CHAPTER

8

Moving Beyond DTI: High Angular Resolution Diffusion Imaging (HARDI)

There is no doubt that diffusion tensor imaging has had a significant impact in neuroscience within a relatively short amount of time. Nonetheless, it does suffer from a number of limitations, many of which are now increasingly recognized as being serious confounds. These include in particular the so-called "crossing-fiber" problem, as well as the non-Gaussian nature of diffusion, both of which cannot be adequately described using the diffusion tensor model. This has prompted the development over the last decade of numerous methods to avoid these limitations. In this chapter, we outline these limitations and review some of these alternative models.

8.1 LIMITATIONS OF DIFFUSION TENSOR IMAGING

As previously highlighted, the diffusion tensor model is a gross over-simplification of the actual anatomy. It is inherently based on the assumption of Gaussian free diffusion, which is clearly not the case even in coherently oriented white matter bundles. Even in the ideal case, there are many barriers to the movement of water molecules, including axonal membranes and myelin sheaths, so that diffusion in this environment is far from "free". There are two situations where these limitations become apparent. One occurs when the diffusion signal is measured over multiple *b*-values, in which case the signal attenuation curve is clearly non-monoexponential (Stanisz et al., 1997; Clark and Le Bihan, 2000; Beaulieu, 2002; Jensen et al., 2005.). The other situation occurs in crossing fiber voxels, in which case the simple diffusion tensor model is a poor fit to the observed DW signal, and cannot provide an adequate characterization of the fiber arrangement (Tuch et al., 2002; Wedeen et al., 2005; Alexander et al., 2001). This last issue is particularly problematic for DTI applications, for the reasons outlined below.

8.1.1 The Issue of Crossing Fibers

The fact that the diffusion tensor is affected by crossing fibers has been known since its very invention: in 1996, Pierpaoli and Basser attributed the observation of low FA values in many deep white matter regions to incoherence in the fiber orientations – in other words, "crossing fiber" effects (Pierpaoli and Basser, 1996). The term "crossing fibers" really refers to any situation in which the fiber orientation is not unique, including interdigitating fibers and adjacent fiber bundles brushing past each other, but also other situations such as bending or diverging fiber bundles, where fibers are not crossing as such, but the fiber orientation is not unique nonetheless. A few examples of these situations are illustrated in Fig. 8.1.

In such situations, the diffusion-weighted signal is not actually well described by the tensor model, as shown in Fig. 8.2. It's easy to see in these examples that there are features in the signal that are simply too complex for the simple tensor model. It is no longer possible to estimate the fiber orientations correctly, and the anisotropy of the tensor will also be strongly affected. Both of these effects have serious implications for both tractography and for the interpretation of anisotropy as a marker of white matter "integrity".

8.1.2 Crossing Fibers are Endemic in Diffusion Imaging

The effect of crossing fibers can easily be seen in FA maps, particularly when compared with the directionally-encoded color (DEC) maps. In Fig. 8.3, low

© 2014 Elsevier B.V. All rights reserved.

FA values can be seen in many areas of pure white matter; by looking at the DEC map, it is clear that these areas correspond to voxels bordering two distinct orientations. Voxels in these regions will contain contributions from both fiber tracts due to partial volume averaging, and this will cause a drop in anisotropy. These effects are inevitable given the low resolution typical of diffusion-weighted imaging (of the order of 2 mm), and the relatively small size of most white matter tracts: even the cortico-spinal tract is only 8 mm wide at its thickest point in the internal capsule, and 3 mm thick in subcortical regions (Ebeling and Reulen, 1992).

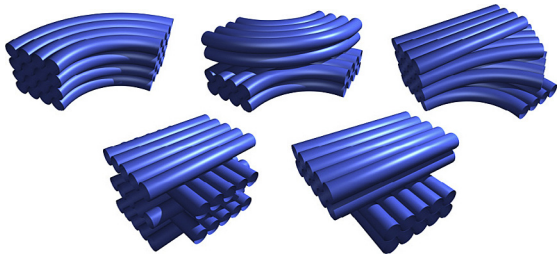

FIGURE 8.1 Different fiber configurations that contain multiple fiber orientations, which the diffusion tensor model cannot characterize adequately, and therefore fall into the "crossing fiber" category. These include actual crossing fibers (bottom row), such as interdigitating fibers (left) or simply two distinct fibers bundles brushing past each other (right). Even though they do not contain crossing fibers as such, curving fibers are also problematic since they contain a range of orientations (top left). Any combination of these configurations also lead to problems, for example two curving bundles brushing past each other (top center) will contain both a range of orientations (since the fibers are curving) and actual crossing fibers (where the two bundles have different orientations). Diverging or "fanning" fibers also contain a mixture of these effects (top right). As can be appreciated, there is a large range of configurations that would give rise to "crossing fiber" effects, and straight fibers really form only a small subset of the full range of configurations likely to be found in reality.

These effects were initially believed to be restricted to a few problem areas, with the rest of the white matter being adequately modeled with the tensor. However, it is now clear that voxels containing "crossing fibers" are very widespread throughout the white matter. In fact, these effects can be detected in 90% of white matter voxels (Jeurissen et al., 2012), as shown in Fig. 8.4. This makes sense given the anatomy under consideration, and strongly argues for a move away from the traditional single tensor model, since it is now clear that it cannot provide the type of information that most people expect to get from it: fiber orientations, and a marker of fiber bundle microstructure or "integrity."

8.1.3 Crossing Fibers have a Serious Impact on Anisotropy

The impact of crossing fiber on tensor-derived anisotropy measures is profound, and easy to see on FA maps such as those shown in Fig. 8.3. In fact, it is not uncommon to see regions of deep white matter with near-zero anisotropy values. When looking at these regions with more advanced higher-order models, it is clear that this is simply due to the presence of two or more fiber bundles within the same voxel. An example of this is shown in Fig. 8.5.

Clearly, knowing that crossing fibers are present completely changes our interpretation of differences in FA. In fact, a number of studies have shown elevations in FA in cases where a reduction in FA would have been expected given the pathology or condition (Douaud et al., 2011; Tuch et al., 2005). The interpretation was that this was probably due to the expected reduction in FA for one fiber bundle, with no changes in other fiber bundles crossing through that region; this is illustrated in Fig. 8.6.

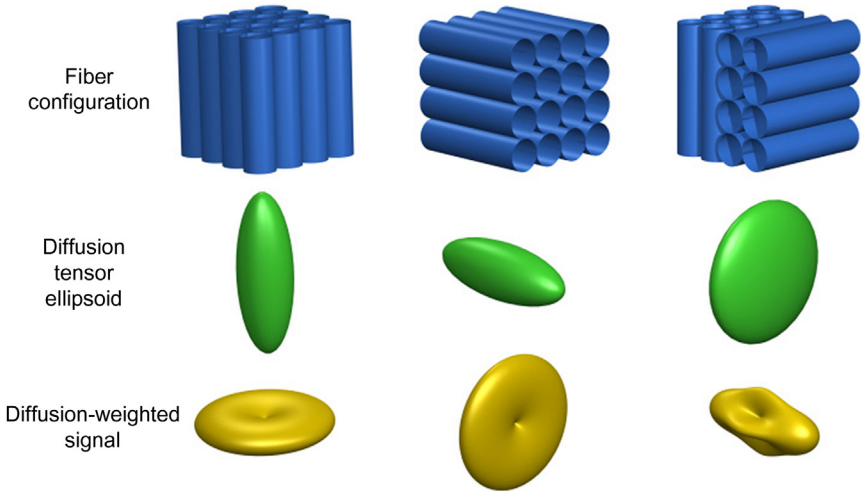

FIGURE 8.2 In straight, coherent "single fiber" bundles (left and center), the DW signal (yellow) is well described by the tensor model (green). When crossing fibers co-exist within the same imaging voxel, the observed DW signal is essentially the sum of the DW signals from each individual bundle. In this example, the DW signal in the crossing fiber case (bottom right) is the sum of the corresponding DW signal for the two crossing fiber cases (bottom left and centre). This contains features that the tensor model cannot represent: the diffusion tensor can only represent an ellipsoid, which is effectively a sphere that has been stretched or compressed in various directions. It cannot capture the sharp features that can be seen in the DW signal, which are essential for detecting and resolving crossing fibers.

The take-home message from these studies is that while a given fiber bundle can have a reduction in its intrinsic FA, this will often be observed as an *elevation* of FA in crossing fiber regions that the bundle travels through. This can sometimes be clarified to some extent by the observation of a reduction in FA in a different, single-fiber part of the same tract, in which case the simplest explanation is that this tract has reduced intrinsic FA throughout its length, which manifests as an increase in crossing fiber regions. Unfortunately, this is not always possible, as elevations in FA might be the *only* significant observation in a particular study. Obviously, this complicates the interpretation of any observed FA differences, and the simplistic interpretation of FA as a marker of white matter "integrity" is simply flawed.

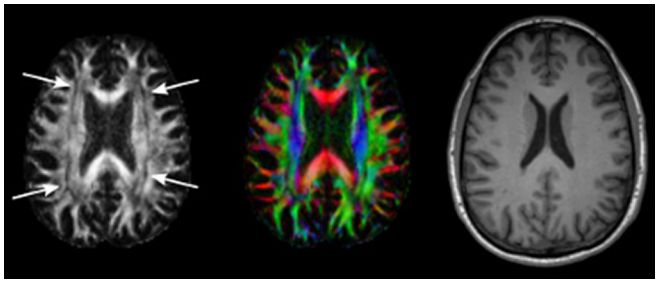

FIGURE 8.3 The effects of crossing fibers can be seen on a simple FA map. Bands of low anisotropy can readily be seen in FA maps (highlighted by the arrows, left) in regions of deep white matter remote from any grey matter (right). When compared with the corresponding directionally-encoded color (DEC) map, it is clear that many of the regions affected correspond to regions bordering distinct colors – i.e. fiber bundles with different orientations. The most likely explanation for the reduced FA in these regions is therefore that it is due to the presence of crossing fibers.

8.1.4 Crossing Fibers have a Serious Impact on DTI Fiber-Tracking

The impact of crossing fibers on tractography results is primarily due to inaccuracies in the estimated fiber orientations. The way that streamlines are propagated in most tractography algorithms (see Chapter 9) is based on the local estimate of the fiber orientation. Obviously, even a single bad orientation estimate is enough to cause the algorithm to veer off-course. Even if the error is small locally, it might be enough for the algorithm to now be following an adjacent, but otherwise completely unrelated tract, as illustrated in Fig. 8.7. Furthermore, the very fact that the error is due to crossing fiber effects implies that there must be other tracts in the immediate vicinity that the algorithm might choose to follow; in other words, the

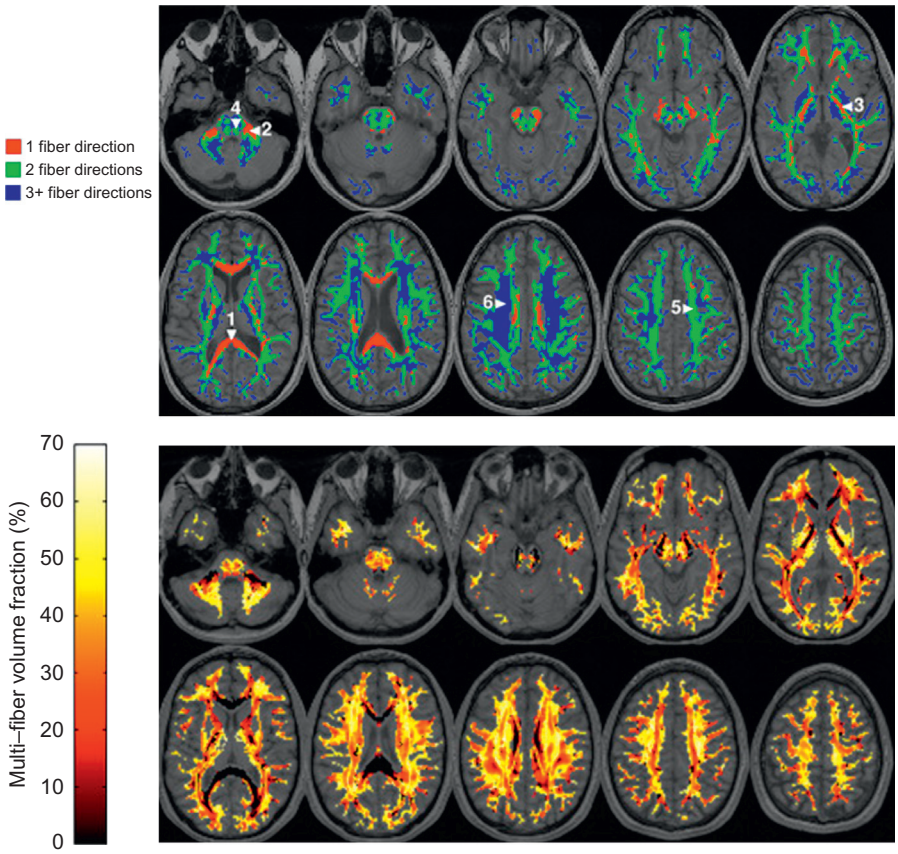

FIGURE 8.4 The crossing fiber problem has recently been shown to be very widespread, as shown here. In the top panel, white matter voxels where only one fiber population can be detected are shown in red, those where two distinct orientations can be detected are in green, and voxels where three or more can be detected are shown in blue. Single-fiber regions are essentially restricted to the mid-sagittal portion of the corpus callosum, the corticospinal tracts, and the body of the middle cerebellar peduncles. The bottom panel shows the proportion of each voxel that is taken up by non-dominant fiber orientations; in other words, that is contaminated by crossing fibers. This demonstrates that these crossing fibers are not just detectable, but that the level of contamination is far from negligible. (*Images reproduced from Jeurissen et al., Human Brain Mapping 2012, doi: http://dx.doi.org/10.1002/hbm.22099, with permission.*)

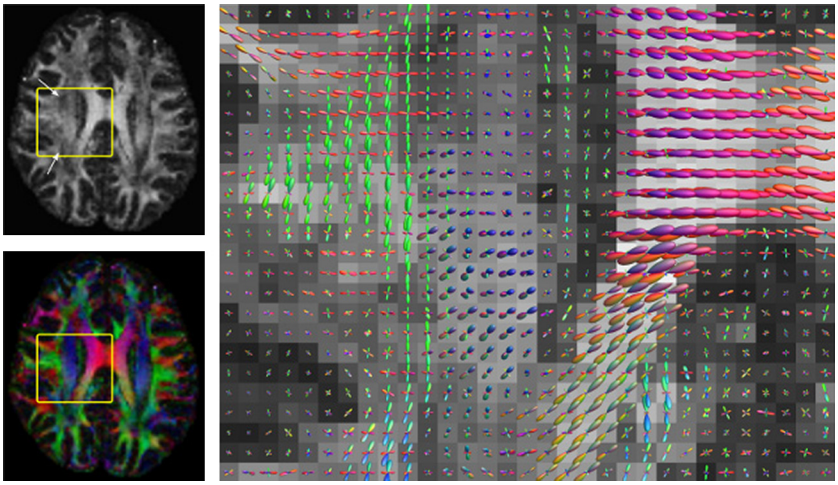

FIGURE 8.5 When looking at regions of low anisotropy using higher-order models, crossing fibers can readily be observed. In this example, the region of the centrum semiovale highlighted on the axial FA map (top left) can be seen to contain regions of low FA, which correspond to interfaces between bundles with distinct fiber orientations (as can be seen from the DEC map, bottom left). The main panel shows the fiber orientation distributions estimated from the same data using constrained spherical deconvolution (described later) (Tournier et al., 2007). Crossing fibers can be seen particularly in the same low FA regions.

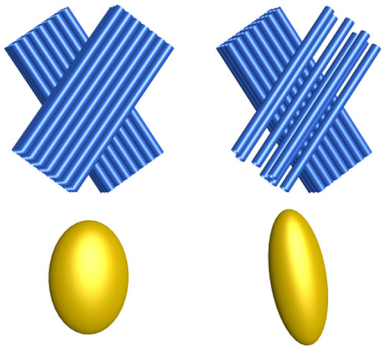

FIGURE 8.6 A reduction in anisotropy in one bundle (due for example to degeneration) can lead to increased overall anisotropy for that region. This can happen when the region originally contained crossing fibers of comparable volume fraction (top left), and one fiber bundle is selectively affected (top right) so that the intrinsic anisotropy of that bundle is reduced. The diffusion ellipsoid in the original case was not fully anisotropic, with an intermediate FA value due to the crossing fibers (bottom left). If one of the fiber bundles degenerates, the amount of crossing fiber contamination is reduced, and the overall FA consequently goes up, with a corresponding more elongated tensor ellipsoid (bottom right).

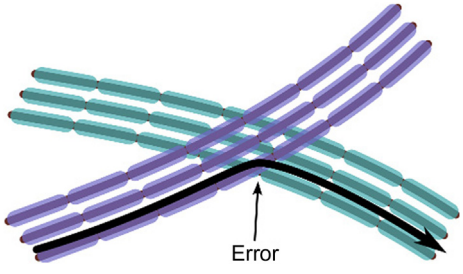

FIGURE 8.7 An error at one point along the tract can cause the algorithm to veer into an adjacent fiber pathway, and so infer connectivity to the wrong region. Since this will in most cases be due to crossing fiber artifacts, the point at which the error is most likely to occur is precisely where the fiber-tracking algorithm can switch to an adjacent fiber bundle, amplifying the impact of an error that might otherwise have been negligible.

regions where the largest errors are likely to occur are also the regions where such errors have the most dramatic impact. Obviously, this can easily amplify a relatively trivial, local inaccuracy to produce a completely different (and obviously incorrect) delineation of a white matter pathway.

In fact, the impact of crossing fibers on tractography is even more insidious, since fiber-tracking involves propagation through the image. To delineate a pathway from one region to the other might require tracking over 10–100 voxels, any one of which might contain crossing fibers. Clearly, even a small proportion of affected voxels would be enough to affect a disproportionately large fraction of the generated tracks. Even if only a third of white matter voxels were affected, it would still be very difficult, if not impossible, to find even a single fiber pathway that bypassed all of them, and hence remained free from crossing-fiber artifacts.

The most obvious consequence of these tracking errors is the appearance of false-positives in the results: the algorithm produces tracks that are obviously not anatomically plausible. A classic example of such a false positive is shown in Fig. 8.8. While it is possible (indeed desirable) to impose prior anatomical information to "filter out" such obvious false positives (see Chapter 9), this can only be done when such prior information exists and is reliable. This makes it particularly difficult to interpret unexpected fiber-tracking results, for example where an unknown pathway is observed in a particular cohort. Given the number of possible confounding factors, such a finding will more than likely be an artifact of the fiber-tracking. Only when robust corroborating evidence can be found supporting its existence can such a finding be given serious consideration.

A less obvious, but potentially far more devastating consequence is the occurrence of false negatives. This is

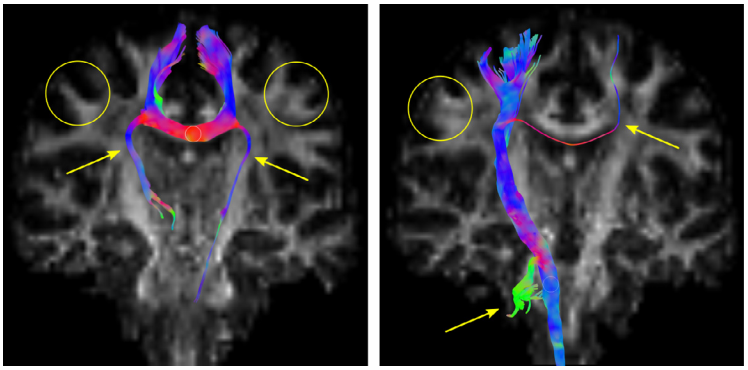

FIGURE 8.8 Fiber-tracking results obtained using DTI by tracking from the mid-body of the corpus callosum (left) and from the cortico-spinal tract at the level of the pons (right)—the seed region is shown by the small white circles. False positives are highlighted by the yellow arrows, and consist of artifactual streamlines from the corpus callosum down into the cortico-spinal tracts (left), and from the cortico-spinal tract into the ipsilateral middle cerebellar peduncle (right). These results are also affected by false negatives, highlighted by the yellow circles; these consist of the lateral projections from the corpus callosum (left) and the lateral projections from the cortico-spinal tract (right).

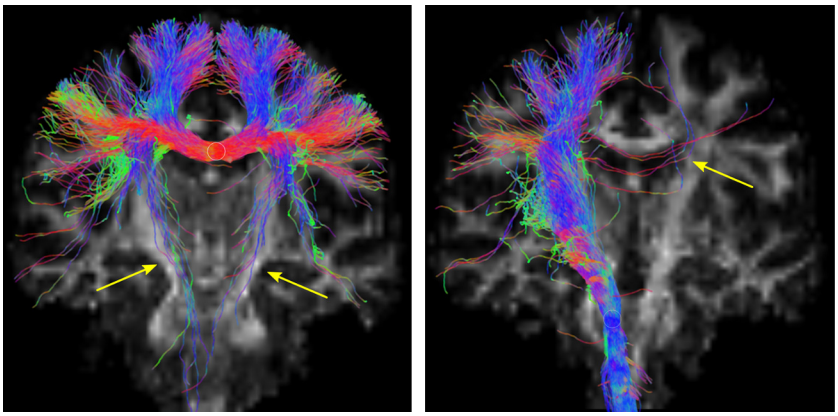

FIGURE 8.9 Fiber-tracking results obtained using a more advanced high-order model approach combined with a probabilistic streamlines algorithm, using the same dataset and parameters as for Fig. 8.8. This approach produces a much richer reconstruction, in which the lateral projections of the corpus callosum (left) and cortico-spinal tract (right) can readily be observed. Note that false positives are still present in these results, as highlighted by the yellow arrows. In this case however, since a probabilistic fiber-tracking algorithm was used, the low density of these false positive connections can be interpreted as indicative of low probability.

particularly problematic for the thinner parts of white matter tracts, since in crossing fiber regions the tensor-derived fiber orientation will be driven by the other, more dominant tracts, making it impossible to delineate these smaller tracts. Such false negatives are a very frequent problem, and can easily be observed even in major white matter pathways such as the cortico-spinal tracts and the corpus callosum, as shown in Fig. 8.8. In fact, results in these kinds of pathways are particularly misleading, since the algorithm is able to delineate parts of the tracts; this makes it very easy to overlook the fact that the delineation is only partial. For example, it very rare to observe any lateral projections of either the cortico-spinal tracts or of the corpus callosum with tensor-based fiber-tracking, even though these are known to exist and are readily observable using fiber-tracking based on higher-order models (see Fig. 8.9) (Farquharson et al., 2012). This becomes a serious issue when such results are used clinically, in particular for neurosurgical planning, where the question is typically one of establishing where the important tracts are *not* present; if parts of the tract are not identified by the fiber-tracking, the obvious interpretation that it is safe to operate in these regions is clearly flawed. In fact, sole reliance on DTI-based fiber tracking for neurosurgical planning has been reported to result in post-operative loss of function that could otherwise have been avoided (Kinoshita et al., 2005).

8.2 CONCEPTS FOR DEALING WITH CROSSING FIBERS

It is now clear that crossing fibers deeply affect DTI results, and at the very least this must be taken into account when interpreting the results. However, in many cases it will be impossible to make a robust interpretation of the data, since there may be many different explanations for the results, each of which may be equally plausible. Fiber-tracking results are particularly difficult to interpret since there is no information as to which pathway an artifact-free method might have produced – all that is known is that errors are likely to exist, and that these errors might lead to potentially drastically different results.

To address these issues, it is clear that alternative methods are required that are robust to crossing fiber effects. Over the last decade, a large number of such methods have been proposed, many of which are now being used clinically with promising results. In this section, we describe some of the concepts and theory behind these methods.

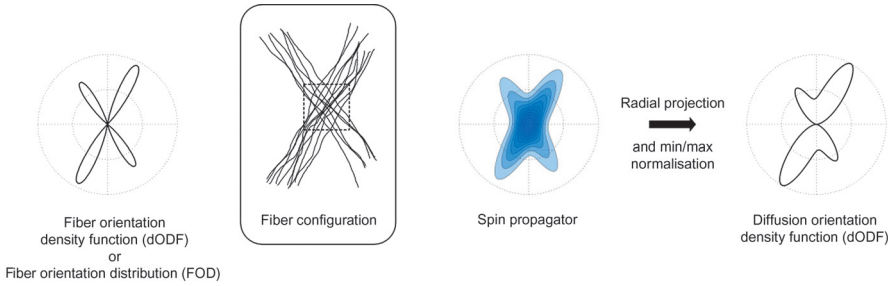

FIGURE 8.10 The relationship between the fiber configuration, the spin propagator, diffusion ODF, and fiber ODF. Consider a voxel containing crossing fibers (center left, dotted square). In this environment, water molecules will be more likely to diffuse by larger distances along the fiber directions. This is represented by a diffusion propagator with elongated ridges along these directions (center right). To extract and display fiber orientation information from the spin propagator, it is common practice to collapse this information down onto the orientation domain by radial projection, to form the diffusion orientation density function (ODF) (right). This function essentially represents the "mass" of the propagator as a function of direction, and therefore contains peaks along the directions of the ridges in the spin propagator. For display purposes, it is typically scaled so that its minimum and maximum values are zero and one respectively. An alternative representation for the fiber configuration is by the fiber ODF (left). This function represents the "amount" (i.e. partial volume) of fibers aligned along the direction of interest, and is therefore more closely related to the information typically required, e.g., for use in fiber-tracking.

The simplest situation to consider is what happens when two fiber bundles cross within the same imaging voxel. As already shown in Fig. 8.2, in this case the DW signal has sharper features than can be described using the tensor model. The question is now one of finding a suitable way to represent the information that we need to extract, and modeling the way this information relates to the measured signal. There are two main ways of looking at the problem; one focuses on the properties of the diffusion process, the other explicitly models the microstructure of interest.

8.2.1 Diffusion in Complex Fiber Arrangements: The Spin Propagator

It is possible to estimate the properties of the diffusion process without making any attempt at relating this to the microstructure explicitly. The information provided can then be used to infer more relevant features in subsequent steps. Conceptually, the advantage of this type of approach is that it does not rely on an explicit model of the microstructure, and so is theoretically less prone to biases introduced by any assumptions inherent in such a model – and given the complexity of the cellular microstructure, such assumptions are inevitable. These methods are therefore often said to be *model-free*, although this is a slight misrepresentation since subsequent processing steps needed to derive microstructural information from the spin propagator will inevitably require a model, with all the caveats that come with it (discussed below). Nonetheless, the relationship between the diffusion-weighted signal and the displacement of the water molecules is well-known, and is given by the theory of q-space (Callaghan, 1996).

Under the q-space framework, the diffusion properties are characterized by an entity known as the *spin propagator* (also referred to as the spin displacement probability density function (PDF)). This function essentially quantifies the relative amount of water molecules (spins) that have diffused away from their starting positions by a given displacement. More formally, the spin propagator can be written as $P(x, r)$, which essentially translates to: the proportion of molecules at position x that have moved by a distance r within the fixed diffusion time of the experiment. Typically, x refers to the voxel location, with resolution of the order of millimeters, while r refers to the displacement of water molecules, with resolution of the order of microns. As illustrated in Fig. 8.10, the spin propagator efficiently captures all the diffusion-related information, and is clearly dependent on the tissue microstructure: spins will tend to diffuse larger distances along fiber orientations, forming elongated "ridges" in the spin propagator.

In practice, the spin propagator will typically be collapsed down to provide a more efficient, reduced representation: the diffusion orientation density function (ODF), often also referred to as the dODF, or simply the ODF (this last acronym is not recommended since it could cause confusion with the fiber ODF – see below). This is done by a process of radial projection: conceptually, this involves computing the proportion of spins that have moved along any given direction, regardless of the distance traveled. It can be written more formally as $\Psi(x, \hat{u})$, providing the proportion of water molecules that have moved along the unit direction $\hat{u}$ within the voxel at location x. As illustrated in Fig. 8.10, the dODF provides a reduced representation of the spin propagator, efficiently capturing the information related to its angular dependence.

Most fiber-tracking algorithms that make use of such data will typically use the orientations of the peaks of the dODF as the estimate of the fiber orientation. While this generally results in vastly improved estimates compared to the diffusion tensor, using the peaks does have its limitations, due to the fact that the dODF typically contains broad peaks. This makes it difficult (if not impossible) to resolve fiber orientations that are aligned at relatively shallow crossing angles, and has also been shown to introduce a bias in the orientation estimates when the crossing angle is not 90° (Tournier et al., 2008; Zhan and Yang, 2006). The other disadvantage of using the peak orientations is that this process implicitly assumes that the fiber orientations are discrete and coherently oriented, which makes it impossible to characterize curving or diverging fiber bundles.

8.2.2 The Fiber Orientation Density Function Provides a More Direct Description of Fiber Arrangement

An alternative way to look at the problem is to focus directly on the desired information, typically the fiber orientations. The question then is: if we know what fiber orientations are present and in what quantities, can we devise a model to predict the DW signal intensities? Obtaining the fiber orientations then reduces to finding the set of model parameters (fiber orientations, volume fractions, etc.) that best explain the measured data – in other words, solving the inverse problem. The advantage of these methods is that by focusing directly on the desired information, it becomes possible to estimate the parameters of interest using much fewer measurements. In other words, since the model only contains parameters that we are interested in, we only need to acquire as many measurements as are required to get a suitable estimate of these parameters, without needing to characterize the whole spin propagator. The other advantage is that we get the information we need directly, with no further processing required.

There are various ways to represent fiber orientation information. The simplest is simply as a list of orientations and their associated volume fractions. This however requires the method to also estimate how many fiber orientations there are in each voxel, which is another potential source of error. It also implicitly assumes that each fiber bundle is completely straight and coherent, with no spread about its orientation. In other words, it cannot account for curving or diverging fiber bundles, which would exhibit a range of orientations. A more general representation of the fiber orientations is provided by the fiber orientation density function (fODF), also known as the fiber orientation distribution (FOD). This entity represents the amount of fibers aligned with any particular direction, as a continuous distribution over the sphere. More formally, it can be written $f(x, \hat{u})$, providing the density of fibers along the unit direction $\hat{u}$ within the voxel at location x. As illustrated in Fig. 8.10, the FOD efficiently captures all the information related to the arrangement of the fiber bundles.

Note that the FOD does not contain any information related to the diffusion or microstructural properties of the fiber bundles. This is a consequence of the fact that most HARDI methods are geared towards identifying fiber orientations, and it is much easier to estimate them when the other variables are not considered. This potential limitation will be discussed in more detail in later sections.

8.3 PRACTICAL APPROACHES TO DEALING WITH CROSSING FIBERS

We need to consider a number of aspects to estimate the fiber orientations in crossing fiber situations. First, the data acquisition will need to be modified to ensure that the data acquired actually contains the most relevant information possible. This will typically be dependent on the approach taken to estimate the orientations. The next stage is to supply these data to the appropriate algorithm for processing. Each of these stages will require careful consideration to ensure the results are as good as they can be within the amount of time available for scanning.

8.3.1 HARDI Data Acquisition: Diffusion Gradient Directions and *b*-Values

The first issue is that of data acquisition. While a lot of work has gone into optimizing the acquisition protocol for DTI, it is less clear what the best parameters are for higher-order models. This is partly due to the fact that these approaches are more recent, but also because there are many different algorithms available, each of which includes different number of parameters, and each of which might produce output that is different either in kind (e.g., spin propagator vs. fiber orientations) or in representation (e.g., discrete orientations vs. continuous ODF). Moreover, different methods might require different types of data acquisitions. This makes it very difficult to provide a single set of optimal parameters suitable for all methods.

Nonetheless, it is possible to make some general statements about the data acquisition. The following considerations apply to the HARDI family of acquisition protocols, although many of the concepts will also apply to the more advanced multi-shell HARDI and DSI acquisition protocols, which are discussed

separately below. In essence, HARDI is simply an extension of the simple DTI acquisition protocol: it consists of a number of DW image volumes, each sensitized to diffusion along a given direction with the same *b*-value, with the DW directions distributed uniformly over the half-sphere, along with a set of non-diffusion-weighted, $b = 0$ image volumes. The parameters that need to be considered beyond those of standard DTI (voxel size, number of slices, EPI bandwidth, parallel imaging acceleration factor, etc.) are: how many DW directions, and what *b*-value.

Before discussing these parameters, we should justify why we might want to use HARDI rather than any other type of protocol. We could be acquiring multiple *b*-values along the same directions, or even multiple repeats of the same directions, or any other mixture of different *b*-values along different directions. Essentially, the main reason is one of time efficiency: scan times in clinical practice are typically short, and every effort must be made to optimize the amount of information gathered within this narrow time frame. The aim is to characterize the angular features in the DW signal that are not apparent with DTI, such as those shown in Fig. 8.2. To capture those more jagged features, we need to sample the signal along a larger number of distinct DW directions. This simple observation already dictates an acquisition with a large number of DW directions.

Typically, there is no *a priori* information about the orientations of the white matter structures we might be interested in. This is especially true when the target application is tractography, where tracts will be propagated throughout the white matter, along pathways with a whole range of orientations. This immediately means that the protocol we use must not introduce any bias between structures aligned along different directions; in other words, the results must be rotationally invariant. This implies that the DW directions must be distributed as uniformly as possible, but also that the *b*-values used must be the same for every DW direction.

The final issue is how many *b*-values to use. The main limitation here is one of scan time: given that we already need to acquire a large number of DW directions, acquiring even a single *b*-value already demands a relatively long scan time. All other thing being equal, acquiring more *b*-values would require a further doubling or tripling of the scan time. While there may be distinct advantages to this type of acquisition, the best protocol now depends on which combination provides the most information for the target application. Some studies may be interested in estimating microstructural features, for example using approaches such as diffusion kurtosis imaging (DKI; discussed later) that require multiple *b*-values (Jensen et al., 2005). In this case, the best approach may be to use fewer DW directions with multiple *b*-values, at the expense of the orientation information. However, most studies will focus on fiber-tracking, where orientation information is paramount: in these cases, it makes sense to sample orientation space as densely as possible using a single *b*-value, since this is how the most relevant information is likely to be obtained; this is the HARDI approach.

HARDI is therefore arguably the most efficient protocol that can be devised to obtain the orientation information required to resolve crossing fibers. The remaining questions to be addressed are: how many directions, and what *b*-value? These issues are still to some extent unresolved, although there is a general consensus that as many DW directions should be acquired as time allows; most studies will use at least 50 DW directions, some research studies use more than 150. The optimal *b*-value is less clear, with many studies still using DTI-optimal values of $b \sim = 1000\ \mathrm{s/mm^2}$; in many cases, this choice of *b*-value is probably influenced by the desire to ensure the data can also be processed using DTI. It is also probably influenced by the fact that high *b*-value DW images clearly have reduced signal-to-noise ratio, due to the higher attenuation induced by the diffusion weighting, but also the longer echo time needed to accommodate the long DW gradients needed to achieve that *b*-value. However, what is required to resolve crossing fibers is a high *contrast-to-noise* ratio in the angular domain, and this is best achieved with higher *b*-values, as illustrated in Fig. 8.11. Alexander & Barker recommend a value of $b \sim 2{,}500\ \mathrm{s/mm^2}$ based on simulations (Alexander and Barker, 2005); others suggest values of $b = 3000\ \mathrm{s/mm^2}$ (Tournier et al., 2004; Jones et al., 2012). In general, *b*-values in the region of $2{,}000-3{,}000\ \mathrm{s/mm^2}$ are likely to be broadly optimal in terms of fiber orientation estimation.

There are however some disadvantages to these high *b*-values. While they undoubtedly improve the conspicuity of fiber orientations, the low SNR of these images, coupled with the strong contrast introduced by the high *b*-value, makes it difficult to perform robust motion correction or eddy-current correction on these data, since most approaches to these problems rely on image registration methods. Thankfully, a few methods have been proposed recently to address this problem, and will likely become available in the near future. The other issue introduced by the low SNR is that of magnitude or *Rician* bias: MR images are typically provided as magnitude data, explicitly preventing negative values from appearing – this is unavoidable in DWI because the phase of the image is corrupted by microscopic motion in unpredictable ways. This means that if the noise is large compared to the signal, it will tend to increase the measured signal's intensity. Since DWI and particularly HARDI are low SNR techniques, this

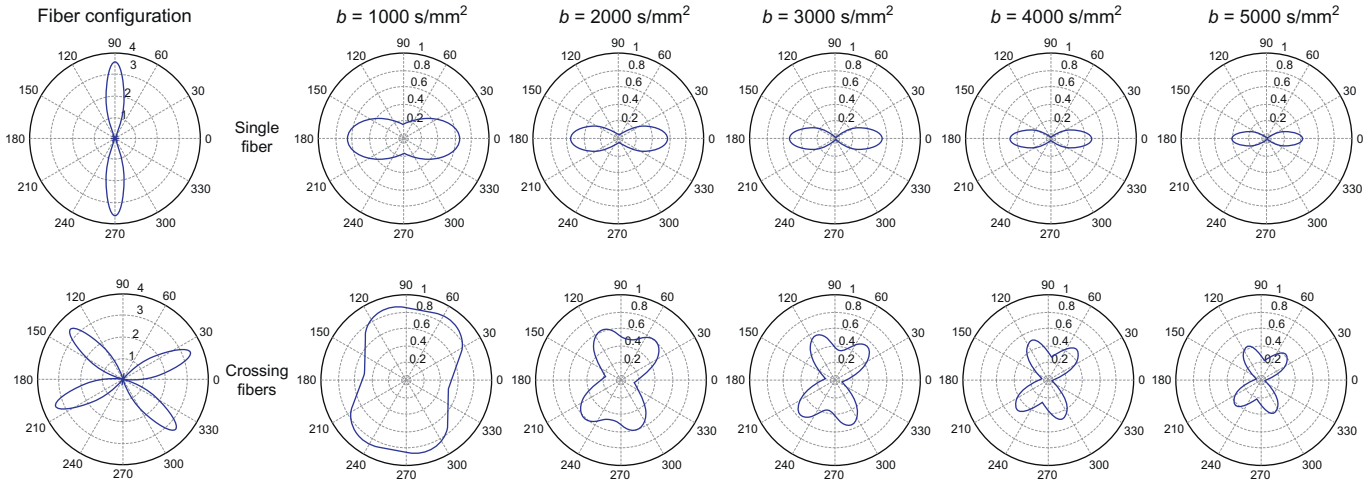

FIGURE 8.11 Dependence of the DW signal on b-value, for a single fiber direction (top) and two fibers crossing at 67° (bottom). At low b-values, the DW signal varies smoothly as a function of orientation, and this essentially "blurs" the angular features when crossing fibers are present. As the b-value is increased, the overall DW signal clearly decreases, but also becomes sharper, with much more pronounced loss of signal along the fiber direction than across it. This means that when crossing fibers are present, the angular features are much more pronounced, and the different orientations become much more conspicuous and can be resolved more readily. At very high b-values, the signal is reduced further, with little increase in the sharpness of the features of interest. Based on these results, the best b-value to use is therefore in the vicinity of $b = 2{,}500{-}3{,}000\ \mathrm{s/mm^2}$.

introduces a bias in the data that can give rise to artifacts such as the appearance of noisy peaks in the results. Again, this is a topic of ongoing research and methods to address these issues are likely to become available in the near future. In the meantime, these issues can be minimized by increasing the voxel size to ensure adequate SNR, at the expense of spatial resolution.

8.3.2 Q-Space Describes the Relationship between the Diffusion Signal and the Spin Propagator

Once suitable data have been acquired, the next issue is how to derive fiber orientation estimates from them. The first class of approaches is aimed at estimating the spin propagator (or reduced representations of it) based on the theory of q-space, which relates the DW signal to the spin propagator via a simple 3D Fourier transform, in the same way that k-space relates to MR images. The challenge here is to acquire data over q-space in such a way that the Fourier transform can be performed. At this point, it is helpful to define what is really meant by q-space. In simple terms, a position in q-space refers to a particular diffusion-encoding gradient vector, according to the relationship $q = \gamma\delta G$, where γ is the proton magnetogyric ratio, δ is the diffusion pulse duration, and G is the gradient vector (direction and amplitude). To perform the 3D Fourier transform directly, data need to be acquired

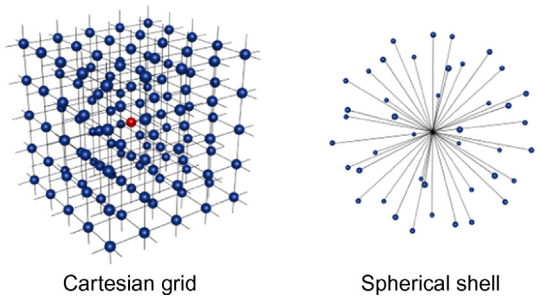

FIGURE 8.12 In Cartesian q-space sampling (as used for example in diffusion spectrum imaging (DSI)), a large number of images need to be acquired, each sensitized using a different diffusion direction and gradient amplitude, so that the set of gradient vectors used covers q-space uniformly (left). On the other hand, high angular resolution diffusion imaging (HARDI) methods rely on dense sampling on a spherical shell in q-space (i.e. using a constant b-value). This involves the acquisition of an intermediate number of images, each sensitized using the same b-value but different diffusion gradient directions, so that the set of gradient vectors covers the sphere uniformly (right).

over a grid of points in q-space – in other words, the end-points of the gradient vectors need to form a regular grid, as illustrated in Fig. 8.12.

This amounts to a potentially very large number of images, with correspondingly large scan times. Nonetheless, this is in essence what the **diffusion spectrum imaging** method (DSI) does. 515 DW image volumes are acquired over a grid in q-space (i.e. different gradient vectors), and the spin propagator is then obtained for each imaging voxel by a straightforward Fourier transform. This technique provides arguably

the most complete characterization of diffusion that can be achieved, in the form of the full spin propagator. However, to extract information from the spin propagator that can be used for applications such as fiber tracking, subsequent steps are typically performed to estimate the diffusion ODF, followed by peak extraction (See Section "Diffusion in complex fiber arrangements: the spin propagator" above). This amounts to a huge reduction in the amount of information, from 515 measurements per voxel down to ~9 to represent three fiber orientations. A criticism of this approach in the context of fiber orientation estimation is therefore that it seems very inefficient to acquire so much data to estimate so few parameters. Nonetheless, the acquisition can be performed in ~25 minutes on a standard clinical scanner, which is within reach of research studies.

Subsequent methods based on *q*-space were developed to estimate the diffusion ODF directly, without needing the full spin propagator, using much shorter HARDI acquisitions. The problem facing these methods is how to perform the 3D Fourier transform when the data have only been collected over a sphere in *q*-space. To get around this, these methods assume a particular radial dependence for the spin propagator (or the DW signal itself), allowing them to perform the Fourier transform while focusing on the angular dependence of the signal. The first such approach is **Q-ball imaging** (QBI), which was designed to provide the diffusion ODF directly from a HARDI acquisition via the so-called Funk-Radon transform (Tuch, 2004). This method has since been refined using a spherical harmonic framework (Hess et al., 2006; Descoteaux et al., 2007), and more recently using solid-angle considerations (Aganji et al., 2010), yielding much improved estimates of the diffusion ODF. Other HARDI methods based on *q*-space include **persistent angular structure (PAS) MRI**, which assumes the distance of water displacements is fixed to perform the 3D Fourier transform (Jansons and Alexander, 2003), and the **diffusion orientation transform** (DOT) method, which assumes mono- or bi-exponential decay of the DW signal as a function of *b*-value to perform its 3D Fourier transform (Özarlan et al., 2006).

8.3.3 Mixture Models Allow us to Extract the Fiber Orientations Directly

Another class of approaches is aimed at estimating the fiber orientations directly. These methods are typically based on so-called *mixture models*, and rely on the assumption of *slow exchange*: water molecules don't have enough time to move between the different fiber bundles within the diffusion time of the experiment. This greatly simplifies the problem, since it implies that the DW signals from the different fiber bundles are independent and simply add up linearly, rather than interacting with each other in complex, non-linear ways. The issue then reduces simply to modeling what the DW signal should be for each fiber bundle, and adding up the different contributions to form the full mixture model.

The assumption of slow exchange is probably safe in most cases, since the displacement of water molecules is of the order of 5 μm within a typical diffusion time, and so exchange can only occur in places where two fiber bundles are directly adjacent. Moreover, there is good evidence that water within the axons exchanges relatively slowly with the extra-cellular compartment, and since ~80% of the water is intracellular, the amount of exchange even with adjacent axons will probably be negligible within typical diffusion times.

The other assumptions commonly made relate to the DW signal for each fiber bundle. Almost all methods will assume that the DW signal is axially symmetric: it does not change if the fiber bundle is rotated about its axis. This is considered a safe assumption, since it is difficult to envisage a biologically realistic situation that might cause such a systematic bias. A more debatable assumption that is commonly made is that the "shape" of the diffusion signal is constant over all fiber bundles; in other words, all fiber bundles have the same intrinsic anisotropy. Obviously, this completely changes our way of looking at diffusion-weighed imaging: we are now attributing all observed variations in anisotropy *entirely* to partial volume effects and crossing fibers – FA and related measures no longer have anything to do with "white matter integrity". While this may at first sight seem like an implausible (even absurd) assumption, it is after all entirely consistent with Basser and Pierpaoli's original interpretation of anisotropy variations in white matter as differences in the coherence of the fiber orientations (Pierpaoli and Basser, 1996). Furthermore, there are other very good reasons to make such assumptions. First, even relatively large variations in anisotropy do not change the shape of the DW signal as much as might have been anticipated, as illustrated in Fig. 8.13. This means that even if anisotropy was not constant, its effect on our estimate of fiber orientation will be negligible. Another reason is that it drastically improves the stability of the parameter estimates: for example, a small change in the intrinsic anisotropy of one fiber bundle causes changes to the measured DW signal that look almost identical to changes in its volume fraction; by making anisotropy constant, the volume fraction estimate becomes much more robust.

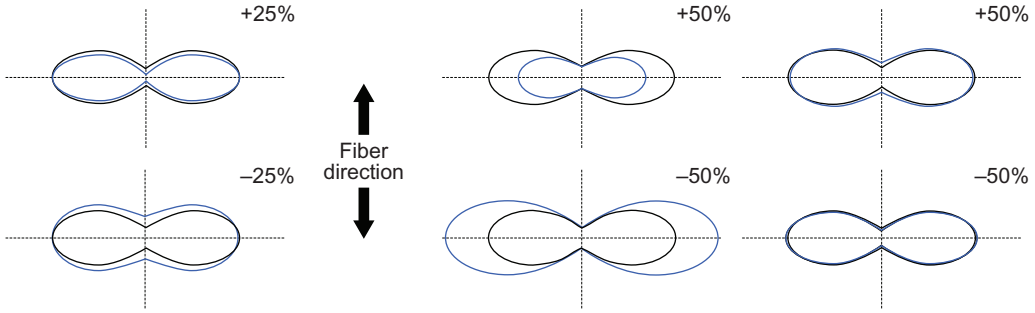

FIGURE 8.13 The impact of small variations in axial (left) and radial (right) diffusivity on the expected DW signal. Changes in axial diffusivity of the order of 25% lead to a fattening or thinning of the DW signal along the fiber axis. While such changes can be detected, the overall shape of the signal remains comparable, and the effect of the estimated fiber orientation is therefore relatively minor. Changes in radial diffusivity have a large effect on the amplitude of the DW signal in the plane perpendicular to the fiber axis (middle column). However, once scaled down to match their original size, it is clear that the overall shape of the DW signal is largely unaffected. This means that changes in radial diffusivity are essentially indistinguishable from changes in the partial volume fraction of the corresponding fiber bundle, while the estimated fiber orientation is unaffected.

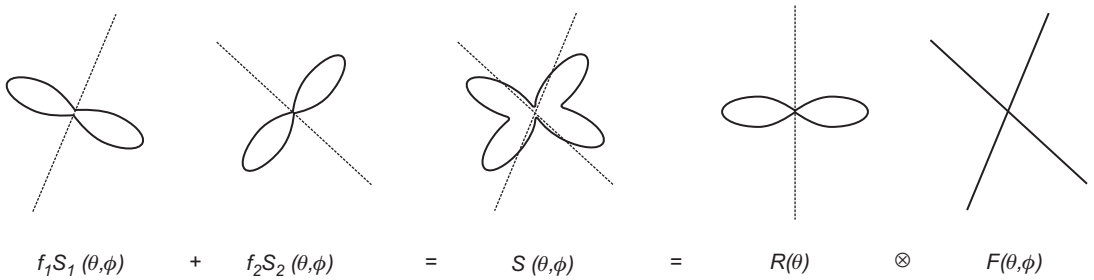

FIGURE 8.14 If we can assume that two fiber bundles with distinct orientation have the same intrinsic DW signal, rotated to line up with the fiber axis (first and second column), the DW signal that would be measured for the combination of the two bundles will simply be the sum of their respective DW signals (center). This simple observation is the central concept behind approaches based on a mixture model. This relationship can be extended to the case of a general distribution, the fiber orientation density function (fODF, denoted $F(\theta,\varphi)$ in this example), based on the concept of a spherical convolution, using the per-fiber bundle DW signal as the convolution kernel (otherwise known as the response function). *(Images reproduced from Tournier et al., NeuroImage 23:1176–1185, 2004, with permission.)*

The simplest mixture model algorithm is the **multi-tensor fitting** approach: the DW signal for each discrete fiber orientation is modeled according the diffusion tensor model, and these are then added up to give the expected DW signal, as illustrated in Fig. 8.14. More formally, the DW signal $S(\hat{u}, b)$ measured with diffusion sensitization direction $\hat{u}$ with b-value b is given by:

$$S(\underline{\hat{u}}, b) = S_0 \sum_i^N f_i e^{-b\underline{\hat{u}}^T \boldsymbol{D}_i \underline{\hat{u}}}$$

where S_0 is the non-diffusion-weighted signal intensity and N is the number of fiber orientations included in the model, each of which is characterized by its volume fraction f_i and intrinsic diffusion tensor $\boldsymbol{D}_i$. The problem is typically solved using optimization approaches, to identify the set of parameters $\{S_0, N \times [f_i, \boldsymbol{D}_i]\}$ that best fit the measured data. As written, this model includes $7N$ parameters, i.e. 14 parameters for the two fiber orientation case. While solving this problem is theoretically possible, the results of the fitting process are unstable, for reasons outlined previously. To obtain a reliable fit, most implementations will restrict the individual tensors to be axially symmetric (i.e. $\lambda_2 = \lambda_3$), and typically also fix all the eigenvalues, so that the tensors all have the same intrinsic anisotropy. For example, in the ball and sticks model (Behrens et al., 2007), the anisotropy is modeled as FA = 1, corresponding to $\lambda_2 = \lambda_3 = 0$; in other words, zero radial diffusivity. While this is clearly not biologically accurate, the benefits of simplifying the model and the vast improvement in the robustness of the model fit far outweigh the relatively small errors introduced by this approximation.

Related methods include more advanced models such as the **composite hindered and restricted model of diffusion** (CHARMED) (Assaf et al., 2004; Assaf and Basser, 2005). This models the DW signal from each fiber bundles using a more biologically realistic model of restricted diffusion in cylinders, and also

includes a conventional diffusion tensor component to model the extra-cellular compartment. This model unfortunately requires a relatively lengthy multi-shell HARDI sequence to obtain enough information for the model fit, which is still reasonable for research studies, but not for routine clinical use.

One limitation of these multi-compartment fitting approaches is the need to estimate N, the number of fiber orientations to include in the model for each voxel. As can be expected, the results will differ substantially when the wrong number if used. In some approaches, an estimate of this number is estimated in a prior pre-processing step (Parker and Alexander, 2003), other approaches use model comparison techniques to find which number provides the best fit to the data (Hosey et al., 2005, 2008), and others use an automatic relevance detection framework to drive fiber bundles with negligible volume fractions to zero (Behrens et al., 2007).

These multi-compartment approaches can be generalized to represent the fiber orientation information not as a discrete set, but as a continuous fiber orientation distribution (FOD – see above), as illustrated in Fig. 8.14. The DW signal is then no longer modeled as the sum of each (discrete) compartment, but by the spherical convolution of the FOD with the expected DW signal for a single fiber orientation. The process of extracting the FOD then consists of performing the inverse **spherical deconvolution** (Tournier et al., 2004).

Spherical deconvolution (like all deconvolution methods) is very sensitive to noise effects. A number of methods have been proposed to improve the robustness of the results, most based on a non-negativity constraint (Tournier et al., 2007; Jian and Vemuri, 2007; Dell'Acqua et al., 2007). With these improvements, this approach is robust enough to provide high quality fiber orientation estimates from relatively modest, clinically feasible HARDI acquisitions (typically <10 minutes), as shown in Fig. 8.5.

8.4 CAN HARDI PROVIDE MORE INFORMATION THAN JUST THE FIBER ORIENTATIONS?

The aim of most higher-order models proposed to date is to provide more robust estimates of the fiber orientations. In many cases, these models make the explicit assumption that all fiber bundles look alike when measured using diffusion MRI. As discussed previously in Fig. 8.13, this is a good approximation when estimating fiber orientations, but does not provide an alternative scalar measure of "white matter integrity." The question now is: can we derive any microstructural information related to each fiber bundle using these methods?

The first approach proposed to characterize anisotropy in HARDI data was the generalized FA (GFA) (Tuch, 2004). This is essentially the normalized variance of the diffusion ODF as computed using Q-ball imaging (although other ODFs could equally be used). The idea is that an isotropic ODF will have little or no variance over orientations, whereas a strongly directional ODF will have large variance. While simple to understand and compute, it is not a fully quantitative measure, in that the ODF will depend on the b-value used, the amount of noise in the data, and the particular parameters of the reconstruction. More importantly, it is not easily interpreted from a biological point of view, since as in the tensor case variations in GFA are due predominantly by crossing fiber effects.

Another approach that is becoming popular is **diffusion kurtosis imaging (DKI)**[4]. The idea behind this method is to estimate how non-Gaussian the diffusion actually is – in simple terms, how poor is the diffusion tensor model at modeling the DW signal within each voxel? The diffusion tensor model is based on the assumption of free diffusion, characterized by a normally-distributed (Gaussian) spin propagator. Kurtosis is a general measure of the deviation of a distribution from the normal distribution, and is therefore the logical metric to quantify non-Gaussian diffusion. In essence, it adds an extra quadratic term to the otherwise linear relationship between log-signal and b-value, and so quantifies non-monoexponential behavior based on the amplitude of this second-order term. This is clarified in Fig. 8.15. DKI requires that data be acquired using a multi-shell HARDI sequence, although a limited number of directions can typically be used, making it suitable for routine clinical practice. One of the outstanding issues with DKI is that the values produced are still sensitive to crossing fibers effects. Nonetheless, it has the potential to provide information that would not otherwise be available, and in some cases this may prove diagnostically useful, for example in assisting with tumor grading (Van Cauter et al., 2012).

More advanced methods attempt to estimate biologically relevant parameters directly. Recent work has focused on estimating **axon diameters and densities** (Assaf et al., 2008; Alexander et al., 2010), something that would clearly be of neurological interest. The problem to overcome with these methods is that the DW signal is relatively insensitive to changes in axonal diameters. To overcome this, multi-shell HARDI approaches are used with custom diffusion gradient waveforms, with correspondingly long scan times. Despite this, it remains very difficult to obtain robust results within a realistic timeframe, and the axon diameter estimates tend to be higher than expected.

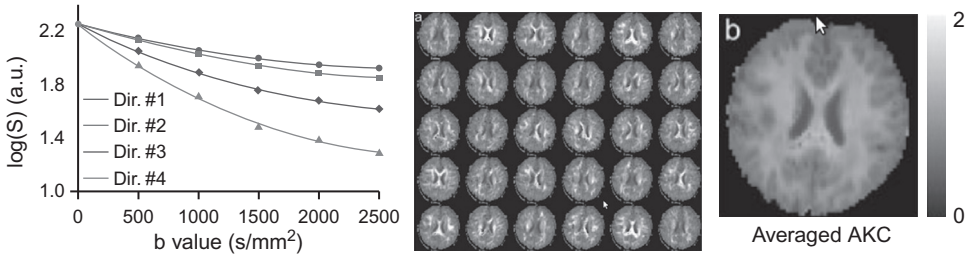

FIGURE 8.15 Diffusion kurtosis imaging (DKI) attempts to characterize the deviation of the DW signal from what would expected for Gaussian (free) diffusion. Essentially, this means that the relationship between the log of the DW signal and the *b*-value, typically assume to be linear in DTI, is extended to include a quadratic term (left). The amplitude of this quadratic term provides an estimate of the kurtosis, and can be used to produce maps of the diffusion kurtosis for each DW direction (center). A map of the average kurtosis can also be generated by averaging over all DW directions (right). *(Images reproduced from Lu et al.*, NMR Biomed. *19: 236–247, 2006, with permission.)*

Furthermore, these approaches are generally only applicable in single fiber regions such as the corpus callosum, since the presence of crossing fibers makes the problem to all intents and purposes intractable.

Another recent approach is to estimate the **apparent fiber density (AFD)**, by exploiting the insensitivity of diffusion MRI to axonal diameters (Raffelt et al., 2012; Dell'Acqua et al., 2012). It is relatively simple to show in simulations that the FOD amplitude as estimated by spherical deconvolution (or partial volume fraction when using multi-tensor approaches) is proportional to the intra-axonal volume fraction of the fiber aligned along that direction. This relationship is valid particularly for *b*-values $\geq 3000\ \mathrm{s/mm^2}$, and the relatively long diffusion times typically used on clinical systems. Based on this observation, it becomes possible to obtain independent measures of the AFD (i.e. the FOD amplitude) for each fiber orientation within a voxel. Coupled with advanced image registration and FOD reorientation and modulation strategies, this makes it possible to perform voxel-based analyses of AFD, and observe tract-specific group differences even within crossing fiber regions. While still relatively immature, such methods may in the future provide valuable information in a way that is robust to crossing fiber effects.

References

Aganj I, Lenglet C, Sapiro G, Yacoub E, Ugurbil K, Harel N. Reconstruction of the orientation distribution function in single- and multiple-shell q-ball imaging within constant solid angle. Magn Reson Med 2010;64:554–66.

Alexander AL, Hasan KM, Lazar M, Tsuruda JS, Parker DL. Analysis of partial volume effects in diffusion-tensor MRI. Magn Reson Med 2001;45:770–80.

Alexander DC, Barker GJ. Optimal imaging parameters for fiber-orientation estimation in diffusion MRI. NeuroImage 2005;27:357–67.

Alexander DC, Hubbard PL, Hall MG, Moore EA, Ptito M, Parker GJM, et al. Orientationally invariant indices of axon diameter and density from diffusion MRI. Neuroimage 2010;52:1374–89.

Assaf Y, Freidlin RZ, Rohde GK, Basser PJ. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. Magn Reson Med 2004;52:965–78.

Assaf Y, Basser PJ. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. NeuroImage 2005;27:48–58.

Assaf Y, Blumenfeld-Katzir T, Yovel Y, Basser PJ. Axcaliber: a method for measuring axon diameter distribution from diffusion MRI. Magn Reson Med 2008;59:1347–54.

Beaulieu C. The basis of anisotropic water diffusion in the nervous system – a technical review. NMR Biomed 2002;15:435–55.

Behrens TEJ, Berg HJ, Jbabdi S, Rushworth MFS, Woolrich MW. Probabilistic diffusion tractography with multiple fibre orientations: what can we gain? NeuroImage 2007;34:144–55.

Callaghan PT. NMR imaging, NMR diffraction and applications of pulsed gradient spin echoes in porous media. Magn Reson Imaging 1996;14:701–9.

Clark CA, Le Bihan D. Water diffusion compartmentation and anisotropy at high b values in the human brain. Magn Reson Med 2000;44:852–9.

Dell'Acqua F, Rizzo G, Scifo P, Clarke RA, Scotti G, Fazio F. A model-based deconvolution approach to solve fiber crossing in diffusion-weighted MR imaging. IEEE Trans Biomed Eng 2007;54:462–72.

Dell'Acqua F, Simmons A, Williams SCR, Catani M. Can spherical deconvolution provide more information than fiber orientations? Hindrance modulated orientational anisotropy, a true-tract specific index to characterize white matter diffusion. Human Brain Map 2012. http://dx.doi.org/10.1002/hbm.22080.

Descoteaux M, Angelino E, Fitzgibbons S, Deriche R. Regularized, fast, and robust analytical Q-ball imaging. Magn Reson Med 2007;58:497–510.

Douaud G, Jbabdi S, Behrens TEJ, Menke RA, Gass A, Monsch AU, et al. DTI measures in crossing-fibre areas: increased diffusion anisotropy reveals early white matter alteration in MCI and mild Alzheimer's disease. Neuroimage 2011;55:880–90.

Ebeling U, Reulen HJ. Subcortical topography and proportions of the pyramidal tract. Acta Neurochir (Wien) 1992;118:164–71.

Farquharson S, Tournier J-D, Calamante F, Fabinyi G, Schneider-Kolsky M, Jackson GD, et al. White matter fibre tractography: why we need to move beyond DTI. J Neurosurg 2012;118(6):1367–77.

Hess CP, Mukherjee P, Han ET, Xu D, Vigneron DB. Q-ball reconstruction of multimodal fiber orientations using the spherical harmonic basis. Magn Reson Med 2006;56:104–17.

Hosey T, Williams G, Ansorge R. Inference of multiple fiber orientations in high angular resolution diffusion imaging. Magn Reson Med 2005;54:1480–9.

Hosey TP, Harding SG, Carpenter TA, Ansorge RE, Williams GB. Application of a probabilistic double-fibre structure model to diffusion-weighted MR images of the human brain. Magn Reson Imaging 2008;26:236–45.

Jansons KM, Alexander DC. Persistent angular structure: new insights from diffusion magnetic resonance imaging data. Inverse Probl 2003;19:1031–46.

Jensen JH, Helpern JA, Ramani A, Lu H, Kaczynski K. Diffusional kurtosis imaging: the quantification of non-gaussian water diffusion by means of magnetic resonance imaging. Magn Reson Med 2005;53:1432–40.

Jeurissen B, Leemans A, Tournier J-D, Jones DK, Sijbers J. Investigating the prevalence of complex fiber configurations in white matter tissue with diffusion magnetic resonance imaging. Human Brain Map 2012. http://dx.doi.org/10.1002/hbm.22099.

Jian B, Vemuri BC. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. IEEE Trans Med Imaging 2007;26:1464–71.

Jones DK, Knösche TR, Turner R. White matter integrity, fiber count, and other fallacies: the do's and don'ts of diffusion MRI. NeuroImage 2012;73:239–54.

Kinoshita M, Yamada K, Hashimoto N, Kato A, Izumoto S, Baba T, et al. Fiber-tracking does not accurately estimate size of fiber bundle in pathological condition: initial neurosurgical experience using neuronavigation and subcortical white matter stimulation. NeuroImage 2005;25:424–9.

Özarslan E, Shepherd TM, Vemuri BC, Blackband SJ, Mareci TH. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). NeuroImage 2006;31:1086–103.

Parker GJM, Alexander DC. Probabilistic Monte Carlo based mapping of cerebral connections utilising whole-brain crossing fibre information. Inf Process Med Imaging 2003;18:684–95.

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

Raffelt D, Tournier J-D, Rose S, Ridgway GR, Henderson R, Crozier S, et al. Apparent Fibre Density: a novel measure for the analysis of diffusion-weighted magnetic resonance images. Neuroimage 2012;59:3976–94.

Stanisz GJ, Szafer A, Wright GA, Henkelman RM. An analytical model of restricted diffusion in bovine optic nerve. Magn Reson Med 1997;37:103–11.

Tournier JD, Calamante F, Gadian DG, Connelly A. Direct estimation of the fiber orientation density function from diffusion-weighted MRI data using spherical deconvolution. NeuroImage 2004;23:1176–85.

Tournier JD, Calamante F, Connelly A. Robust determination of the fibre orientation distribution in diffusion MRI: non-negativity constrained super-resolved spherical deconvolution. NeuroImage 2007;35:1459–72.

Tournier JD, Yeh C-H, Calamante F, Cho K-H, Connelly A, Lin C-P. Resolving crossing fibres using constrained spherical deconvolution: validation using diffusion-weighted imaging phantom data. NeuroImage 2008;42:617–25.

Tuch DS, Reese TG, Wiegell MR, Makris N, Belliveau JW, Wedeen VJ. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. Magn Reson Med 2002;48:577–82.

Tuch DS. Q-ball imaging. Magn Reson Med 2004;52:1358–72.

Tuch DS, Salat DH, Wisco JJ, Zaleta AK, Hevelone ND, Rosas HD. Choice reaction time performance correlates with diffusion anisotropy in white matter pathways supporting visuospatial attention. Proc Natl Acad Sci USA 2005;102:12212–7.

Van Cauter S, Veraart J, Sijbers J, Peeters RR, Himmelreich U, De Keyzer F, et al. Gliomas: diffusion kurtosis MR imaging in grading. Radiology 2012;263:492–501.

Wedeen VJ, Hagmann P, Tseng W-YI, Reese TG, Weisskoff RM. Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. Magn Reson Med 2005;54:1377–86.

Zhan W, Yang Y. How accurately can the diffusion profiles indicate multiple fiber orientations? A study on general fiber crossings in diffusion MRI. J Magn Reson 2006;183:193–202.