18

Artifacts in Diffusion MRI

CARLO PIERPAOLI

Summary Box

- This chapter focuses on artifacts that are relevant for diffusion imaging studies of the brain in a clinical research setting.
- For typical clinical studies, the most relevant sources of "physiological noise" artifacts are subject motion and cardiac pulsation, and the most common sources of "system-related artifacts" are eddy current distortions.
- Artifacts not only increase the variability of diffusion MRI measures but also, most importantly, introduce bias, resulting in potentially false-positive findings.
- Physiological noise artifacts affect the statistical properties of diffusion-derived measures in a regionally heterogeneous manner. Some anatomical regions are severely affected, while others are almost immune.
- For diffusion tensor imaging (DTI), there are several data-processing approaches that can successfully remove effects of artifacts and improve the quality of the computed diffusion metrics.
- Identification and removal of artifactual data points in high *b*-value /HARDI applications is still a largely unmet challenge.
- As technology improves and scanning is performed at higher magnetic fields, the relative importance of thermal noise in determining the quality of diffusion MRI measurements is decreasing, while that of physiological noise is increasing.
- Contrary to what is commonly assumed, if correctly implemented, cardiac gating can be performed in a time-efficient way with no or very little increase in total scan time.
- Practical suggestions for designing optimal acquisition protocols aimed at minimizing the effects of artifacts are provided in this chapter.

18-Jones-Chap-18.indd 303 8/4/2010 2:45:13 PM

This chapter reviews artifacts and confounds that could affect clinical diffusion MRI studies. In diffusion tensor imaging (DTI) (Basser et al., 1994) and in other diffusion MRI processing approaches, including high angular resolution diffusion Imaging (HARDI) methods (Frank, 2002; Tuch et al., 2002; Jansons and Alexander, 2003; Özarslan and Mareci, 2003; Wedeen et al., 2005), diffusion quantities are computed on a voxel-by-voxel basis from a series of diffusion-weighted images (DWIs) acquired with different magnitude and orientation of diffusion sensitization. Artifacts may be originally present in the DWIs and in turn get propagated to diffusion quantities calculated from them. In other cases, the DWIs may be artifact free, but the computed diffusion quantities may contain artifacts originating from errors in the way that diffusion-weighted signals are combined. For example, voxel-wise computation of the diffusion tensor may be corrupted by misregistration of the DWIs or by errors in the calculation of the b-matrix (Basser et al., 1994) value assigned to each data point.

Artifacts in diffusion MRI can be classified according to different criteria: their origin, mechanisms of manifestation, consequences, approaches that can be used for correction, effects on statistical analysis of diffusion data, or practical relevance in clinical studies, among others. A classical distinction is between system-related artifacts (e.g., eddy current effects, thermal radiofrequency noise) and subject-related artifacts (e.g., motion and cardiac pulsation). From the standpoint of statistical analysis, it is important to know if an artifact would increase the variance of computed diffusion metrics, bias their average value, or both. Artifacts originating from different sources may have a common approach for correction; for example, misregistration artifacts originating from eddy currents or subject motion can both be corrected in post-processing. In this chapter, we are not trying to review systematically all possible artifacts occurring in diffusion MRI or to follow a specific classification. Instead, our aim is to focus on artifacts that are more relevant for diffusion imaging studies of the brain in a clinical research setting and to provide practical suggestions for designing an image acquisition protocol and processing pipeline that would minimize the impact of artifacts and confounds that we may encounter in diffusion MRI.

Artifacts Originating from Image Distortion and Subject Motion

In Chapter 17, Andersson and Skare examined in detail the causes of echo-planar imaging (EPI) image distortions originating from static B_0 field inhomogeneities, concomitant fields, and eddy currents and comprehensively reviewed strategies for correction during data collection and post-processing. Here we will examine the consequences of image distortions on diffusion-computed quantities and will review image registration approaches that could be used to remediate the effects of image distortion artifacts in conjunction with correcting for subject motion.

From the standpoint of the effect on EPI images, static B_0 field inhomogeneities, concomitant fields, and eddy currents all result in image distortion, but from the standpoint of the effect on computed diffusion tensor parameters, the consequences of B_0 field inhomogeneities and concomitant fields distortions are quite different from that of eddy currents distortions. B_0 field inhomogeneities and concomitant fields (Bernstein et al., 1998; Zhou et al., 1998) affect DWIs independently of their degree of diffusion sensitization; therefore, they do not produce misregistration within a DWI data set that would affect the voxel-wise computation of diffusion tensor quantities. On the contrary, eddy currents originating from diffusion gradients produce distortions that affect differently images of the DWI data set, resulting in misregistration of different DWIs, which is problematic for DTI computation and, indeed, for any other analysis method that involves multiple DWIs.

EPI Distortions Originating from B_0 Field Inhomogeneities and Concomitant Fields

Figure 18.1 shows the effects that EPI distortions can have on typical EPI acquisitions used in diffusion MRI. Images in Figure 18.1 have identical acquisition parameters except for the direction of phase encoding. Distortions are more evident when the phase-encoding direction is oriented right–left (RL), because the "natural" RL symmetry of the brain is disrupted. However, although less aesthetically disturbing, distortions are also present for images acquired with anterior–posterior (AP) direction. AP is the customary orientation of phase encoding for clinical DTI.

Although EPI distortions originating from B_0 field inhomogeneities and concomitant fields can severely degrade the anatomical accuracy of DWIs, they do not produce major artifacts on diffusion-computed parameters. This is because all DWIs are equally distorted at each voxel location, and signal intensities for all DWI signals used to compute the tensor at that location are scaled by the same factor. This affects the intercept but not the slope of the tensor fitting, resulting in computed diffusivity values that can be accurate even if the anatomy is distorted. Obviously, this does not apply for extreme cases where the distortion is so severe that it results in massive signal accumulation or rarefaction, or where the susceptibility gradients are so strong that they have significant interactions with the diffusion-sensitizing gradients. The fact that DTI-derived metrics are relatively insensitive to the effect of EPI B_0 field inhomogeneities and concomitant field-related distortions could explain why DWIs are not generally corrected for these distortions prior to tensor computation.

It should be noted that in order to correct for EPI distortion, additional data need to be acquired, either additional images for B_0 mapping (Jezzard and Balaban, 1995) or a structural image with proper contrast to be used as a target for registration (see Chapter 17 for alternative methods). Moreover, with few exceptions, software packages for DTI processing neglect to include an EPI-distortion correction step. We believe, however, that there are reasons for implementing EPI-distortion correction in diffusion MRI applications.

First, EPI distortions produce images that are anatomically inaccurate, precluding direct matching of corresponding regions between diffusion and structural images. Second, EPI distortions may vary across different scanners and may vary in the same scanner over time, introducing an undesirable element of variability for longitudinal and multicenter population studies. Third, local EPI distortions are related to the cranial anatomy of each subject (for example, the position, size, and patency of the sinuses), again introducing elements of intersubject variability or even bias in clinical studies. Figure 18.2 presents an example from the work of Wu et al. (2008), showing high variability of fractional anisotropy (FA) in specific regions of the brain for data sets acquired with different phase-encoding directions. This effect is reduced if DWIs are corrected for EPI distortion prior to tensor computation. In this example, two methods have been used for computing the deformation field to be applied for distortion correction: B_0 field mapping (implemented using tools provided by the Oxford

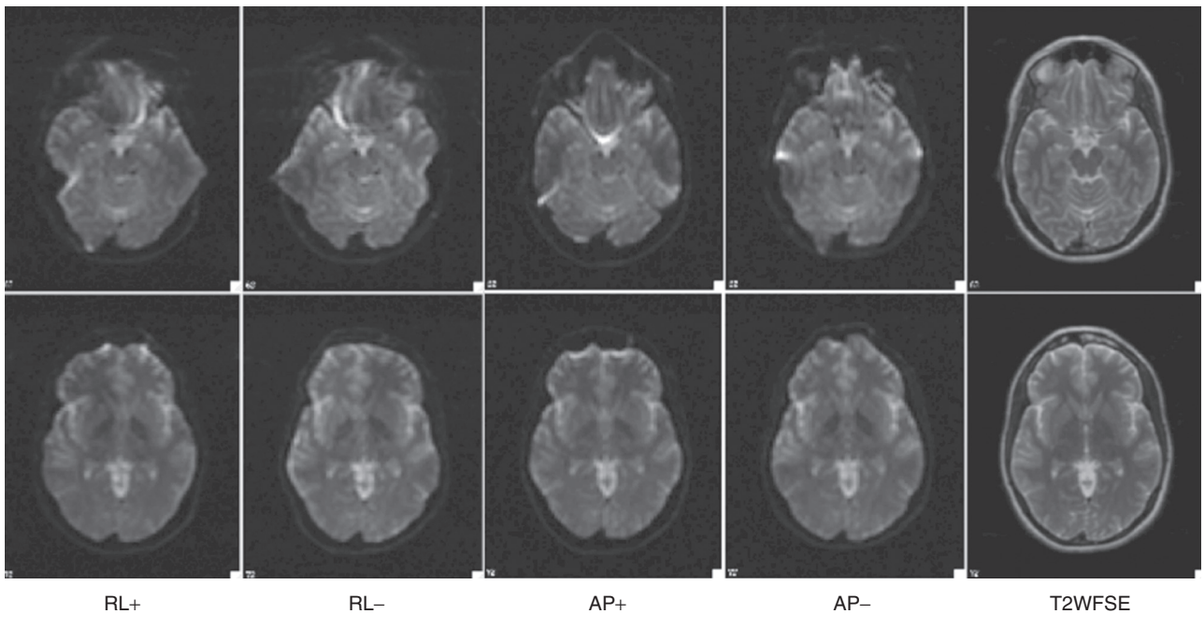

Figure 18.1 Shown are $B = 0$ EPI axial slices at two anatomical levels in the brain acquired with different phase encoding directions. RL indicates phase encoding oriented right–left with either positive (RL^+) or negative (RL–) blips; AP indicates phase encoding oriented anterior–posterior with either positive (AP^+) or negative (AP–) blips. Images on the right column are the corresponding anatomically undistorted structural images acquired with a fast spin echo sequence (T2WFSE). Distortions are more evident for images in the top row where susceptibility variations are more pronounced. However, distortion is also detectable for images acquired in regions of the brain that are relatively immune from susceptibility variations (bottom row) (see Wu et al., 2008).

Center for Functional Magnetic Resonance Imaging of the Brain, FMRIB, FSL, at: http://www.fmrib.ox.ac.uk/fsl) and b-spline registration to a T_2-weighted structural image as target (Wu et al., 2008) (implemented in the TORTOISE software package for diffusion MRI processing, available at: http://tortoisedti.org). The b-spline approach shows better performance at the top of the brain, while correction using B_0 mapping best corrects distortions in the brainstem and in the temporal lobes. These regional differences in performance between the two methods may be related to the ability of the b-spline method to correct for concomitant field (Bernstein et al., 1998; Zhou et al., 1998) distortions, whereas B_0 mapping may be more robust in characterizing deformations in regions with large susceptibility variations. Other groups have also reported successful correction of susceptibility-induced distortions in DWI without B_0 mapping using structural MRIs as target (Huang et al., 2008).

A fourth reason for correcting EPI distortions in diffusion MRI applications is that they could affect diffusion-based tractography results, regardless if performed with probabilistic or deterministic methods and even with HARDI approaches. The effect of EPI distortion on diffusion MRI-based fiber tractography has hardly been investigated (Andersson et al., 2004), but it is reasonable to expect that tractography results would be affected by EPI distortions. EPI distortions do not affect the estimated tensor, so at each location the tensor will reflect the correct orientation of fibers. However, the tensor field from which tract trajectories are computed will be inconsistent with the distorted anatomy. Because of the distorted anatomy, fiber tracts computed from EPI diffusion images in typical clinical scans are likely to suffer from systematic inaccuracy resulting in anatomically incorrect fiber trajectories if EPI distortions are left uncorrected. Figure 18.3 shows an example illustrating this situation. The data are from a typical DTI acquisition at 3T. Fiber tracts in cingulum bundle traced from diffusion tensor data computed from images that have not been corrected for EPI distortion are unable to capture the anatomy of the structure. Tractography results are significantly improved after applying EPI distortion correction to the DWIs prior to tensor computation. It should be noted that the underlying tensor is exactly the same in the two data sets, the only difference is the correction of the distorted anatomy.

Misregistration Artifacts: EPI Distortions Originating from Eddy Currents, and Subject Motion

While the effect of B_0 susceptibility distortions has been somewhat neglected in diffusion MRI, the effects of eddy currents have received much more attention since the early applications of diffusion-weighted EPI in the human brain (Turner et al., 1990). Other chapters in this book address the physical underpinnings of eddy currents and review methods to reduce their impact in diffusion MRI. Here we review the most common manifestations of eddy current–related artifacts in clinical diffusion MRI. Eddy currents can corrupt diffusion measurements via two main mechanisms: 1) contributions to diffusion sensitization that are not accounted for in the calculation of the *b*-value (or b-matrix), and 2) image distortions. While image-distortion correction can be accomplished reliably for clinical data in post-processing, trying to quantify and correct for the extra diffusion sensitization produced by eddy currents is impractical in a clinical setting.

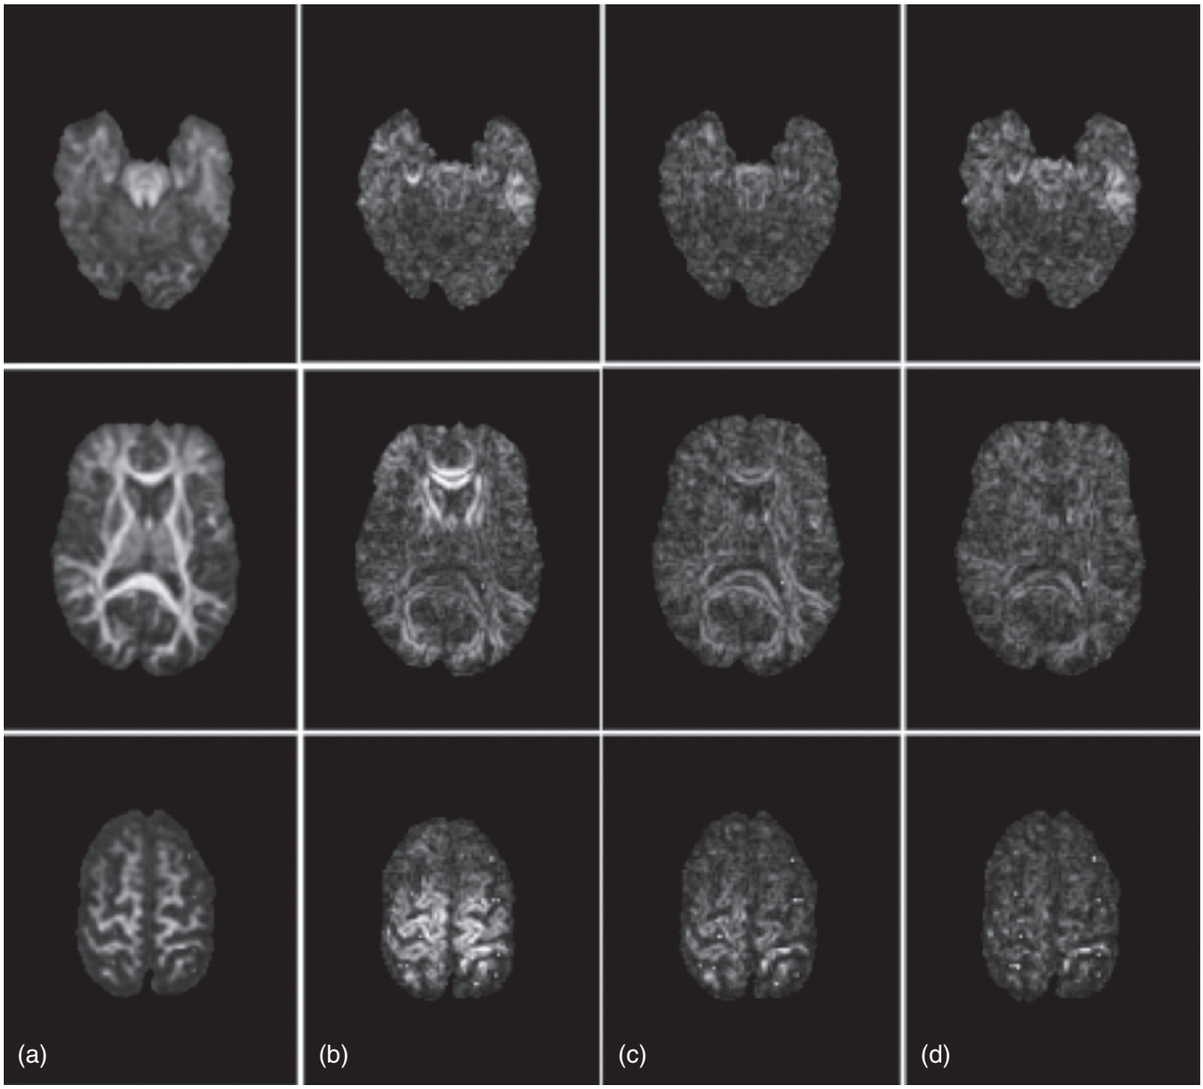

Figure 18.2 **a)** Axial FA maps computed at three levels in the brain of a normal volunteer. **b)** The corresponding standard deviations of FA for four coregistered maps computed in the same subject from data acquired with different directions of phase encoding. Regions in the temporal lobes, genu of the corpus callosum, and at the top of the brain show high variance of FA consequent to EPI-related distortions. The standard deviation of FA computed for the same data sets is reduced when DWIs have been corrected using **c)** a standard B_0 mapping approach or **d)** a b-spline registration to a T_2-weighted structural image as target as described by Wu et al. (2008).

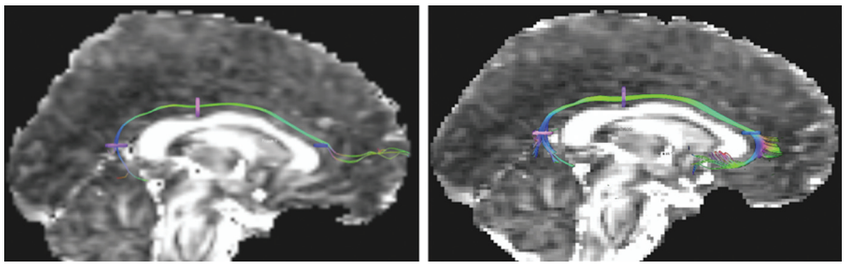

Figure 18.3 Fiber tractography of the cingulum bundle at 3T. In the left panel, DWI data have not been corrected for EPI distortion, whereas in the right panel, the same data have been registered to an undistorted T_2-weighted image with a b-spline volume-based registration algorithm prior to tensor computation. The tracts computed from images that have been distortion-corrected are more consistent with the known anatomy of the cingulum bundle.

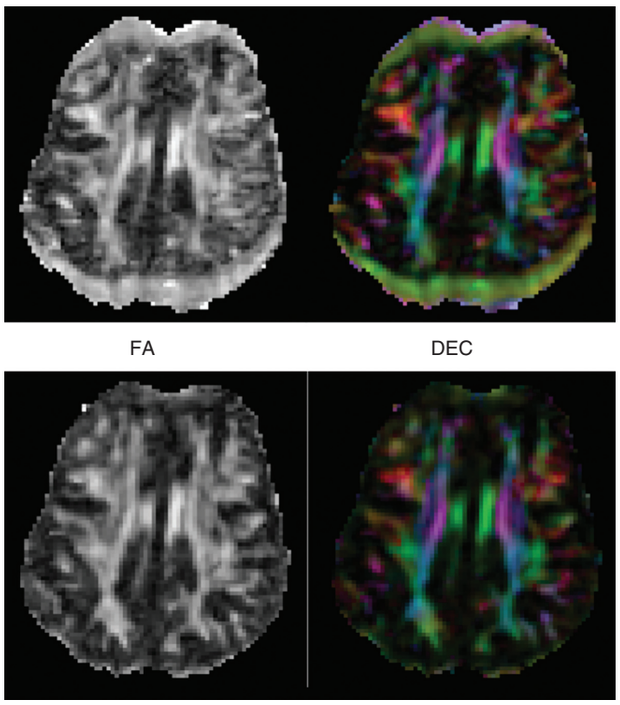

Figure 18.4 Top: Fractional anisotropy (FA) map (left) and direction-encoded color (DEC) map (right) computed from diffusion-weighted images with severe eddy current distortions. Bottom: The same quantities computed after distortion correction.

Eddy currents should be minimized in the acquisition by using good hardware design (see Chapter 11 by Yadav and colleagues), proper calibration of the pre-emphasis of the gradients, and diffusion pulse waveforms that are less likely to produce eddy current artifacts (Reese et al., 2003).

A symptom of eddy current distortion in DTI data is bands of increased anisotropy at the periphery of the brain (Fig. 18.4). Sometimes a small rim (one voxel thick) of increased anisotropy may be present at the periphery of the brain even for data that do not suffer from eddy current distortions or other misregistration artifacts. This is because diffusion anisotropy is artifactually increased at low signal-to-noise ratio (SNR) (Pierpaoli and Basser, 1996) (see Chapter 20 by Pajevic for more details) and, at the periphery of the brain, there are voxels with partial volume with air that have intrinsically low SNR. In the case of eddy current distortions, however, the vector field in the bands of increased anisotropy shows a consistent orientational bias that can be appreciated in direction-encoded color (DEC) maps (Pajevic and Pierpaoli, 1999) (Fig. 18.4). Directional bias would not be present if the high-anisotropy band originated simply from low SNR. Misregistration of DWIs caused by patient motion may also sometimes result in a rim of high anisotropy with a coherent orientation of the vector field.

Artifacts from subject motion, however, can be differentiated from those produced by eddy currents, because eddy current effects are only detectable in the direction of phase encoding (usually anterior-posterior in axial images of the brain), whereas misregistration artifacts due to subject motion will also appear on the sides of the brain. Images of the off-diagonal elements of the diffusion tensor are also very sensitive to misregistration artifacts (Fig. 18.5), so inspection of images of the off-diagonal elements would be a useful additional step in routine quality control of diffusion MRI data. (For more recommendations on quality control the reader is referred to Chapter 19 by Dowell and Tofts, and Chapter 21 by Leemans.)

Subject motion is another source of artifacts in diffusion MRI data. In multi-shot acquisitions, the main consequence of motion is severe ghosting in the images (see Chapter 14 by Bammer and colleagues). Clinical diffusion images, however, are acquired almost exclusively with single-shot EPI sequences, which are much less prone to image degradation from motion than multi-shot acquisitions. With single-shot diffusion MRI, the main consequence of motion is image misregistration, which will corrupt diffusion-derived quantities computed on a voxel-by-voxel basis from different DWIs. Motion-related misregistration artifacts have more heterogeneous consequences than eddy current misregistration artifacts. In some cases they may result in regions of spurious anisotropy at the periphery of the brain that are easy to identify on anisotropy maps or DEC maps. In other cases they may produce a widespread increase or decrease in anisotropy and abnormalities in the measured diffusivity. Figure 18.6 shows an

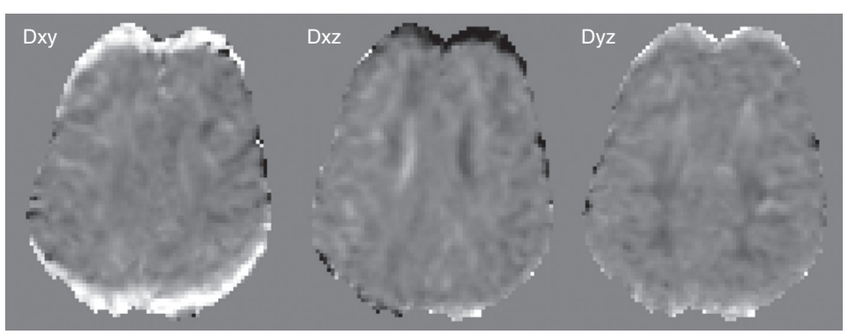

Figure 18.5 Images of the off-diagonal elements of the diffusion tensor for the same data set presented in Figure 18.4 before distortion correction. Values of the off-diagonal elements can be positive or negative and are zero in isotropic regions. In this figure, the gray background is set to zero, positive values are bright, and negative values are dark. Misregistration artifacts can be easily detected in these images as bright or dark regions at the periphery of the brain.

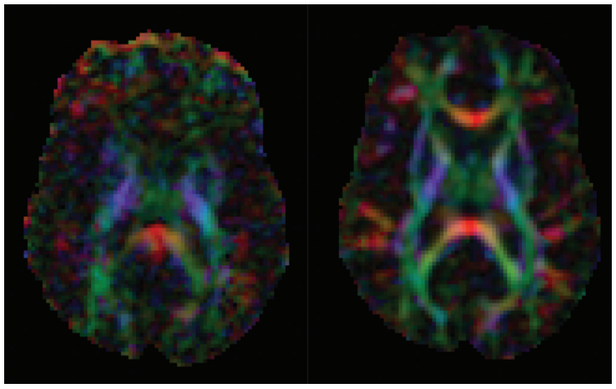

Figure 18.6 Left: Direction-encoded color (DEC) map computed from DWIs acquired on a subject who moved his head repeatedly during the acquisition. Right: DEC map computed from the same data after coregistering the various DWIs' volumes and rotating the b-matrix to account for the rotation of the head in the magnet coordinate system. Data processing and motion correction were performed with the software package TORTOISE (www.tortoisedti.org).

example of a DEC map computed from DWIs that are severely misregistered because of large head-movements. Although the DEC map appears noisy and of poor quality, it is not immediately clear that the data set is corrupted by motion. The effects of subject motion may have different manifestations depending on the type of motion. The effects of an isolated abrupt movement of the head would be different from those of a progressive small drift in the head position during the scan. Moreover, experimental design aspects, such as the order of acquisition (Dubois et al., 2006; Cook et al., 2007) of different directions and *b*-values, will have an impact on the consequence of subject motion (increased variability or biased results).

Possible confounding effects of uncorrected motion on diffusion results should be kept in mind when studying subjects with different levels of compliance during the scan. There may be groups of subjects that, for various reasons, may be more prone to moving during the scan, for example, elderly people, who may be uncomfortable during long scanning sessions; uncooperative patients; and children.

Post-Acquisition Processing Pipeline for Correction of DWIs

Several strategies for correction of eddy current distortions have been proposed in the literature; Andersson and Skare review them in detail in Chapter 17. For clinical diffusion MRI, correction of eddy current distortion using image-based registration algorithms is practical and effective. Software packages for image-based registration of DWIs are becoming readily available, although the affine registration model that is generally employed in most software packages for correction of eddy current distortions is inadequate. Full correction of eddy current distortions should include a quadratic term in addition to the linear terms contained in the affine model (Rohde et al., 2004). Image registration is also commonly employed in diffusion MRI to correct for DWI misregistration caused by subject motion.

Ideally, a comprehensive post-acquisition processing pipeline for clinical DWIs should accomplish the following goals: 1) EPI distortion correction, 2) eddy current distortion correction, 3) subject motion correction, 4) reorientation of images from native space into a standardized space, and, eventually 5) fully deformable spatial standardization into a common target. All transformation and deformations should be applied at once to the original images to ensure minimal degradation of image quality that would occur with multiple interpolation steps.

One complicating factor in dealing with spatial transformation in diffusion MRI is the reorientation of the coordinate reference frame of diffusion sensitization that must accompany some of the rotations applied to the DWIs. For example, if a subject rotates the head in the magnet during the acquisition of the diffusion MRI dataset and the DWIs are coregistered in post-processing, the orientation of the diffusion-sensitizing gradients for each DWI must also be recalculated to account for the rotation applied (Rohde et al., 2004). In diffusion tensor MRI, this is accomplished by rotating the b-matrix associated with each DWI. Neglecting to perform this step can lead to incorrect computation of diffusion tensor metrics, resulting in erroneous tractography findings (Leemans and Jones, 2009).

Not all spatial deformations applied to DWIs should be accompanied by a rotation of the b-matrix. Deformations applied to correct for EPI distortion induced by B_0 field inhomogeneities and concomitant fields, as well as deformations applied to correct for eddy current distortions, should not be accompanied by a rotation of the b-matrix. In general, only rotations that affect the actual object being scanned should be applied to the b-matrix, not rotations consequent to a correction of spatially mismapped MRI signals. Steps 3, 4, and 5 of the processing pipeline mentioned earlier (subject motion correction, reorientation of images from native space into a standardized space, and fully deformable spatial standardization into a common target) would need a reorientation of the b-matrix. Step 5 would actually require a reorientation of the b-matrix independently for each voxel.

Open Challenges in DWI Registration

Early attempts to register DWIs directly to $b = 0$ images were not very successful because the cross-correlation cost function used for the registration could not handle the large difference in contrast between these two types of images (Bastin, 1999; Haselgrove and Moore, 1996). Mutual information-based registration methods are well suited for registering images with very different contrast (Maes et al., 1997; Studholme et al., 1996; Wells et al., 1996), and their introduction into diffusion MRI has allowed correction of eddy current distortion in diffusion-weighted MRI volumes acquired with the diffusion sensitization level ($b = 1000$ s/mm^2) and resolution ($2 \times 2 \times 2$ mm) typically used in clinical studies (Rohde et al., 2004).

Registering DWIs acquired at very high *b*-values, however, is a challenge. Figure 18.7 shows images acquired at *b*-values of 0, 900, and 9000 s/mm^2, all displayed at the same window level. It can be seen that the SNR of the $b = 9000$ s/mm^2 images is very low, the contour of the brain is no longer visible, and the images acquired with different gradient orientations are too dissimilar from each other to expect image-based registration to work. The inability of coregistering images acquired at a very high *b*-value is a practical obstacle to achieving a reliable implementation of high *b*-value HARDI approaches in a clinical setting. Without the possibility of coregistering the DWIs directly, the effects of eddy current distortions could be estimated using separate measurements on phantoms (Horsfield, 1999) and the computed deformation then applied to the DWIs. This approach is not very practical

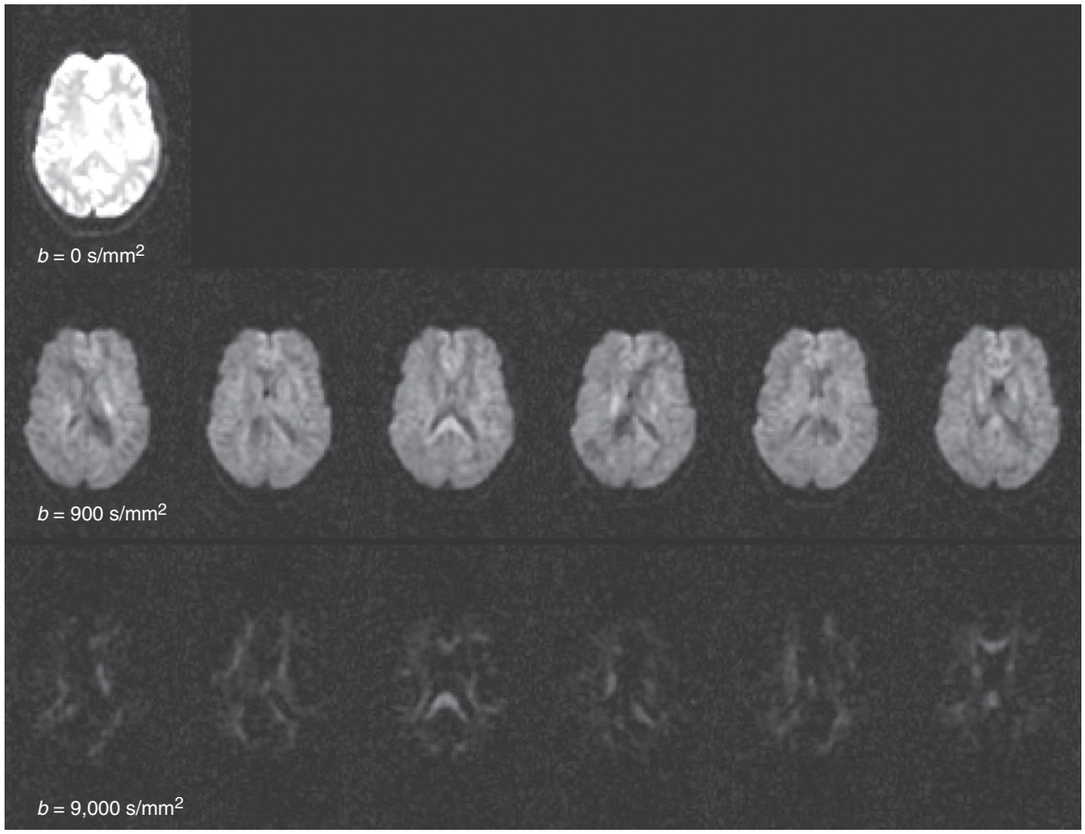

Figure 18.7 Images acquired with b-value = 0 s/mm^2 (top row), b-value = 900 s/mm^2 with six different gradient orientations (middle row), and b-value = 9000 s/mm^2 with the same six gradient orientations (bottom row). All images are displayed at the same window level. Mutual information-based registration algorithms can successfully coregister b = 900 s/mm^2 images among themselves and to the b = 0 s/mm^2 image, but would fail to register the b = 9000 s/mm^2 images.

in a clinical setting because it requires time-consuming initial calibration scans for the various diffusion protocols available on the scanner and periodic monitoring of the reproducibility of eddy current effects. In theory, one should be able to calculate the eddy current for arbitrary directions and b-values from a limited set of experimental measurements. In practice, however, drifts in gradient response due to heating and other factors affecting the performance of the gradients require a larger set of experimental measures.

Even if correction of eddy current distortion could be achieved using external calibration, one would still have to find a suitable approach for correction of subject motion. External camera motion detection systems (Tremblay et al., 2005; Zaitsev et al., 2006) could be employed. Alternatively, a practical trick would be to monitor the position of the organ being scanned by acquiring b = 0 volumes interspersed with the high b-value DWIs and extract the information about subject motion by registering only the b = 0 images. The correction could then be applied to the DWIs, assuming, for example, constant velocity motion between temporally contiguous pairs of b = 0 images. Obviously, this approach increases scan time, because it requires the acquisition of a large number of b = 0 images that do not contribute diffusion information.

All things being considered, available strategies for correcting motion and distortion in high b-value DWIs are suboptimal, and those that are available are rarely applied. Most studies in the literature that employed high b-value acquisitions do not describe the steps taken to correct for subject motion and eddy current distortion. Given that uncorrected eddy current distortion and subject motion are known to introduce severe artifacts in DTI, similar problems can be expected in high b-value HARDI studies, perhaps with the aggravating factors that eddy current distortions and motion are expected to be more severe with the long diffusion gradient pulses and the lengthy acquisition sessions of these studies. Given the growing interest in clinical applications of high b-value HARDI applications, it is hoped that solutions to these issues will be found soon.

Another application that presents unique challenges related to image registration is fetal DTI. Diffusion MRI is emerging as a very promising technique for investigating fetal development and for diagnosing fetal abnormalities (Righini et al., 2003; Jiang et al., 2007; Kasprian et al., 2008; Kim et al., 2008). Acquiring data suitable for voxel-wise computation of the diffusion tensor, however, is challenging given that fetal motion can be large, frequent, and superimposed on respiratory motion of the mother. Conventional fetal MRI acquisitions are usually performed during maternal breathhold, but this approach is unsuitable for DTI given the amount of data to be collected. Image registration algorithms that work well for data collected in adults and pediatric subjects are not designed to handle the complex motion patterns found in

fetal MRI because they generally coregister DWI "volumes" (sets of slices acquired with identical acquisition parameters covering the imaged organ without gaps) and assume that motion during the time necessary to acquire a DWI volume is negligible. In fetal DWI acquisitions, this assumption is almost constantly violated, so alternative strategies need to be developed. Jiang et al. (2007, 2009) proposed a very clever approach in which individual slices can be coregistered independently and the corresponding gradient directions rotated accordingly to account for motion. In the presence of significant motion, the resulting set of data for tensor computation will have different diffusion directions and different amounts of data at each voxel location. However, oversampling will ensure that all voxels will have enough data to enable computing the diffusion tensor. DTI results obtained with this processing approach appear very promising (Jiang et al., 2009).

Physiological Noise

In the early days of functional MRI (fMRI), Fourier analysis of fMRI time series acquired during resting state revealed the presence of signal fluctuations with frequencies consistent with that of the cardiac and respiratory cycles (Weisskoff et al., 1993). Cardiac-related fluctuations were found to be spatially varying with larger magnitude in the cerebrospinal fluid (CSF) and sagittal sinus, followed by gray matter, with lower but still detectable effects in white matter. Since these early findings, the term *physiological noise* has been used in the fMRI community to refer to signal variability related to cardiac and respiratory effects (Kim et al., 1994; Friston et al., 1995). More broadly, the term *physiological noise* could also be used to include all sources of signal variability that are related to scanning a living subject and cannot be explained by thermal noise and system instabilities, for instance, bulk motion effects addressed in the previous section.

Thermal noise is unrelated to the intensity of the true signal whereas physiological noise generally scales with the underlying signal (Kruger et al., 2001). This is an important difference between these two types of sources of signal variability, with relevant practical consequences. The tendency toward scanning human subjects at higher static magnetic fields, from 1.5T to 3T, and recently from 3T to 7T, has been in part driven by the theoretically justified assumption that higher field strengths ensure better SNR of the acquired images. The SNR improvement with higher field, however, applies only to thermal noise, not to physiological noise (Kruger et al., 2001; Triantafyllou et al., 2005). For images acquired at the typical spatial resolution of fMRI and diffusion MRI, physiological noise becomes more relevant than thermal noise as a source of signal variance in EPI images at 3T. The effects of thermal noise in diffusion MRI have been extensively investigated (Pierpaoli and Basser, 1996; Bastin et al., 1998; Skare et al., 2000; Anderson, 2001; Koay et al., 2006; Farrell et al., 2007), whereas physiological noise has not received the same level of attention. It is to be expected, however, that with the transition of clinical diffusion studies to higher field magnets, there will be renewed interest in analyzing the effects of physiological noise not only for fMRI applications but also in diffusion MRI.

Effects of Cardiac Pulsation and Respiration on Brain MRI

Pulsations during the cardiac cycle cause significant motion of the brain parenchyma, with peak velocities of 1–2 mm/s and maximum displacements of about 0.5 mm (Poncelet et al., 1992). The timing of the effects of brain pulsation during the cardiac cycle on the brain was characterized several years ago. Enzmann and Pelc (1992) studied brain motion during the cardiac cycle using a pulse oxymeter on a finger and found that brain motion is relatively small during the first 25% of the cardiac cycle. Poncelet et al. (1992) and Feinberg and Mark (1987) studied the relationship between brain motion and cardiac cycle using an electrocardiogram (EKG) trigger. These studies found that the largest movement of the brainstem occurs 100–200 ms after the R-wave. The apparent discrepancy between the findings of Enzmann, Poncelet, and Feinberg is related to the fact that Enzmann used peripheral triggering whereas the other authors used EKG triggering. The pulse wave takes about 200–300 ms after the R-wave to arrive in the brachial artery (Hasegawa and Rodbard, 1979) and about 250 ms to be detected at the pulse oxymeter on the index finger (Pierpaoli et al., 2003).

Cardiac-induced brain motion is not a rigid-body type motion, as different brain regions have velocity and displacement profiles that differ in magnitude, direction, and time course. Tissue strain is most pronounced in the brainstem and midbrain with an observed craniocaudal elongation. The magnitude of the local strain can be 2 μm/mm, with an intravoxel strain of about 4 μm for isotropic 2 mm voxels (Reese et al., 2002). Brain tissue velocity at the level of the pons has been reported as having a temporal profile with almost a Gaussian shape with values being still close to zero at 100 ms after the R-wave, peaking to 2 mm/s at 200 ms, and going back to zero at 300 ms (Wirestam et al., 1996).

The effect of the respiratory cycle on functional magnetic resonance images has been investigated (Raj et al., 2001), but to our knowledge there are no studies specifically addressing the effects of respiration on diffusion MRI. Reported effects of respiration on EPI include shift of the images in the phase-encoding direction and shading. These effects seem to be caused by spatially constant and spatially varying modulations of the main magnetic field due to changes in the geometry of the chest cavity during respiration (Raj et al., 2001).

Effects of Cardiac Pulsation on Diffusion Tensor–Derived Quantities

Cardiac pulsation was identified as a source of artifacts in single-shot EPI-based clinical diffusion MRI studies about 20 years ago (Chien et al., 1990). Typical manifestations were severe signal dropouts in the periventricular areas in the center of the brain and thalamic regions (Chien et al., 1990). Subsequent studies have confirmed that the main effect of cardiac pulsation in the DWIs is exaggerated signal loss in certain regions (Skare and Andersson, 2001; Nunes et al., 2005). Signal dropouts affect mainly images acquired at the peak of the systolic phase (100–200 ms after the R-wave) (Pierpaoli et al., 2003) and are most pronounced in regions known to have large velocity variations during the cardiac cycle. While the exact origin of the signal loss is not known, it may be due to intravoxel pseudo-diffusion effects related to local strain of the tissue (Pierpaoli et al., 2003) occurring in the systolic phase. Other authors have pointed out the importance of image reconstruction methods in making DWIs prone to cardiac pulsation artifacts. According to Robson and Porter (2005), Margosian reconstruction is more prone than zero padding in accentuating the artifacts due to pulsatile motion.

Cardiac pulsation affects eigenvalues and eigenvectors and other quantities derived from the diffusion tensor, such as Trace(D) and FA, via two main mechanisms. The first is the

corruption of DWIs by signal dropouts, and the second is local misregistration from the brain moving during the cardiac cycle. Given that the maximum displacement of the brain parenchyma due to brain pulsation is reported be no larger than 0.5 mm, local misregistration is likely to be most pronounced at boundaries of regions having different diffusion properties.

Figure 18.8 shows regions of the brain in which Trace(D) was found to be significantly affected by the cardiac cycle in a healthy volunteer. This result is representative of findings in five healthy young subjects. Regions where trace(D) is significantly affected by cardiac pulsation include the ventricles and other CSF-containing spaces, the brainstem, medial portions of the cerebellum, hypothalamic and thalamic regions, the medial portion of the genu of the corpus callosum, and cortical regions at the interface between tissue and CSF. Generally, the affected regions appear to involve 1) areas that are known to be subjected to significant pulsatility during the cardiac cycle with consequent tissue strain; 2) cavities containing CSF, which is likely to present convection and pseudo-diffusion effects; and 3) regions at the periphery of the brain that are not known to be highly mobile during the cardiac cycle but in which there is a sharp difference in diffusivity properties (e.g., tissue–CSF interfaces). In regions

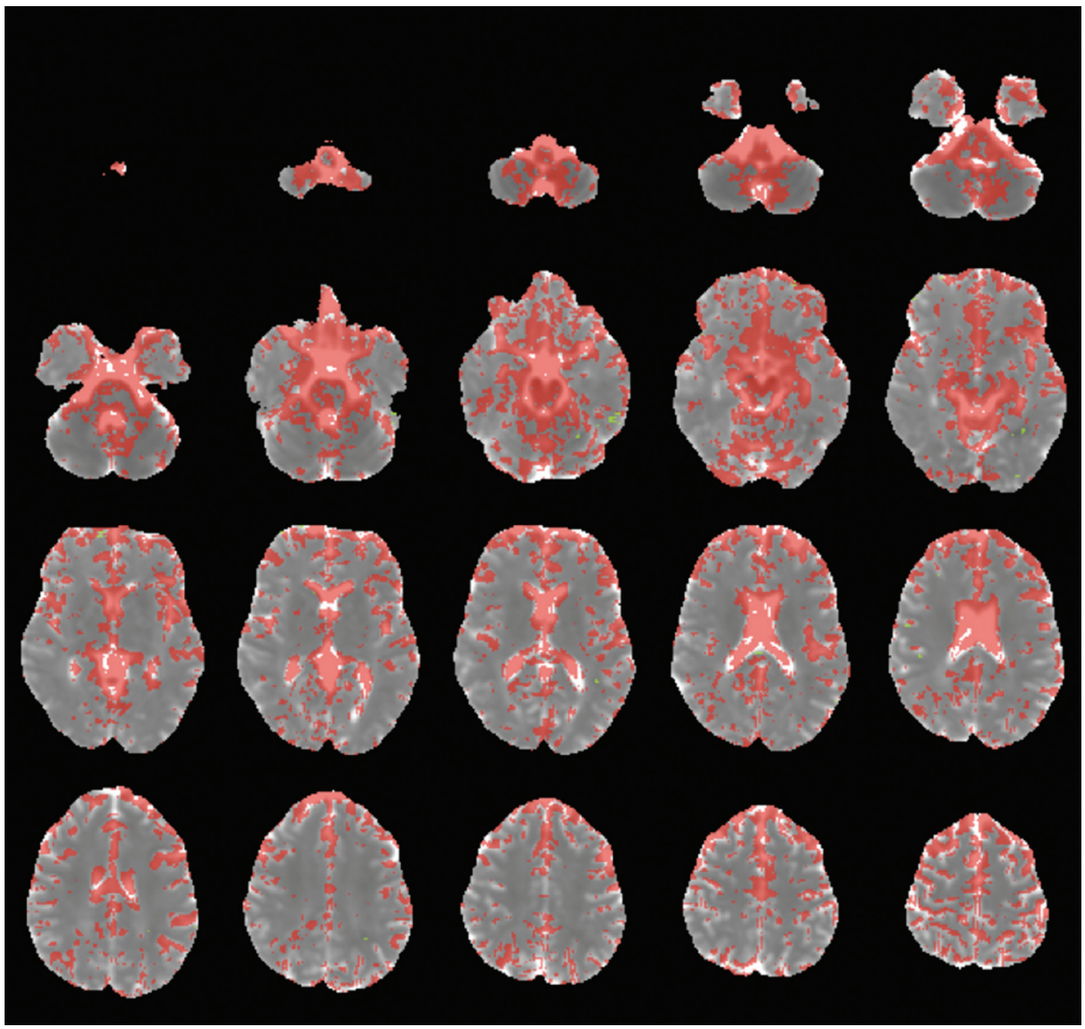

Figure 18.8 Axial images of mean diffusivity <D> (1/3 trace(D)) with superimposed in red areas where <D> is most likely to be affected by cardiac pulsation. Data for this analysis are from 81 whole-brain DT-MRI data sets acquired in a healthy volunteer with nine different trigger delays stating at 40 ms after the onset of the R-wave with 70-ms increments (i.e., 40, 110, 180, 250, 320, 390, 460, 530, and 600 ms after the R-wave). For each trigger delay nine replicates were acquired. Each replicate consisted of one image with no diffusion sensitization and six DWIs acquired with $b = 1100$ s/mm^2 and the gradient direction scheme proposed in *Radiology* 201:637–648, 1996. The effect of cardiac pulsation was tested by grouping data by trigger-delay values and performing a Kruskal-Wallis test, which is a nonparametric analog of the analysis of variance. In red are reported voxels that are significant at $p < 0.0001$ uncorrected for multiple comparisons. In order to test for the presence of false-positive results at this significance level, the analysis was repeated by grouping data by replicate number (i.e. swapping rows and columns in the design matrix). Voxels significant at $p < 0.0001$ for this analysis are in green. Hardly any green voxels are visible, indicating that the red voxels in the trigger delay analysis are unlikely to be false-positive results.

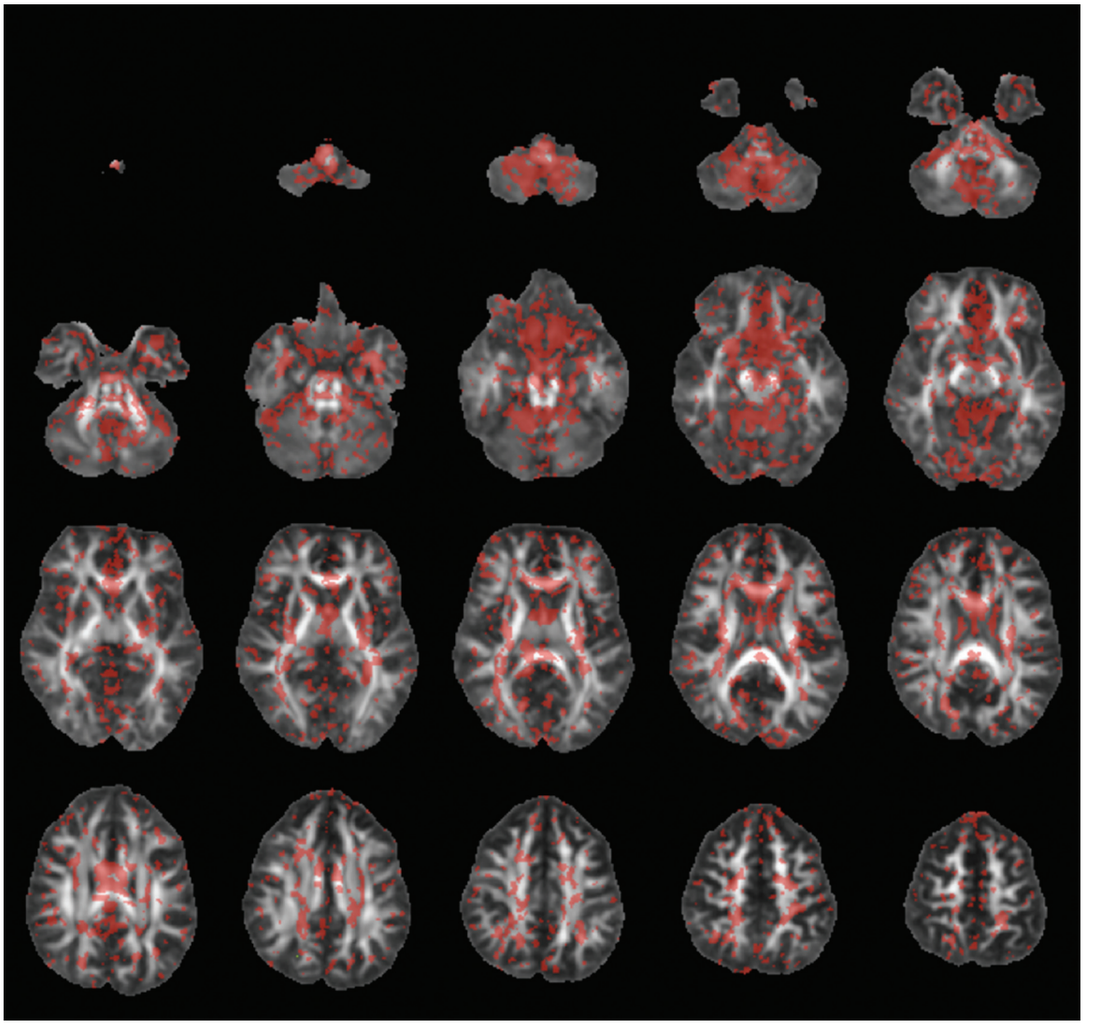

Figure 18.9 Axial images of fractional anisotropy (FA) with superimposed in red areas where FA is most likely to be affected by cardiac pulsation. Data are from the same study reported in Figure 18.8.

with sharp differences in diffusivity, the effects of a small tissue displacement could be magnified. Figure 18.9 shows the same analysis for FA. It is interesting that some of the affected regions overlap with those found for trace(D), while others do not. For example, the cortical boundaries of tissue–CSF do not appear to be affected by cardiac pulsation in the FA analysis, but the interfaces between gray matter and white matter show a significant effect.

Artifactual changes in diffusivity during the cardiac cycle also affect the apparent orientation of fibers (Pierpaoli et al., 2003; Jones and Pierpaoli, 2005). Figure 18.10 shows DEC maps (Pajevic and Pierpaoli, 1999) computed from DWIs collected at different delays from the R-wave. Data collected during the peak systolic phase show spurious anisotropic structures with inferior–superior orientation (blue hue) in the cerebellum. The genu of corpus callosum, by contrast, shows an artifactual decrease in anisotropy at the midline that would result in a reduction of the measured transcallosal "connectivity" (Jones and Pierpaoli, 2005).

It is easy to understand the potential confounding effects of these artifacts when comparing diffusion results between subjects with different heart rates. At a faster heart rate, the diastolic period is shortened and the likelihood of collecting data during systole is increased; therefore, the effects of cardiac pulsation artifacts will be more prominent. For example, in a hypothetical diffusion MRI experiment in which one found that presentation of stressful stimuli in the scanner results in lower transcallosal connectivity and increased cerebellar connectivity, a possible artifactual origin of this finding (increased heart rate during stress) should be considered before attributing it to brain plasticity. The most prominent cardiac-related artifacts occur in systole; however, there is evidence that some tracts, such as the fornix may show larger variability due to cardiac pulsation in diastole (Jones and Pierpaoli, 2005).

Few studies have investigated the potential effects of physiological noise on the statistical analysis of diffusion tensor MRI results in populations (Peterson et al., 2008; Walker et al., 2008, 2009). All studies that have approached this task have compared

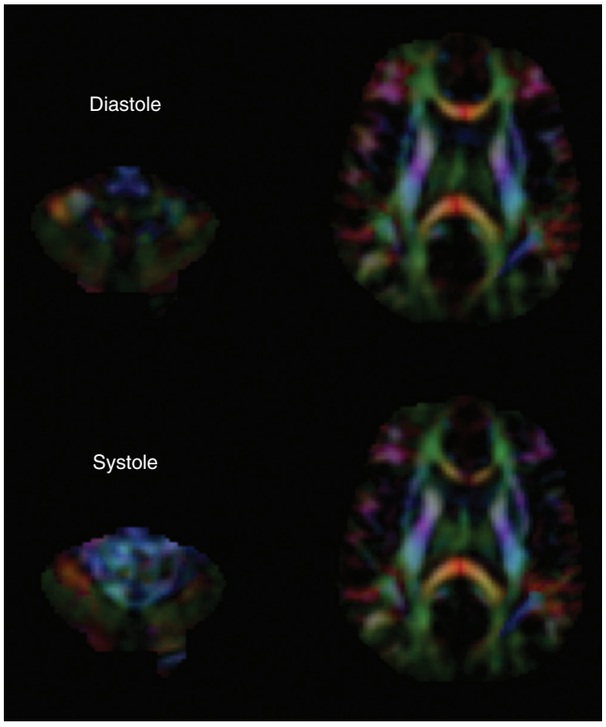

Figure 18.10 Directional-encoded color (DEC) maps computed from DWIs acquired at 420 ms (diastole) and 120 ms (systole) after the R-wave. On the left, the slice is at an anatomical level just inferior to the pons, and on the right, at the level of the thalamus. The appearance of spurious anisotropic structures with inferior–superior orientation (blue hue) can be seen in systole in the cerebellum. The genu of corpus callosum, by contrast, shows an artifactual decrease in anisotropy during systole.

results obtained from data collected without cardiac gating and processed with both conventional tensor fitting and robust tensor fitting using the RESTORE (robust estimation of tensors by outlier rejection) algorithm (Chang et al., 2005). RESTORE allows artifactual data points to be identified as outliers and removed on a voxel-by-voxel basis from contributing to the estimation of the tensor (see next section, Robust Tensor Fitting and Cardiac Gating). It has been shown that outlier data points are regionally consistent across a population and are heterogeneously distributed throughout the brain (Peterson et al., 2008; Walker et al., 2008). Outliers are more densely represented in regions known to be affected by cardiac pulsation, such as brainstem, cerebellum, CSF spaces, medial portions of the diencephalon, and the genu of the corpus callosum (see Fig. 18.11). In addition, a large concentration of outliers was found in regions not been previously reported as being strongly affected by cardiac pulsation, such as ventral portions of the temporal lobes, the splenium of the corpus callosum, fornix, and insular regions (Walker et al., 2008).

There is evidence that the presence of outliers due to physiological noise affects the statistical analysis of diffusion results in a population (Peterson et al., 2008; Walker et al., 2009). However, these initial findings need to be confirmed and expanded. The presence of outliers could bias the value of diffusion tensor–derived metrics, causing their systematic increase or decrease, but it could also increase their variance, modulating the statistical power in a heterogeneous way across the brain. Also, different tensor-derived quantities may be differently affected by cardiac pulsation, as we have seen in the examples reported for single-subject analysis in Figures 18.8 and 18.9.

Robust Tensor Fitting and Cardiac Gating

There are three relatively simple approaches for reducing the likelihood that cardiac pulsation artifact may corrupt diffusion tensor MRI measures: 1) robust tensor estimation and 2) retrospective cardiac gating for ungated acquisitions, and 3) cardiac gating during acquisition.

In conventional nonlinear tensor fitting, each data point has identical weight (Basser et al., 1994). This is a correct model assuming that only thermal noise contributes to the measurement error. The use of robust tensor estimation methods (Mangin et al., 2002; Chang et al., 2005) has been proposed to account for DWIs possibly being affected by artifacts that violate the assumption of constant errors for all data points. Robust fitting methods assume that artifactual data points are less numerous than good data points in a data set and use the residuals of the fitting to identify artifactual data as outliers and reduce their contribution to the estimation of the fitted parameters. The classical Geman-McClure M-estimator approach, applied to diffusion MRI by Mangin et al. (2002), uses an iterative reweighting procedure in which experimental data with large residuals are progressively weighted less in the fitting. The approach proposed by Chang et al. (2005), called RESTORE, uses iterative reweighting to identify outliers, but the exclusion criteria and the final fitting take into account the expected error due to thermal noise.

One advantage of robust tensor estimation methods is that they are effective in reducing the contribution of all types of sporadic artifacts, not just those originating from cardiac pulsation. For example, if a few images in a data set are corrupted by spike noise (Chavez et al., 2009) or by severe bulk motion, their contribution will be reduced by robust estimation without requiring manual exclusion. This makes robust estimation very appealing for automated data-processing pipelines. Another advantage of robust fitting is its ability to exclude artifacts on a voxel-by-voxel basis. If an image is only partially corrupted, robust tensor fitting will exclude voxels in regions affected by the artifact but will still make use of the information in the voxels that do not contain artifactual signals. Drawbacks of robust tensor fitting approaches are the increased computational burden and the requirement of redundancy in the acquired data. Moreover, robust fitting by definition requires a model to fit the data to, so it may not be directly applicable for artifact remediation if the data set is not suitable for DTI analysis, such as HARDI acquisitions at very high *b*-values.

Retrospective gating is another appealing method for removing potentially artifactual data points caused by cardiac pulsation. It requires synchronized recording of the EKG or the peripheral pulse waveform during image acquisition, but not direct gating of the acquisition itself. Images being acquired during "bad" periods of the cardiac cycle (e.g., 100–300 ms after the R-wave) could be discarded. It is conceivable to perform this analysis in real time during the scanning session and repeat the acquisitions that had to be discarded. Compared to robust tensor fitting, retrospective gating has the potential disadvantage of not being able to correct for non–cardiac-related artifacts, but it would not require the same amount of data redundancy to be effective.

Finally, direct cardiac gating should be considered. The use of cardiac gating was recommended to reduce the incidence of cardiac pulsation artifacts since the early clinical diffusion MRI

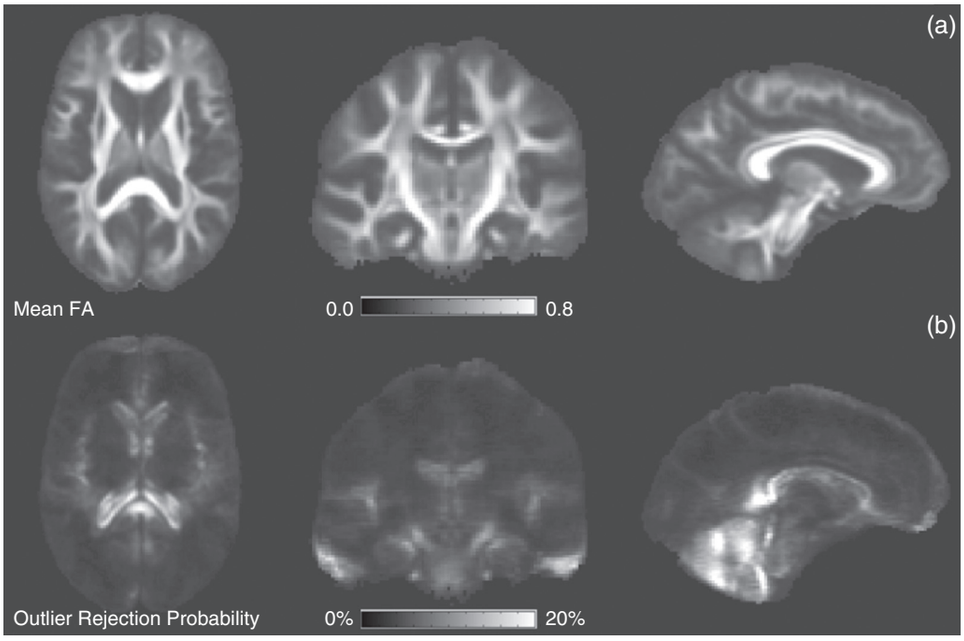

Figure 18.11 **a)** Average FA map computed from diffusion tensor data from 20 healthy subjects, obtained by fully deformable tensor-based registration of the FA maps of each subject to one brain chosen as a template. **b)** Corresponding outlier rejection probability map computed by averaging the percentage of outliers data points identified by RESTORE fitting in each subject. Note the regional consistency of the spatial distribution of outlier data points within the population See Walker et al. Proc. Intl. Soc. Mag. Reson. Med. 2008, p. 139. with permission from John Wiley & Sons, Inc.

studies (Chien et al., 1990; Turner et al., 1990), and there is now a large body of evidence that cardiac gating is advantageous in reducing effects of cardiac pulsation in single-shot EPI diffusion measurements (Skare and Andersson, 2001; Pierpaoli et al., 2003; Jones and Pierpaoli, 2005; Nunes et al., 2005; Gui et al., 2008). However, cardiac gating is not used in most clinical diffusion studies. The reason most widely reported in the literature for not using cardiac gating is that gating increases scan time. While it may be true that time-efficient gating for diffusion MRI may not be currently available on commercial scanners, the time penalty for gating can be minimal if gating is correctly implemented. Ideally, it would be desirable to be able to trigger the acquisition of an arbitrary number of slices after each trigger with an arbitrary intersequence delay between consecutive slices and with a new trigger at each cardiac cycle. The trigger delay should be set equal to zero for peripheral gating and the intersequence delay between subsequent acquisitions should be set to the shortest possible value for time efficiency (i.e., zero, if decaying eddy currents from the previous acquisitions do not represent a problem). With EKG gating, the trigger delay of the first acquisition should be set to zero, then it should follow a pause of about 200 ms, and finally, data could continue to be collected during diastole. In this way, the acquisition could be very efficient and optimized for different heart rates.

There is some evidence that cardiac-related variability is not completely eliminated even for data collected in diastole. The period toward end diastole shows the lowest variance. One study found 500 ms after the peripheral trigger to be the optimal delay, although all data acquired in a time window from 18 ms to 500 ms had lower variance than ungated acquisitions (Skare and Andersson, 2001). However, there is the risk of ending up sampling in the next cardiac cycle if the delay is too prolonged. Nunes et al. (2005), for example, sampled cardiac delays of 0,200,400, 600,800 ms from peripheral trigger and found increased signal variance for data acquired with 600 and 800 ms delays, probably because these delayed acquisitions were collected during the systolic period of the next cardiac cycle.

To avoid the worst period in cardiac cycle, one has to exclude not more than 200 ms per heart beat, which represent 20% of the total available time with a heart rate of 60 beats/min, leaving 80% of the time available for data collection. For most clinical magnets, the acquisition speed of diffusion-weighted images is limited by the duty cycle of the gradients to about 30%–40% of total time, so it is evident that cardiac gating could be accomplished with no time penalty. Newer clinical magnets are approaching a 100% effective duty cycle for DWI data collection. On these magnets cardiac gating will result in an acquisition speed penalty but, if correctly implemented, the time penalty of cardiac gating would be limited to about 25% longer scan time.

Strategies for Designing Optimal Acquisition Protocols Accounting for Physiological Noise

In Chapter 15, Jones considered optimal acquisition strategies from an MR physics, mathematical, and theoretical background, yielding some recommendations for optimal designs. However, when using diffusion MRI in a clinical setting, there are further considerations to be taken into account. The first step in defining a diffusion acquisition protocol is to consider the biological questions we want to ask and to ensure that our experiment design is adequate for the task. Clearly, a protocol appropriate for clinical diagnosis of acute stroke will be different from the protocol one would need to implement to investigate the hypothesis that mental activity produces subtle changes in white matter architecture.

In general, with MRI we need to find a reasonable compromise between data quality and duration of the scan. Diffusion MRI, however, is more demanding than conventional MRI because the number of parameters we need to investigate is larger. One should avoid trying to answer very ambitious questions with very short diffusion MRI studies that do not allow an adequate experimental design.

Once I know how much time can be assigned to the diffusion scan, how do I determine the total number of images, resolution, and slice thickness I can use?

The first information needed is the number of DWI slices that the scanner is able to acquire per unit time. At the time of this writing, this number varies from about 3 to 8 slices per second for clinical scanners of different manufacturers. There are no large differences between 1.5T and 3T units. For some magnets, the main speed-limiting factor for DWI acquisition is the duty cycle of the gradients. As technology improves, gradient duty cycle is improving and is becoming less relevant in limiting acquisition speed. Another factor that could limit the rate of data acquisition is the actual time needed to play out the sequence. Currently, it takes about 100 ms to play out the gradient waveform to acquire a DWI slice, so even with no duty cycle limitation we could not acquire more than 10 slices per second. These figures are likely to improve in the future. In EPI DWIs, a good part of the gradient waveform is occupied by the diffusion gradients and the EPI readout gradients. Parallel imaging has already enabled a shorter EPI readout, and stronger gradient will enable shorter diffusion encoding times.

From the total imaging time for the diffusion scan and the number of DWI slices that the scanner can acquire per unit time one can calculate the total number of slices that can be collected during a scanning session. Experimental parameters such as the total number of gradient directions, the slice thickness, and the desired coverage need to be negotiated within the constraint of total number of slices that can be acquired in the scanning session.

The number of diffusion directions and *b*-values one should acquire are again related to the type of analysis one would like to perform. For diffusion tensor analysis, the minimum number of directions is 6, plus a non–diffusion-weighted image, although a larger number is recommended to improve SNR, ensure rotational invariance of the statistical properties of the computed quantities (Jones, 2004) (see Chapter 15 for more details), and enable artifact removal using robust tensor fitting methods. Previous work has indicated that 20–30 directions or more will be necessary to ensure rotational invariance of the most commonly used tensor quantities, such as trace, fractional anisotropy, and eigenvectors (Jones, 2004). Although redundant sampling is necessary for robust methods, the degree of data redundancy for optimal performance of robust tensor fitting methods is unknown.

In theory, robust methods require the "good data" to be more than 50% of the total amount of data. Given that artifactual images are usually a small fraction of the total number of images collected, it would seem that a small degree of data redundancy would suffice to enable robust tensor estimation with methods like RESTORE. However, it is easy to understand that if a small number of artifactual images happen to be sampling a similar diffusion direction, their removal from the tensor calculation will result in a very unstable or even ill-conditioned diffusion-sampling scheme (Chang et al., 2009). This aspect needs to be further investigated; however, based on our experience, we are reluctant to use RESTORE on data sets containing less than about 8 times the minimum number of images required for DTI (i.e., 48 high *b*-value volumes).

Another aspect that requires additional research is the differential effect of physiological noise on various diffusion directions and *b*-values. Cardiac pulsation artifacts are more prominent for diffusion sensitization in the caudorostral direction. Although gradient optimization approaches based on theoretical consideration and Monte Carlo simulation have consistently recommended homogeneously sparse diffusion-sampling directions, it would make sense to oversample directions that are most likely to yield images corrupted by artifacts. A similar reasoning would also apply for the optimal ratio of low *b*-value versus high *b*-value images. Previous theoretical work reported an optimal ratio of about one low-*b* image to six high-*b* images (see Chapter 15). However, low *b*-value images may be affected by physiological noise factors that are not present in heavily diffusion-weighted images. For example, the effect of signal fluctuations due to spins in the vasculature are present at low *b*-values but not at high *b*-values. More experimental work is needed, but evidences suggest that brain images acquired without diffusion sensitization ($b = 0$) have higher signal variance than that of heavily diffusion-weighted images (Chang et al., 2009). In this case, it would be advisable to increase the ratio of low-*b* images versus high-*b* images to a higher level than what theoretical considerations would recommend.

In terms of coverage, in general, for brain studies we strongly recommend covering the whole brain, from the foramen magnum to the apex of the head, adding a few extra slices at the top to include the skull. Having volumes containing the whole head is very important for correction of motion and distortion during post-processing using image registration methods. Full brain coverage is also desirable because it makes the resulting data more suitable for inclusion into brain databases and enables registration to brain atlases.

With the total number of slices dictated by the scan time, the number of diffusion directions defined by experimental design considerations, and the requirement of full brain coverage, one can determine the slice thickness to use.

Slice thickness in turn determines in plane resolution. The criterion here is simple: one should aim at isotropic resolution, i.e., cubic voxels. In radiology there is a long-standing tradition of acquiring thick slices with very high in-plane resolution. In part, this approach is motivated by the fact that acquiring thinner slices increases the scan time, while the time penalty for achieving high in-plane resolution is less severe. Thick slices with high in-plane resolution give images of apparent very good quality, and radiologists have gotten accustomed to extracting diagnostic information from these types of images in their clinical practice.

At first glance, the time-saving approach of acquiring thick slices would benefit EPI DWI. Actually, in EPI there is essentially no time penalty to increasing in-plane resolution (although image quality may suffer because of the longer readout), so cutting down the total scan time by increasing the slice slickness is tempting.

Why should we resist this temptation? A fundamental issue with diffusion MRI is that the measured diffusion properties in each voxel are dependent on the architectural features of the tissue contained in the voxel. The effect of white matter architecture on measured diffusion properties, in particular on diffusion anisotropy, has been reported since the early DTI studies. In fact,

it can be argued that the underlying architecture of white matter is the most significant factor accounting for regional difference in anisotropy in the healthy adult human brain (Pierpaoli et al., 1996). If white matter structures included in a voxel have inconsistent orientation, the measured diffusion anisotropy is low, whereas if fibers are organized in bundles with parallel orientation, the measured diffusion anisotropy is much higher. Intravoxel orientational coherence affects the shape of the diffusion displacement profile and the apparent orientation of its principal eigenvector and, therefore, has an impact on DTI tractography. In the brain, white matter pathways with different orientation often run adjacent to each other. Examples include the cingulum bundle and the corpus callosum, the transverse pontine fibers and the motor and sensory projection pathways in the pons, and association bundles and projection fibers in the centrum semiovale. Image resolution will affect the apparent local organization of fibers within each voxel because with large voxels, individual bundles of fibers having different orientation are more likely to be included in the same voxel, resulting in a more isotropic diffusion displacement profile. Although the dependence on image resolution is unavoidable (a brain scanned at low resolution will have anisotropy value significantly lower than those of the same brain scanned at high resolution), if we use cubic voxels we ensure that the averaging effect is not biased toward a particular direction. Figure 18.12 further illustrates this concept by showing anisotropy maps computed form a high-resolution DWI data set resampled at a lower resolution to simulate data acquired at low resolution with either cubic voxels or with high in-plane resolution and thick slices. The FA map computed from the DWI data set resampled with high in-plane resolution and thick slices seems of higher quality than the FA map computed from the DWI data set resampled with cubic voxels; it is less pixilated and seems to have higher spatial resolution. It is also evident, however, that the apparent higher pictorial quality is accompanied by a dramatic artifactual decrease in measured diffusion anisotropy and severe directional bias.

Several diffusion analysis methods that aim at extracting the contributions to the diffusion displacement profile of multiple fiber populations in a voxel have been proposed. They are known

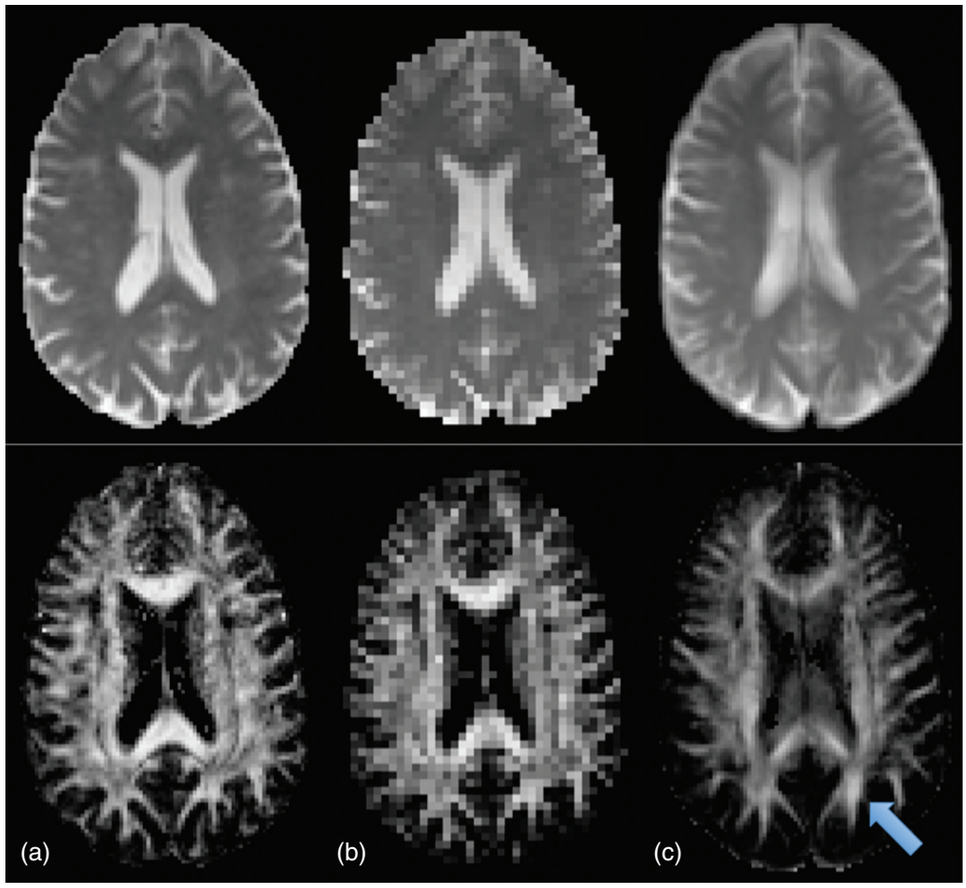

Figure 18.12 Effects of image resolution on measured anisotropy. **a)** Amplitude image ($b = 0$) (top), and the corresponding FA map (bottom) computed from a DWI data set acquired at a resolution of $1.7 \times 1.7 \times 1.7$ mm (voxel volume = 4.9 mm^3). **b)** The DWIs were downsampled prior to tensor computation to a resolution of $3.4 \times 3.4 \times 3.4$ mm (voxel volume = 39.3 mm^3). **c)** The in-plane resolution was maintained at 1.7 mm, but the slice thickness was increased to 13.6 mm in order to obtain the same voxel volume of 39.3 mm^3 as in b). The FA maps in a), b), and c) are displayed with the same window level. It can be seen that FA maps in b) and c) have lower values than in a). Moreover, in c) anisotropy is affected in an heterogeneous way, appearing dramatically reduced in some areas, such as the corpus callosum, and relatively preserved in other areas, such as the subcortical posterior white matter region, indicated by the arrow.

as high angular resolution diffusion imaging (HARDI) methods (see Chapter 27 by Alexander and Seunarine and Chapter 28 by Tournier). HARDI methods should in theory be less susceptible to intravoxel orientational averaging effects. Unfortunately, HARDI methods require a very large number of sampled diffusion directions and high *b*-value images that consequently have much lower SNR than that of DTI. Therefore, for a given scan duration, images used for HARDI analysis are usually acquired with lower spatial resolution than are images acquired for DTI analysis. The requirement of using larger voxels for HARDI analysis, however, to some extent offsets the theoretical advantage of HARDI acquisitions with respects to problems caused by intravoxel orientational averaging. There is no free lunch in diffusion MRI!.

In summary, we would give the following practical suggestions for designing a good acquisition protocol for brain DTI scanning:

1. Use isotropic resolution (cubic voxels).
2. Cover the entire organ.
3. Oversample, i.e., acquire a larger number of volumes than suggested on the basis of theoretical considerations. We would recommend at least 50 volumes with high-diffusion sensitization (*b* about 1100 s/mm^2) to enable using robust tensor estimation methods, such as RESTORE, for removing artifacts originating from cardiac pulsation.
4. Acquire an adequate number of low *b*-value images, with a ratio of at least 1 low *b*-value volume to every 5–6 high *b*-value volumes.
5. Use the highest angular resolution possible (i.e., do not acquire replicates of the same direction).
6. Use optimal sampling schemes (see Jones, 2004, and Chapter 15) as a starting point in determining the set of gradient orientations to be sampled; however, consider oversampling particular directions that appear to be most affected by physiological noise or other artifacts.
7. If replicates of the same direction are acquired, do not average them at the scanner (i.e., use NEX = 1), because doing so would preclude independent correction of misregistration caused by motion in post-processing and the removal of replicates corrupted by artifacts using robust tensor estimation methods.
8. Use optimal gradient orientation ordering (see, for example, Dubois et al., 2006; Cook et al., 2007) to ensure that usable data are available if the scan needs to be cut short. Optimal ordering reduces the risk of systematic errors due to gradient heating and subject motion.
9. Do not acquire all low *b*-value images (or $b = 0$ images) at the beginning or at the end of the scan, but rather intersperse them with the high *b*-value acquisitions. This is done for reasons identical to those of the previous point, avoiding systematic bias as the scanner performance may change during acquisition.
10. If images are too noisy to be coregistered and corrected for motion and distortion in post-processing, consider collecting them at lower resolution.

Acknowledgments

We would like to thank Alan Barnett, Derek Jones, Okan Irfanoglu, Stefano Marenco, Joelle Sarlls, Lindsay Walker, and Minjie Wu, who have been actively involved in several projects relevant to the topics addressed in this chapter and have provided material for figures and text.

References

Anderson AW (2001). Theoretical analysis of the effects of noise on diffusion tensor imaging. *Magn Reson Med* 46:1174–1188.

Andersson JL, Richter M, Richter W, Skare S, Nunes RG, Robson MD, Behrens TE (2004). Effects of susceptibility distortions on tractography. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 11:87.

Basser PJ, Mattiello J, LeBihan D (1994). MR diffusion tensor spectroscopy and imaging. *Biophys J* 66:259–267.

Bastin ME (1999). Correction of eddy current-induced artefacts in diffusion tensor imaging using iterative cross-correlation. *Magn Reson Imaging* 17:1011–1024.

Bastin ME, Armitage PA, Marshall I (1998). A theoretical study of the effect of experimental noise on the measurement of anisotropy in diffusion imaging. *Magn Reson Imaging* 16:773–785.

Bernstein MA, Zhou XJ, Polzin JA, King KF, Ganin A, Pelc NJ, Glover GH (1998). Concomitant gradient terms in phase contrast MR: analysis and correction. *Magn Reson Med* 39:300–308.

Chang LC, Jones DK, Pierpaoli C (2005). RESTORE: robust estimation of tensors by outlier rejection. *Magn Reson Med* 53:1088–1095.

Chang LC, Walker L, Pierpaoli C (2009). Making the robust tensor estimation approach: "RESTORE" more robust. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 17:3557.

Chavez S, Storey P, Graham SJ (2009). Robust correction of spike noise: application to diffusion tensor imaging. *Magn Reson Med* 62:510–519.

Chien D, Buxton RB, Kwong KK, Brady TJ, Rosen BR (1990). MR diffusion imaging of the human brain. *J Comput Assist Tomogr* 14:514–520.

Cook PA, Symms M, Boulby PA, Alexander DC (2007). Optimal acquisition orders of diffusion-weighted MRI measurements. *J Magn Reson Imaging* 25:1051–1058.

Dubois J, Poupon C, Lethimonnier F, Le Bihan D (2006). Optimized diffusion gradient orientation schemes for corrupted clinical DTI data sets. *MAGMA* 19:134–143.

Enzmann DR, Pelc NJ (1992). Brain motion: measurement with phase-contrast MR imaging. *Radiology* 185:653–660.

Farrell JA, Landman BA, Jones CK, Smith SA, Prince JL, van Zijl PC, Mori S (2007). Effects of signal-to-noise ratio on the accuracy and reproducibility of diffusion tensor imaging-derived fractional anisotropy, mean diffusivity, and principal eigenvector measurements at 1.5 T. *J Magn Reson Imaging* 26:756–767.

Feinberg DA, Mark AS (1987). Human-brain motion and cerebrospinal-fluid circulation demonstrated with MR velocity imaging. *Radiology* 163: 793–799.

Frank LR (2002). Characterization of anisotropy in high angular resolution diffusion-weighted MRI. *Magn Reson Med* 47:1083–1099.

Friston KJ, Holmes AP, Poline JB, Grasby PJ, Williams SC, Frackowiak RS, Turner R (1995). Analysis of fMRI time-series revisited. *Neuroimage* 2:45–53.

Gui M, Tamhane AA, Arfanakis K (2008). Contribution of cardiac-induced brain pulsation to the noise of the diffusion tensor in Turboprop diffusion tensor imaging (DTI). *J Magn Reson Imaging* 27:1164–1168.

Hasegawa M, Rodbard S (1979). Effect of posture on arterial pressures, timing of the arterial sounds and pulse-wave velocities in the extremities. *Cardiology* 64:122–132.

Haselgrove JC, Moore JR (1996). Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magn Reson Med* 36:960–964.

Horsfield MA (1999). Mapping eddy current induced fields for the correction of diffusion-weighted echo planar images. *Magn Reson Imaging* 17:1335–1345.

Huang H, Ceritoglu C, Li X, Qiu A, Miller MI, van Zijl PC, Mori S (2008). Correction of B_0 susceptibility induced distortion in diffusion-weighted images using large-deformation diffeomorphic metric mapping. *Magn Reson Imaging* 26:1294–1302.

Jansons KM, Alexander DC (2003). Persistent angular structure: new insights from diffusion MRI data. Dummy version. *Inf Process Med Imaging* 18:672–683.

Jezzard P, Balaban RS (1995). Correction for geometric distortion in echo planar images from B_0 field variations. *Magn Reson Med* 34:65–73.

Jiang S, Xue H, Counsell S, Anjari M, Allsop J, Rutherford M, Rueckert D, Hajnal JV (2007). In-utero three dimension high resolution fetal brain diffusion tensor imaging. *Med Image Comput Assist Interv* 10:18–26.

Jiang S, Xue H, Counsell S, Anjari M, Allsop J, Rutherford M, Rueckert D, Hajnal JV (2009). Diffusion tensor imaging (DTI) of the brain in moving subjects: application to in-utero fetal and ex-utero studies. *Magn Reson Med* 62:645–655.

Jones DK (2004). The effect of gradient sampling schemes on measures derived from diffusion tensor MRI: a Monte Carlo study. *Magn Reson Med* 51:807–815.

Jones DK, Pierpaoli C (2005). The contribution of cardiac pulsation to variability in tractography results. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 13:225.

Kasprian G, Brugger PC, Weber M, Krssak M, Krampl E, Herold C, Prayer D (2008). In utero tractography of fetal white matter development. *Neuroimage* 43:213–224.

Kim DH, Chung S, Vigneron DB, Barkovich AJ, Glenn OA (2008). Diffusion-weighted imaging of the fetal brain in vivo. *Magn Reson Med* 59: 216–220.

Kim SG, Hendrich K, Hu X, Merkle H, Ugurbil K (1994). Potential pitfalls of functional MRI using conventional gradient-recalled echo techniques. *NMR Biomed* 7:69–74.

Koay CG, Carew JD, Alexander AL, Basser PJ, Meyerand ME (2006). Investigation of anomalous estimates of tensor-derived quantities in diffusion tensor imaging. *Magn Reson Med* 55:930–936.

Kruger G, Kastrup A, Glover GH (2001). Neuroimaging at 1.5 T and 3.0 T: comparison of oxygenation-sensitive magnetic resonance imaging. *Magn Reson Med* 45:595–604.

Leemans A, Jones DK (2009). The B-matrix must be rotated when correcting for subject motion in DTI data. *Magn Reson Med* 61:1336–1349.

Maes F, Collignon A, Vandermeulen D, Marchal G, Suetens P (1997). Multimodality image registration by maximization of mutual information. *IEEE Trans Med Imaging* 16:187–198.

Mangin JF, Poupon C, Clark C, Le Bihan D, Bloch I (2002). Distortion correction and robust tensor estimation for MR diffusion imaging. *Med Image Anal* 6:191–198.

Nunes RG, Jezzard P, Clare S (2005). Investigations on the efficiency of cardiac-gated methods for the acquisition of diffusion-weighted images. *J Magn Reson* 177:102–110.

Özarslan E, Mareci TH (2003). Generalized diffusion tensor imaging and analytical relationships between diffusion tensor imaging and high angular resolution diffusion imaging. *Magn Reson Med* 50:955–965.

Pajevic S, Pierpaoli C (1999). 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* 42:526–540.

Peterson DJ, Landman BA, Cutting LE (2008). The impact of robust tensor estimation on voxel-wise analysis of DTI data. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 16:1823.

Pierpaoli C, Basser PJ (1996). Toward a quantitative assessment of diffusion anisotropy. *Magn Reson Med* 36:893–906.

Pierpaoli C, Jezzard P, Basser PJ, Barnett A, Di Chiro G (1996). Diffusion tensor MR imaging of the human brain. *Radiology* 201:637–648.

Pierpaoli C, Marenco S, Rohde G, Jones DK, Barnett AS (2003). Analyzing the contribution of cardiac pulsation to the variability of quantities derived from the diffusion tensor. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 11:70.

Poncelet BP, Wedeen VJ, Weisskoff RM, Cohen MS (1992). Brain parenchyma motion: measurement with cine echo-planar MR imaging. *Radiology* 185:645–651.

Raj D, Anderson AW, Gore JC (2001). Respiratory effects in human functional magnetic resonance imaging due to bulk susceptibility changes. *Phys Med Biol* 46:3331–3340.

Reese TG, Feinberg DA, Dou JG, Wedeen VJ (2002). Phase contrast MRI of myocardial 3D strain by encoding contiguous slices in a single shot. *Magn Reson Med* 47:665–676.

Reese TG, Heid O, Weisskoff RM, Wedeen VJ (2003). Reduction of eddy-current-induced distortion in diffusion MRI using a twice-refocused spin echo. *Magn Reson Med* 49:177–182.

Righini A, Bianchini E, Parazzini C, Gementi P, Ramenghi L, Baldoli C, Nicolini U, Mosca F, Triulzi F (2003). Apparent diffusion coefficient determination in normal fetal brain: a prenatal MR imaging study. *AJNR Am J Neuroradiol* 24:799–804.

Robson MD, Porter DA (2005). Reconstruction as a source of artifact in non-gated single-shot diffusion-weighted EPI. *Magn Reson Imaging* 23:899–905.

Rohde GK, Barnett AS, Basser PJ, Marenco S, Pierpaoli C (2004). Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magn Reson Med* 51:103–114.

Skare S, Andersson JLR (2001). On the effects of gating in diffusion imaging of the brain using single-shot EPI. *Magn Reson Imaging* 19:1125–1128.

Skare S, Li T, Nordell B, Ingvar M (2000). Noise considerations in the determination of diffusion tensor anisotropy. *Magn Reson Imaging* 18: 659–669.

Studholme C, Hill DL, Hawkes DJ (1996). Automated 3-D registration of MR and CT images of the head. *Med Image Anal* 1:163–175.

Tremblay M, Tam F, Graham SJ (2005). Retrospective coregistration of functional magnetic resonance imaging data using external monitoring. *Magn Reson Med* 53:141–149.

Triantafyllou C, Hoge RD, Krueger G, Wiggins CJ, Potthast A, Wiggins GC, Wald LL (2005). Comparison of physiological noise at 1.5 T, 3 T and 7 T and optimization of fMRI acquisition parameters. *Neuroimage* 26:243–250.

Tuch DS, Reese TG, Wiegell MR, Makris N, Belliveau JW, Wedeen VJ (2002). High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magn Reson Med* 48:577–582.

Turner R, Lebihan D, Maier J, Vavrek R, Hedges LK, Pekar J (1990). Echo-planar imaging of intravoxel incoherent motion. *Radiology* 177:407–414.

Walker L, Chang L, Kanterakis E, Simonyan K, Pierpaoli C, Verma R, Pierpaoli C (2009). Statistical assessment of the effects of physiological noise and artifacts in a population analysis of diffusion tensor MRI data. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 17:459.

Walker L, Yang J, Wu X, Simonyan K, Verma R, Pierpaoli C (2008). Regional distribution of outliers of diffusion MRI in the human brain. *Proc Intl Soc Mag Reson Med Sci Meet Exhib* 16:139.

Wedeen VJ, Hagmann P, Tseng WY, Reese TG, Weisskoff RM (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magn Reson Med* 54:1377–1386.

Weisskoff R, Baker J, Belliveau J, Davis T, Kwong K, Cohen M, Rosen B (1993). Power spectrum analysis of functionally weighted MR data: what's in the noise? In *Proceedings of the Society of Magnetic Resonance, First Annual Meeting*, New York, p. 7.

Wells WM, 3rd, Viola P, Atsumi H, Nakajima S, Kikinis R (1996). Multimodal volume registration by maximization of mutual information. *Med Image Anal* 1:35–51.

Wirestam R, Greitz D, Thomsen C, Brockstedt S, Olsson MBE, Stahlberg F (1996). Theoretical and experimental evaluation of phase-dispersion effects caused by brain motion in diffusion and perfusion MR imaging. *J Magn Reson Imaging* 6:348–355.

Wu M, Chang L-C, Walker L, Lemaitre H, Barnett AS, Marenco S, Pierpaoli C (2008). Comparison of EPI distortion correction methods in diffusion tensor MRI using a novel framework. *MICCAI Proceedings* 2:321–329.

Zaitsev M, Dold C, Sakas G, Hennig J, Speck O (2006). Magnetic resonance imaging of freely moving objects: prospective real-time motion correction using an external optical motion tracking system. *Neuroimage* 31:1038–1050.

Zhou XJ, Du YP, Bernstein MA, Reynolds HG, Maier JK, Polzin JA (1998). Concomitant magnetic-field-induced artifacts in axial echo planar imaging. *Magn Reson Med* 39:597–605.