earplugs before being placed onto the scanner bed in the supine position. The pulse-oximeter is then placed onto the subject's forefinger, and they are given the squeeze-bulb. We take special care to warn the participant that the diffusion tensor imaging part of the protocol is '*louder than the other scans*' and that they '*can expect the bed to vibrate quite a lot*' (see **Note 5**). We also warn the participant that '*there will be irregularly timed knocking noises – and these will appear to move about as the scan progresses. This is fully expected*' (see **Note 5**).
2. *Scanning:* The integrated laser alignment system is used to landmark on the nasion, and the participant slid into the magnet, taking particular care not to trap the squeeze-bulb/pulse-oximeter leads during the process. As an optional extra, we provide the participant with the option of watching a subtitled movie of their choice in the scanner via a rear projection onto a periscope mounted on the head coil (see **Note 6**).

The sequence is a twice-refocused spin-echo EPI sequence (**11**) (see **Note 7**), with a parallel imaging (ASSET) factor of 2 (see **Note 8**). Sixty axially oriented slices are prescribed to cover the entire head (see **Note 9**). The field of view is 230 mm, with an acquisition matrix of  $96 \times 96$  and a slice thickness of 2.4 mm (see **Note 10**). A total of 66 images are acquired (see **Note 11**) at each of 60 slice locations. Six images are acquired with no diffusion-weighting gradients applied, and 60 diffusion-weighted images are acquired at a  $b$ -value of  $1,000 \text{ s/mm}^2$  (see **Note 12**). The diffusion-weighted images are acquired with encoding gradients applied along 60 non-collinear directions (see **Table 6.1** for the configurations – and see **Note 13**). The echo time is 87 ms (see **Note 14**), and the sequence is triggered to the cardiac cycle via a pulse-oximeter placed on the participant's forefinger (see **Note 15**). We select a minimal trigger delay (see **Note 16**), and the effective TR that is set depends on the heart rate estimated from the pulse-oximeter trace. Each image is initially stored in DICOM format. We then convert the separate DICOM images into a 4D data set (with 'time' or 'diffusion-weighted measurement' as the fourth dimension) in the NIFTI imaging format.

**Table 6.1**  
**60 Electrostatically arranged sampling vectors (21)**  
**optimally ordered according to Cook et al. (24)**

0.1706	0.3255	0.9300
0.3933	0.9154	-0.0853
-0.6263	-0.2862	0.7252
-0.6058	0.7763	0.1742
0.9204	-0.0040	0.3911
0.2903	-0.6194	0.7294
-0.7022	0.3532	0.6182
0.3712	0.7702	0.5187
-0.9738	-0.0801	0.2126
0.2158	-0.9371	0.2743
-0.0604	-0.2236	0.9728
0.8147	0.5611	0.1462
-0.2150	0.7211	0.6587
0.7166	-0.5846	0.3804
0.5354	-0.1672	0.8279
-0.2157	-0.7327	0.6454
0.6549	0.3557	0.6668
-0.0601	0.9726	0.2244
-0.7389	-0.5195	0.4290
-0.9432	0.3042	0.1338
-0.4428	0.2032	0.8733
-0.1629	-0.9173	0.3634
0.0988	0.6380	0.7637
0.7879	-0.2745	0.5513
-0.3122	-0.4225	0.8509

3. *Initial Quality Check*: The data are loaded into FSLview (part of the FSL software package) ([www.fmrib.ox.ac.uk/fsl](http://www.fmrib.ox.ac.uk/fsl)) and viewed in all three planes simultaneously – in ‘cine’ mode – looping through the separate volumes to check for any obvious artefacts in the data (see **Fig. 6.1a** and **Note 17**).
4. *Correcting Motion/Distortion*: We register the 2nd through 66th diffusion-weighted volumes to the 1st diffusion-weighted volume using a global, 12 degrees of freedom affine deformation, with normalised mutual information as the cost function (see **Note 18**). We extract the rotational component of the transformation for each volume and apply this rotation to the gradient encoding tables (see **Note 19**).

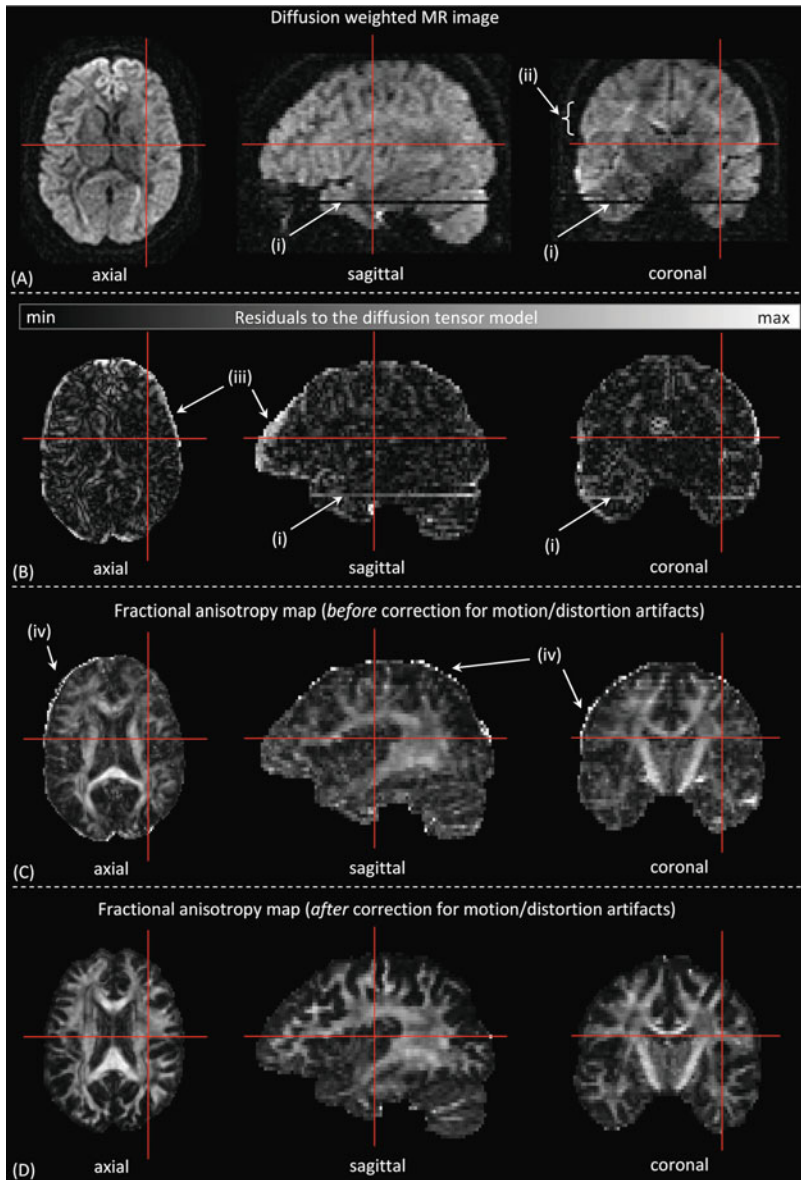


Fig. 6.1. Example of the data quality assessment through visual inspection. **a** Visualisation of the raw diffusion-weighted data. This is done as a cine loop to quickly view all diffusion-weighted volumes. Abnormal data values, such as signal dropouts (i) or hyperintensities (ii), are easily detectable on the planes orthogonal to the acquisition plane (axial), i.e. the coronal and sagittal image views. **b** Looking at the data residuals to the diffusion tensor fit may reveal other data abnormalities, for instance, in the form of hyperintensities around the rim of the brain (iii), which can also be observed in the FA maps (c). This high anisotropy rims suggest image misalignment across the different diffusion-weighted images due to subject motion or geometric distortions. In (d), the FA map is shown after performing this correction procedure.

5. *Re-inspecting the Data*: We re-inspect the data in three orthogonal planes in a cine format to ensure that the motion/distortion correction has been performed correctly and that no additional artefacts have been introduced into the data.
6. *Fitting the Diffusion Tensor*: We perform an initial ordinary least squares fit of a single Gaussian diffusion tensor to the data to generate a starting estimate for subsequent entry into a nonlinear least squares (Levenberg–Marquardt) algorithm to estimate the tensor in each voxel (*see Note 20*).
7. *Further Inspection of the Data*: We generate a 4D data set of the residuals to the tensor fit (**Fig. 6.1b**) and again view these in three orthogonal planes as a movie to look for obvious outliers/artefacts (*see Note 21*).
8. *Computation of Parametric Maps*: The tensor in each voxel is diagonalised to derive the eigenvectors and eigenvalues, and parametric maps of the mean diffusivity and fractional anisotropy (2) are computed (*see Note 22*). We also compute the directionally encoded colour (DEC) map showing fibre orientation (5).
9. *Final Inspection of the Data*: Before passing any maps into further analyses, we visualise the FA, MD and DEC maps in three orthogonal orientations to check for any obvious artefacts, such as rims of high anisotropy at the edge of the brain (*see Fig. 6.1c and Note 23*).

---

## 4. Notes

1. In most research environments for human imaging, the field strength is 1.5 or 3.0 T with a few exceptions. The advantage of higher field is higher signal-to-noise ratio (SNR) per unit time – allowing higher resolution for fixed scan time, shorter scan time for the same resolution or higher precision in the data if all scan parameters are kept the same. The disadvantage is that the standard approach to DT-MRI (1, 2) uses echo-planar imaging (EPI) – which renders the images to be very sensitive to differences in magnetic susceptibility (such as at air–tissue interfaces) – leading to either a stretch or a compression of the image along the phase encode direction (usually aligned with the anterior–posterior axis of the head), and these distortions become worse at higher field strength.
2. The imaging gradients should also ideally have the capability of producing gradient amplitudes of above 20 mT/m. The stronger the gradient amplitudes, the better for

diffusion MRI. The key factors in determining the amount of diffusion weighting, characterised by the ‘*b*-factor’, *b*, are the amplitude (*G*), the duration ( $\delta$ ) and the temporal separation ( $\Delta$ ), given by the so-called Stejskal–Tanner equation (3):

$$b = \gamma^2 G^2 \delta^2 \left( \Delta - \frac{\delta}{3} \right),$$

where  $\gamma$  is the gyromagnetic ratio.

Increasing the gradient amplitude means that smaller values of  $\delta$  and  $\Delta$  can be employed, which in turn means shorter echo times and therefore increased SNR.

3. A multi-channel head coil is preferable for improved SNR and the possibility to employ parallel imaging, which in turn helps to reduce the EPI-based distortions (12). Eight-channel head coils are prevalent in neuroimaging research centres, but again – more channels, if available, are preferred with 12-channel and even 32-channel coils becoming purchasable options. Again – these will boost the SNR – which is most definitely beneficial for DT-MRI. Further, more channels permit the use of higher speed-up factors in parallel imaging acquisition strategies, which is advantageous for DT-MRI (12).
4. The scanner should provide the capability of applying diffusion-encoding gradients in at least six different orientations (1), although more is better (*see Note 13*). The capability to choose one’s own encoding directions (normally facilitated under a research agreement with the manufacturer) is preferred, particularly if the sampling orientations are to be optimally ordered (*see Paragraph 2 of Note 13*).
5. Past experience has shown that it is beneficial to alert the participant to the fact that the vibrations can be quite severe. This is particularly true with the twice-refocused pulse sequence – as reported in (13). Moreover, it is useful to warn the participant that there will be spatially varying irregularly spaced knocking which refers to the fact that the TR will be non-uniform due to the cardiac gating, and the knocking is the gradients being played out in different orientations.
6. The use of a video projection system, at least anecdotally, appears to reduce head motion as it engages the participant and reduces the likelihood of them looking around (which may result in additional head motion) or being distracted by other thoughts over what can be a considerable duration of scan (sometimes up to 30 min). It is not expected that there will be functionally dependent changes in the signal

acquired with a standard DT-MRI protocol that will impact on standard DT-MRI analyses.

7. The most commonly used acquisition scheme for DT-MRI is now the twice-refocused spin-echo EPI sequence (11) and is provided by most manufacturers as part of the DT-MRI package that they sell. The twice-refocused sequence will markedly reduce the effects of eddy current-induced distortions resulting from the rapid switching of the diffusion-encoding gradients.
8. For EPI-based DT-MRI acquisitions, there is a clear benefit to the use of parallel imaging (12) using approaches such as the image domain SENSE approach (14) or  $k$ -space domain approaches such as GRAPPA (15). The choice of which to use normally depends on availability and/or software provided by the manufacturer. A parallel imaging factor of 2 seems to provide a reasonable compromise between SNR, distortion reduction and speed-up factor.
9. Pure axial orientation gives the best quality data in terms of ghosts and distortions.
10. For most purposes, it is highly desirable to have isotropic imaging voxels for DTI so that there is no preferential averaging of fibre orientations along a particular axis. This is particularly important for tractography applications.

The limitations to consider are

(1) The slice profile: On some scanners (e.g. General Electric), the issue of the fat-water frequency shift, particularly problematic in EPI, is addressed through the use of a spatially and spectrally selective pulse – which leaves the fat (in the scalp, for example) unexcited. While this is an alternative to implicitly turning on ‘FatSat’, it does tend to limit the minimal slice thickness that can be achieved (around 2.5 mm is typical).

(2) The signal-to-noise ratio: It is worth bearing in mind that the *average* SNR in the diffusion-weighted image is on the order of 30% of that in the non-diffusion weighted image. If the non-diffusion weighted signal intensity is  $I_0$  then the diffusion-weighted signal is  $I_0 \exp(-bD)$ , where  $D$  is the apparent diffusion coefficient along the axis of the applied encoding gradient. As discussed below, the  $b$ -value is typically on the order of the reciprocal of  $D$  so that  $I = I_0 \exp(-1) = 0.33I_0$ . For isotropic media, all diffusion-encoding directions will have the same attenuation. It is important that the SNR in the diffusion-weighted images does not go below approximately 3:1. This is the domain where the Rician-distributed data begin to look non-Gaussian (16, 17), and one encounters problems with the rectified noise floor that cause underestimation of dif-



fusivities (18) and corrupt estimates of diffusion anisotropy (19), among other problems. For anisotropic media, say with a fractional anisotropy (20) of 0.7, the largest eigenvalue will be twice the mean diffusivity. This, in turn, means that the signal attenuation will be equal to  $\exp(-2bD) = 0.13$ . Therefore to void the rectified noise floor issue, the SNR in the non-diffusion weighted image should be at least  $3:0.13 \sim 22:1$ . It is therefore advisable, when setting up the protocol, to make some SNR measurements (with the coil, parallel imaging factor, etc., for the experiment already chosen and in place) before settling on the final resolution. In summary, one should go for the highest resolution possible that achieves isotropic resolution but ensures that the SNR in white matter in the non-diffusion image is greater than 20. The field of view should be sufficient to cover the entire head and selected in consideration of the image acquisition matrix and slice thickness to ensure isotropic voxels.

11. The more measurements are made, the more accurate and precise the parameters derived from DTI will be. The total number of measurements will therefore be dictated by the time allocated in the protocol. The ratio of measurements made at the higher  $b$ -value to those with no diffusion weighting should be around 8–10:1 (21); it is highly recommended that the images with  $b = 1,000 \text{ s/mm}^2$  are acquired with gradients distributed as uniformly in space as possible (21, 22 – see **Note 13**). If time permits, 30 directions (with two or three non-diffusion weighted scans) will give measures that are statistically rotationally invariant (22) – but if time permits even more measurements to be made, then one should take advantage and acquire more points.
12. It is recommended that just two different diffusion weightings are used: one close to zero and one close to  $b = 1,000 \text{ s/mm}^2$ , the latter being a compromise between ensuring sufficient attenuation for precise estimation of diffusivities and balancing this against the penalty of increasing the echo time to accommodate the diffusion-encoding gradients, which in turn leads to signal loss through  $T_2$  relaxation (21).
13. For estimating the tensor, it is beneficial to space out the sampling gradients as uniformly in space as possible (21, 22). This helps to reduce the dependence of the reproducibility (variance) of measurements on the orientation of the material being imaged (22). The de facto standard, adopted by most manufacturers and research groups, is now to use an ‘electrostatic repulsion’ algorithm (21) to

space out the gradients. This is done by simulating what happens if the gradient sampling vectors are all rods passing through the origin (so that each rod points in the direction of a sampling vector and its antipode) and if there are point charges placed at the ends of each rod. The orientations of the rods are then changed until the sum of repulsive forces between all possible pairs of charges is minimised (this is effectively what happens in the formation of crystals in nature – so it is no surprise that, for appropriate numbers of sampling vectors, the orientations that are produced by this algorithm produce the regular polyhedral arrangements – such as the tetrahedral arrangement of  $sp^3$  hybrid orbitals seen in crystals (23)).

If the scan is run to completion, then the order in which the different sampling orientations are played out is irrelevant – as the complete data set is available and all points go into the estimation of the tensor. However, in subjects where it is anticipated that there may be a curtailment (for example, due to excessive motion/claustrophobia or general non-compliance), stopping a randomly ordered acquisition before completion could lead to, for example, lots of measurements being made along similar orientations – with large portions of the sampling space left out (which is sub-optimal). To counter this, Cook et al. (24) and Dubois et al. (25) have independently proposed ways of re-ordering the directions derived from the electrostatic repulsion algorithm so that, if (for example) a 30-direction sampling scheme is interrupted half-way through, the 15 directions that have been collected are still pretty much uniformly distributed. This seems an eminently sensible strategy and, if the full data set is collected anyway, has no impact on the data since, as stated before, the ordering is immaterial. **Table 6.1** lists the example of 60 electrostatically arranged sampling vectors that have been ordered according to this principle and are used in our protocol. For the sake of completeness, in **Tables 6.2** and **6.3**, we tabulate the optimal point ordering for 18 and 30 unique sampling orientations, respectively. These were obtained with the Camino package (26), which could be used to generate other numbers of ordered point sets.

14. One is always battling with SNR issues in DT-MRI (27). Therefore, one should strive to minimise additional attenuation – consequently, whenever possible, having selected the  $b$ -value on the scanner console, one should select the ‘minimum TE’ setting. This can force the acquisition into a partial  $k$ -space acquisition, with, for example, 8 lines of  $k$ -space acquired before the echo and 48 lines acquired

**Table 6.2**  
**18 Electrostatically arranged and optimally ordered**  
**sampling vectors**

0.7371	-0.5680	0.3662
0.7958	0.4311	0.4253
-0.8225	0.3677	0.4339
0.0006	0.9856	0.1692
0.2290	0.1508	0.9617
-0.4124	-0.7535	0.5120
-0.3586	0.2328	0.9040
-0.8912	-0.4176	0.1768
0.3199	-0.4987	0.8056
0.3099	0.6677	0.6769
0.5797	-0.8070	-0.1124
-0.2096	-0.3585	0.9097
0.9907	-0.1123	0.0774
0.1533	-0.9033	0.4008
0.5302	0.8454	0.0651
-0.2829	0.7167	0.6374
0.7201	-0.0527	0.6919
-0.7339	-0.1786	0.6554

afterwards. In some cases, it may be necessary to increase the number of lines of  $k$ -space before the centre of  $k$  to avoid corruption of the data due to excessive vibration.

15. It has been shown in several studies that diffusion MRI measurements can be severely corrupted by the effect of cardiac pulsation (28–31). As the heart beats – the pressure wave is carried through to the brain – and one gets both local deformation of the tissue (32, 33), which results in local misregistration between successive images, and intra-voxel dephasing due to diffusion (30). The latter will be interpreted as increased diffusion measured at the point a particular encoding gradient is applied – which, in turn, will lead to biases in the estimate of the diffusion tensor, resulting in inaccurate estimates of both anisotropy and fibre orientation (31). To avoid this, it is recommended that the acquisition be timed so as to avoid acquiring data during the time that the brain tissue is susceptible to this pulsation.
16. We have previously found that corruption of data due to cardiac pulsation can be avoided by waiting until at least

**Table 6.3**  
**30 Electrostatically arranged and optimally ordered**  
**sampling vectors**

0.0559	-0.9920	-0.1134
-0.6669	-0.6780	0.3091
0.1637	0.5330	0.8301
-0.4320	-0.0898	0.8974
0.8865	0.2190	0.4076
0.7671	-0.6107	0.1964
0.3840	-0.2940	0.8753
-0.5464	0.6192	0.5639
-0.9500	0.0308	0.3108
0.0973	-0.7119	0.6955
-0.7308	-0.6541	-0.1950
0.0777	0.0950	0.9924
-0.3115	0.9086	-0.2781
0.2961	0.8430	0.4491
-0.7523	-0.3014	0.5858
0.7882	-0.2137	0.5771
0.8458	-0.4786	-0.2358
-0.1722	-0.9086	0.3806
-0.2838	0.3816	0.8797
0.5428	0.1339	0.8291
-0.3989	-0.6000	0.6934
0.5072	-0.8486	-0.1505
-0.9375	-0.3410	0.0689
0.1467	-0.8120	-0.5649
-0.9757	0.1655	-0.1434

220 ms after the onset of the R-wave in the ECG trace recorded from chest leads (30). In turn, we have also found that the arrival of the pulse wave on a peripheral pulse-oximeter placed on the finger is  $249 \pm 17$  ms (mean  $\pm$  SD) after the R-wave of the ECG (30). Consequently, the delay needed after the peak on a peripheral oximeter trace to avoid DT-MRI data corruption is minimal. It is therefore beneficial to use a peripheral oximeter – as there is less ‘dead time’ in each pulse-to-pulse interval, and it is undoubtedly more convenient for the participant. One point to note is that since the repetition time will then vary due to natural variations in the cardiac cycle, it is important that the effective TR is at least five times the  $T_1$  of the tissue

of interest (to avoid partial recovery effects). However, in practice – this condition is almost always satisfied – as one can rarely squeeze in more than three or four slices into an R–R interval and, for a typical 60-slice acquisition, this means an effective TR of 15 or 20 R–R intervals. We use a look-up table to rapidly determine the optimal effective TR (Table 6.4).

**Table 6.4**  
**Prescribed effective TR as a function of participant heart rate**

Participant heart rate	Effective TR
<50 bpm	15 R–R intervals
50–64 bpm	20 R–R intervals
65–95 bpm	30 R–R intervals
>95 bpm	60 R–R intervals

17. The first and most important thing that must be done is to examine the raw data. With pressures to get studies completed quickly, this is an often overlooked step in many laboratories. However, given that one is going to take multiple images and use them to compute the diffusion tensor at each voxel (1) – it is important to check that there are no corruptions in any of the individual data points. It can be extremely informative to view the data in three orthogonal planes simultaneously (see Fig. 6.1a). There are multiple tools available for doing this (for example, the popular *FSLview* from the FSL software library). This allows for rapid identification of slice-to-slice intensity variations (for an axially acquired data set, these will be visible on the coronal and sagittal planes). With the data set stored in a 4D format (the fourth dimension being the number of the diffusion-weighted scan), viewing the data set as a looping movie is a very efficient way of checking the data for unexpected signal dropouts (which will be visible as unusually dark horizontal bands on the sagittal and coronal slices) or other artefacts (see Fig. 6.1a).
18. It is extremely important to correct the data for subject motion and eddy current-induced distortions (e.g. 34, 35). Although the twice-refocused spin-echo sequence will ameliorate much of the eddy currents (11), there may still be residual distortions that need to be taken care of and it is unlikely that the participant will have remained perfectly still. The ‘industry standard’ approach is to use a global affine registration (12 degrees of freedom – with translation, rotation, magnification and shear along each of the

principal axes) to register each diffusion-weighted volume to the volumes collected without any diffusion-weighting applied (the ‘ $b = 0$  s/mm<sup>2</sup> images’). Given that the image contrast is so different between the  $b = 0$  image and those acquired at  $b = 1,000$  s/mm<sup>2</sup>, cost functions for the registration such as cross-correlation tend to fail and much better results are obtained with entropy-based metrics, such as the mutual information index (and its normalised version) (36, 37). Again, there are many software packages that cater to such global affine-based registration requirements, allowing one to specify the cost function.

19. Estimation of the tensor matrix requires exact knowledge of the orientation of the diffusion-encoding gradient with respect to the participant’s head (1). If the participant moves their head during the acquisition – this can be corrected with image realignment methods – as just discussed. However, simple naïve application of ‘off-the-shelf’ registration software will not account for the fact that, during such a rotation, the angle between the participant’s head and the pre-selected gradient sampling vectors will change. Failing to account for this can lead to substantial errors in estimates of anisotropy and of fibre orientation (38). Therefore, it is desirable that, when available, the gradient table that is used as input into the tensor estimation routine is modified accordingly. The amount of diffusion encoding for a particular combination of gradients is characterised by the  $b$ -matrix (1). The effect of rotation can be handled by deriving the rotational part of the transformation required to realign the images and subsequently applying this rotation to the  $b$ -matrix prior to estimation of the tensor. **Figure 6.2** shows the effect of neglecting to perform this step on estimates of anisotropy and fibre orientation.
20. There are three widely used approaches to estimating the diffusion tensor from the  $b$ -matrix and the diffusion-weighted data: ordinary linear least squares (OLLS), weighted linear least squares (WLLS) (1) and nonlinear least squares (NLLS) (39). (Note that there are lots more approaches in the literature – but these are the most common.) It is outside the scope of this chapter to go into the fine detail of these different approaches – but standard mathematical processing packages provide access to all of these. For the first two approaches, i.e. linear fitting, the diffusion-weighted signal intensities are first log transformed (1). In OLLS, each observation contributes equally to the fit – and thus a set of simultaneous equations relating the log of the signal to the unknown elements of the tensor are set up – and a simple matrix




Fig. 6.2. Directionally encoded fibre orientation maps without (a) and with (b) the required  $b$ -matrix rotation prior to estimating the diffusion tensor. The differences in orientation (colour) and fractional anisotropy (intensity) between the images is clearly visible, as shown, for instance, with the enlarged view of the posterior region of the corpus callosum.

inversion yields the unknown elements of the tensor. Given the rapid nature of this approach, it is extremely popular and is employed in several popular software packages (e.g. FSL). The consequence, however, of taking the log transformation is that, although the noise/random errors in the signal prior to the log transformation are uniform (i.e. *homoscedastic*) – after the log transformation, the variance in the (log-transformed) signal becomes a function of the signal itself (1, 40) and so the errors are *heteroscedastic*. To properly address this, a weighted linear regression (WLLS) is required, where one has to compute a covariance matrix – deriving the relationship between the variance in the log-transformed and non-log transformed data – and include it in the regression step (1). Although this makes computation longer, the results are far more robust and WLLS is to be preferred over OLLS (39). In NLLS, on the other hand, there is no log transformation of the signal – and thus the errors remain homoscedastic, so the covariance matrix is a multiple of the identity matrix and can effectively be factored out of the analysis, which is an advantage over the linear framework approaches. While NLLS is attractive in that it fits the model to the data directly, producing results that are superior to WLLS (and therefore to OLLS) (39), the computational time is considerably longer – and special care has to be taken that the fitting algorithm has not been trapped in a local extremum.

21. Once one has fitted the model, it is expected that any differences between the observed signal and those predicted by the model (i.e. the ‘residuals’) should be random and

should not contain any structure. Deviations will rapidly identify points that are corrupt (41) (see Fig. 6.1b). Therefore, before proceeding with further analyses, we recommend inspecting a map of the residuals to ensure that there is nothing unexpected in the data.

22. The final stage in a standard DT-MRI pipeline is to derive parameters of interest from the diffusion tensor. These are invariably the mean diffusivity (or trace of the diffusion tensor) (2), a measure of anisotropy (27) – and the principal diffusion orientation, which can be used for creation of directionally encoded colour (DEC) maps (5) or for fibre-tracking analyses (6–10). With regard to the choice of anisotropy index, there is a plethora of indices in the literature (20). However, by far the most popular is the fractional anisotropy, derived from the three eigenvalues of the diffusion tensor. One can show that this has a better signal-to-noise ratio characteristic than other popular indices such as relative anisotropy (42).
23. It is always prudent to look carefully at the resultant parametric maps before inputting them into any form of analysis. Values of anisotropy greater than 1 are physically non-meaningful, since they are designed to lie between 0 and 1 – and will result when one or more of the eigenvalues is negative. Again, this is non-physically meaningful – but can occur when the diffusion-weighted signal is higher in intensity than the non-diffusion weighted signal. In turn, this may arise in regions that are particularly noisy – or where there is insufficiently corrected misregistration. A similar tell-tale sign for artefact is to examine the rim of the anisotropy maps. A high anisotropy around the rim of the adult mammalian brain is not expected in practice (Fig. 6.1c). Given that high anisotropy means that the DW signal varies rapidly with the direction of the diffusion-encoding gradient, the appearance of such a bright rim would be consistent with misregistration of, for instance, CSF and gray matter with either white matter or perhaps the edge of the brain.

## References

1. Basser, P. J., Mattiello, J., LeBihan, D. MR diffusion tensor spectroscopy and imaging. *Biophys J* 1994;66: 259–267.
2. Pierpaoli, C., Jezzard, P., Basser, P. J., Barnett, A., Di Chiro, G. Diffusion tensor MR imaging of the human brain. *Radiology* 1996;201:637–648.
3. Stejskal, E. O., Tanner, J. E. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *J Chem Phys* 1965;42:288–292.
4. Moseley, M. E., Cohen, Y. C., Kucharczyk, J., Asgari, H. S., Wendland, M. F., Tsuruda, J., Norman, D. Diffusion-weighted MR imaging of anisotropic water diffusion



- in cat central nervous system. *Radiology* 1990;176:439–445.
5. Pajevic, S., Pierpaoli, C. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magn Reson Med* 1999;42:526–540 (Erratum in: *Magn Reson Med* 2000 43:921).
  6. Mori, S., Crain, B. J., Chacko, V. P., van Zijl, P. C. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann Neurol* 1999;45: 265–269.
  7. Jones, D. K., Simmons, A., Williams, S. C., Horsfield, M. A. Non-invasive assessment of axonal fiber connectivity in the human brain via diffusion tensor MRI. *Magn Reson Med* 1999;42:37–41.
  8. Conturo, T. E., Lori, N. F., Cull, T. S., Akbudak, E., Snyder, A. Z., Shimony, J. S., McKinstry, R. C., Burton, H., Raichle, M. E. Tracking neuronal fiber pathways in the living human brain. *Proc Natl Acad Sci USA* 1999;96:10422–10427.
  9. Basser, P. J., Pajevic, S., Pierpaoli, C., Duda, J., Aldroubi, A. In vivo fiber tractography using DT-MRI data. *Magn Reson Med* 2000;44:625–632.
  10. Parker, G. J. M., Haroon, H. A. et al. A framework for a streamline-based probabilistic index of connectivity (PICO) using a structural interpretation of MRI diffusion measurements. *J Magn Reson Imag* 2003;18:242–254.
  11. Reese, T. G., Heid, O., Weisskoff, R. M., Wedeen, V. J. Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magn Reson Med* 2003;49:177–182.
  12. Bammer, R., Auer, M., Keeling, S. L., Augustin, M., Stables, L. A., Prokesch, R. W., Stollberger, R., Moseley, M. E., Fazekas, F. Diffusion tensor imaging using single-shot SENSE-EPI. *Magn Reson Med* 2002;48:128–136.
  13. Hiltunen, J., Hari, R., Jousmäki, V., Müller, K., Sepponen, R., Joensuu, R. Quantification of mechanical vibration during diffusion tensor imaging at 3 T. *Neuroimage* 2006;32: 93–103.
  14. Pruessmann, K. P., Weiger, M., Scheidegger, M. B., Boesiger, P. SENSE: Sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
  15. Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, K., Kiefer, B., Haase, A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47:1202–1210.
  16. Edelstein, W. A., Bottomley, P. A., Pfeifer, L. M. A signal-to-noise calibration procedure for NMR imaging systems. *Med Phys* 1984;11:180–185.
  17. Gudbjartsson, H., Patz, S. The Rician distribution of noisy MRI data. *Magn Reson Med* 1995;34:910–914.
  18. Dietrich, O., Heiland, S., Sartor, K. Noise correction for the exact determination of apparent diffusion coefficients at low SNR. *Magn Reson Med* 2001;45:448–453.
  19. Jones, D. K., Basser, P. J. “Squashing peanuts and smashing pumpkins”: How noise distorts diffusion-weighted MR data. *Magn Reson Med* 2004;52:979–993.
  20. Basser, P. J., Pierpaoli, C. Microstructural and physiological features of tissue elucidated by quantitative-diffusion-tensor MRI. *J Magn Reson B* 1996;111:209–219.
  21. Jones, D. K., Horsfield, M. A., Simmons, A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. *Magn Reson Med* 1999;42: 515–525.
  22. Jones, D. K. The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: A Monte Carlo study. *Magn Reson Med* 2004;51:807–815.
  23. Conturo, T. E., McKinstry, R. C., Akbudak, E., Robinson, B. H. Encoding of anisotropic diffusion with tetrahedral gradients: A general mathematical diffusion formalism and experimental results. *Magn Reson Med* 1996;35:399–412.
  24. Cook, P. A., Symms, M., Boulby, P. A., Alexander, D. C. Optimal acquisition orders of diffusion-weighted MRI measurements. *J Magn Reson Imaging* 2007;25: 51–58.
  25. Dubois, J., Poupon, C., Lethimonnier, F., Le Bihan, D. Optimized diffusion gradient orientation schemes for corrupted clinical DTI data sets. *Magma* 2006;19: 134–143.
  26. Cook, P. A., Bai, Y., Nedjati-Gilani, S., Seunarine, K. K., Hall, M. G., Parker, G. J., Alexander, D. C. (2006) Camino: Open-Source Diffusion-MRI Reconstruction and Processing, 14th Scientific Meeting of the International Society for Magnetic Resonance in Medicine, Seattle, WA, USA, p. 2759, May 2006.
  27. Pierpaoli, C., Basser, P. J. Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med* 1996;36:893–906.
  28. Turner, R., Le Bihan, D., Maier, J., Vavrek, R., Hedges, L. K., Pekar, J. Echo-planar

- imaging of intravoxel incoherent motions. *Radiology* 1990;177:407–414.
29. Skare, S., Andersson, J. L. On the effects of gating in diffusion imaging of the brain using single shot EPI. *Magn Reson Imaging* 2001;19:1125–1128.
  30. Pierpaoli, C., Marengo, S., Rohde, G., Jones, D. K., Barnett, A. S. (2003). Analyzing the contribution of cardiac pulsation to the variability of quantities derived from the diffusion tensor in “Proc. ISMRM 11th Annual Meeting, Toronto” p. 70.
  31. Jones, D. K., Pierpaoli, C. (2005) The contribution of cardiac pulsation to variability in tractography results. In “Proc. ISMRM 13th Annual Meeting, Miami”, p. 225.
  32. Poncelet, B. P., Wedeen, V. J., Weisskoff, R. M., Cohen, M. S. Brain parenchyma motion: Measurement with cine echo-planar MR imaging. *Radiology* 1992;185: 645–651.
  33. Enzmann, D. R., Pelc, N. J. Brain motion: Measurement with phase-contrast MR imaging. *Radiology* 1992;185: 653–660.
  34. Rohde, G. K., Barnett, A. S., Basser, P. J., Marengo, S., Pierpaoli, C. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magn Reson Med* 2004;51:103–114.
  35. Andersson, J. L., Skare, S. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *Neuroimage* 2002;16: 177–199.
  36. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P. Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imaging* 1997;16:187–198.
  37. Studholme, C., Constable, R. T., Duncan, J. S. Accurate alignment of functional EPI data to anatomical MRI using a physics-based distortion model. *IEEE Trans Med Imaging* 2000;19:1115–1127.
  38. Leemans, A., Jones, D. K. The B-matrix must be rotated when motion correcting diffusion tensor imaging data. *Magn Reson Med* 2009;61:1336–1349.
  39. Koay, C. G., Chang, L. C., Carew, J. D., Pierpaoli, C., Basser, P. J. A unifying theoretical and algorithmic framework for least squares methods of estimation in diffusion tensor imaging. *J Magn Reson* 2006;182: 115–125.
  40. Bevington, P. R., Robinson, D. K. *Data Reduction and Error Analysis for the Physical Sciences*, 2nd ed. New York, NY: McGraw-Hill; 1992.
  41. Leemans, A., Evans, C. J., Jones, D. K. (2008). Quality assessment through analysis of residuals of diffusion image fitting. In “Proc. ISMRM 16th Annual Meeting, Toronto”. p. 3300.
  42. Hasan, K. M., Alexander, A. L., Narayana, P. A. Does fractional anisotropy have better noise immunity characteristics than relative anisotropy in diffusion tensor MRI? An analytical approach. *Magn Reson Med* 2004;51:413–417.