19

Brain Connectivity Mapping and Analysis Using Diffusion MRI

Brian G. Booth and Ghassan Hamarneh

CONTENTS

19.1 Diffusion-Weighted Image Acquisition

Diffusion magnetic resonance imaging (dMRI) is a powerful imaging protocol that allows for the assessment of the organization and integrity of fibrous tissue. The imaging works by measuring the diffusion of water molecules within the body. This diffusion is restricted by cell membranes and as such, rates of diffusion are far less in directions perpendicular to fibrous tissue than parallel to the fibers. With enough diffusion measurements along different directions in 3D, we can noninvasively obtain a profile of the diffusion at various points within the imaged subject.

The diffusion profiles obtained from dMRI have had a significant impact on the analysis of neural connectivity within the white matter of the brain. Neural pathways, dubbed fiber tracts, can be traced out using the directional information from the diffusion profiles. This process is known as tractography and, due to noise, motion artifacts, and partial voluming effects, is a computationally difficult problem.

We present here an examination of the current state of tractography and dMRI. In particular, we look at the computational challenges inherent in this area and the open problems that remain.

19.1.1 Biological Basis for Diffusion MRI

The biological basis for dMRI dates back to 1828 when botanist Robert Brown noticed the continuous and random motion of pollen grains suspended in water [29]. What Brown had discovered was later determined to be the motion of water molecules due to thermal agitation [58]. This motion, now known as Brownian motion or diffusion, was later characterized by A. Einstein [41], resulting in Einstein's equation:

$$r^2 = 6dt \tag{19.1}$$

What Einstein's equation characterized was that the square of the average displacement of molecules (r) with a given diffusion rate (d) is proportional to the observation time (t). If we can measure this molecular displacement over a fixed time, we can obtain the diffusion rate of different substances under different conditions.

As the majority of the human body is water [46], the diffusion phenomenon occurs within us as well. While the diffusion process is random, our cell structures can restrict or hinder the motion of water molecules [19]. As such, the diffusion of water molecules in our body depends on the microstructure of our tissues. Fast molecular diffusion occurs within and around a cell as there are few microstructures to inhibit motion. Diffusion through the cell however is slower as the cell membrane and other structures (e.g., myelin sheaths in the brain white matter) restrict molecular motion.

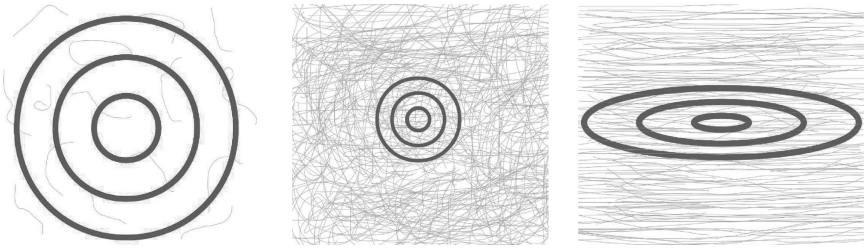

FIGURE 19.1
Synthetic examples of the diffusion seen in CSF (a), gray matter (b), and white matter (c) within the brain. The diffusion rates for various directions are shown in bold. (Adapted from Alexander, D.C., *Visualization and Processing of Tensor Fields*, Chapter 5: an introduction to computational diffusion MRI: the diffusion tensor and beyond, Springer Berlin Heidelberg, Heidelberg, Germany, pp. 83–106, 2006.)

Since the diffusion of water within the body is dependent on local cell structure, we can discuss how different organizations of these structures affect diffusion rates. Consider, for example, the human brain where functional regions (gray matter) are connected by a collection of neural pathways (white matter). Figure 19.1 presents diffusion measures for the brain's corticospinal fluid (CSF), gray matter, and white matter, respectively. When the cell structure is minimal as in CSF, we see fast isotropic diffusion. More complex cell structure that is not consistently organized, such as gray matter, shows slower, but still isotropic, diffusion. Yet if the local cell structure is organized in a consistent orientation, as it is in white matter, the diffusion rates become anisotropic, that is, they vary with regard to direction [19].

These diffusion differences within the brain are potentially useful cues in analyzing brain structure and function. For example, measuring the average diffusion rate or the anisotropy of a tissue can give us significant information about the tissue's organization and integrity [76]. Diffusion measurements would be most informative in white matter regions where the orientation of the microstructure can be inferred from the diffusion. This microstructural orientation within the brain's white matter is in turn known to describe the direction of neural pathways in the brain [20]. As a result, by measuring the diffusion using Einstein's equation, we could infer the orientational structure of the brain's white matter and in turn map out the brain's neural pathways. This is precisely what dMRI is used to accomplish.

19.1.2 Diffusion-Weighted Image Acquisition

To understand how diffusion can be measured through MRI, we must first address the basic concepts on nuclear magnetic resonance. All elementary

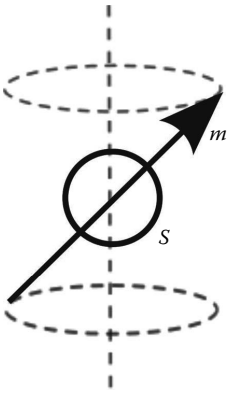

FIGURE 19.2
Nuclear spin *s* generating a magnetic moment *m*. The particle spins around a rotational axis shown here in gray. (Adapted from Lenglet, C. et al., Diffusion tensor magnetic resonance imaging: Brain connectivity mapping. Technical Report 4983, INRIA, October 2003.)

particles possess a physical property of spin (*s*) as seen in Figure 19.2. This spin property rotates the particle around its nucleus, thereby giving the particle a magnetic moment (*m*). This magnetic moment can then be manipulated using nuclear magnetic resonance. As the body is mostly water, the spins of hydrogen atoms within water molecules become a good candidate for MRI.

19.1.2.1 Magnetic Resonance Imaging

MRI is comprised of three principal steps: precession, resonance, and relaxation. We consider each in turn.

Precession: A static magnetic field $\mathbf{B}_0$ is applied to the body. This magnetic field aligns the rotational axis of each spin with its field direction. These spins now rotate, or precess, around the same magnetic axis. Note that roughly an equal number of spins will be aligned with the positive direction of the magnetic axis as with the negative direction and the overall signal generated during precession will be minimal.

Resonance: With the magnetic field $\mathbf{B}_0$ in place, a second, weaker, magnetic pulse is applied to the body in the direction $\mathbf{g}$. This second field results in the magnetic moment *m* of each spin aligning with the pulse direction $\mathbf{g}$. The spin's axis of rotation remains aligned with $\mathbf{B}_0$. The result of the resonance phase is to cause the net magnetism of the spin to veer away from the main magnetic field $\mathbf{B}_0$. This change in the magnetic field induces a small current within the subject.

Relaxation: The second magnetic pulse is removed and the magnetic moments of the hydrogen atoms realign with $\mathbf{B}_0$. As this realignment occurs, the changing magnetic field generated by the realignment of the spins induces

a current in the coil of the MRI scanner. From this current, two common measurements can be taken:

1. Spin–spin relaxation time (T_2): The amount of time it takes for the magnetism in the direction of **g** to reduce to 37% of its original value
2. Spin–lattice relaxation time (T_1): The amount of time it takes for the magnetism in the direction of $\mathbf{B}_0$ to recover 63% of the magnetism it lost when the second gradient was applied in the direction **g**

These relaxation times can be visualized at multiple locations in the brain, resulting in what are known as T_1 and T_2 weighted images.

19.1.2.2 Diffusion-Weighted Imaging

As magnetic resonance imaging depends on the magnetic moments of hydrogen atoms, Stejskal and Tanner were able to develop a sequence of precession, resonance, and relaxation periods that allow MRI to measure the movement of hydrogen atoms over time and in turn the water molecules of which they are a part [95]. Le Bihan and Breton later adapted this MR image sequence to the scanning of the human body [24]. This imaging sequence is summarized in Figure 19.3 for a given angular direction **g**.

The sequence in Figure 19.3 assumes that the magnetic field $\mathbf{B}_0$ has been applied and that the spins are precessing around $\mathbf{B}_0$. In this state, a magnetic pulse is applied at an angle of 90° from the direction of $\mathbf{B}_0$. This pulse aligns the spins that were separately aligned to either the positive or negative $\mathbf{B}_0$ axis. Once the spins are aligned, the 90° pulse is removed and a second pulse, known as a gradient pulse, is applied in the direction **g**. This gradient pulse senses the induced current to a specific angular direction.

A third magnetic pulse in the direction 180° from $\mathbf{B}_0$ follows the gradient pulse, then the gradient pulse is reapplied. The 180° pulse plays a key role in that it flips the spin direction of the atoms to the opposite of what they were during the precession phase. As a result of this flip, the current induced by stationary atoms during the application of the second gradient pulse will cancel out the current induced by the same atoms during the first gradient pulse [67]. Therefore, the resulting signal measured after all gradient pulses have been applied relates solely to the molecules experiencing motion in the direction **g**.

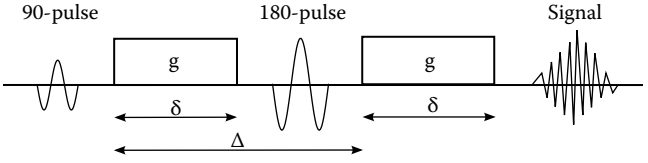

FIGURE 19.3
The Stejskal–Tanner diffusion-weighted imaging sequence. (Adapted from Westin, C.-F. et al., *Med. Image Anal.*, 6, 93, 2002.)

The T_2 relaxation time is then measured from this final signal for multiple locations in the brain and visualized in what are known as diffusion-weighted images (DWIs). Figure 19.4 displays a conventional T_2 image next to sample DWIs for various gradient directions **g**. Note here that rapid diffusion results in fast T_2 relaxation times, resulting in a low intensity in the DWI.

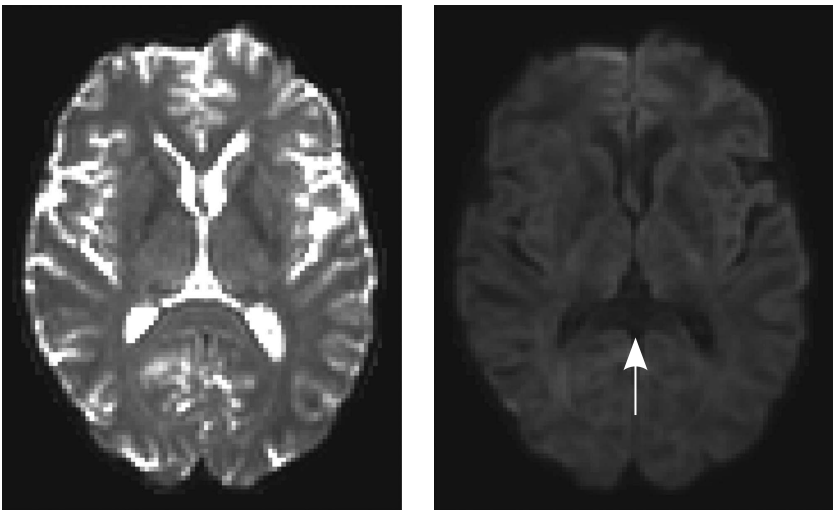

(a)

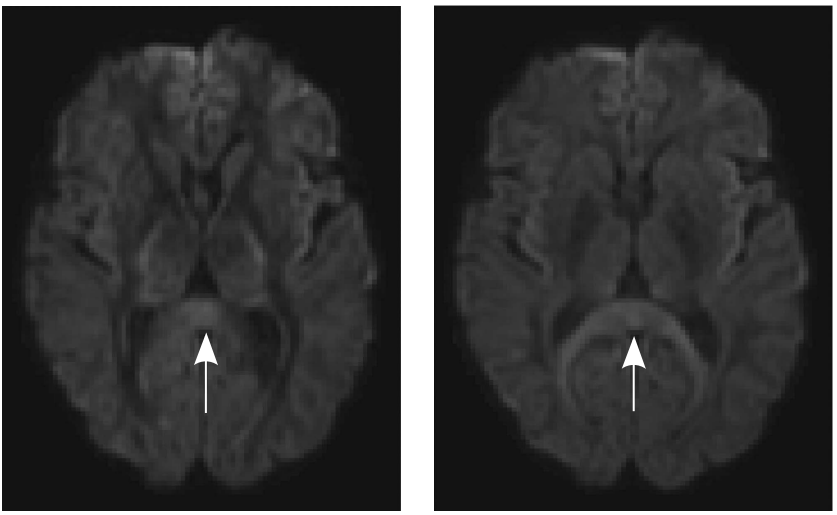

(b)

FIGURE 19.4
Axial slices of (left to right) (a) standard T_2 image and (b) its corresponding DWIs from gradient pulses in the horizontal, vertical, and out-of-plane directions. Note the differences in measured diffusion in the splenium due to gradient direction (highlighted by the white arrows). (Adapted from Jones, D.K., *Cortex*, 44(8), 936, September 2008.)

Further note the different rates of diffusion for different directions within the brain's white matter as pointed out by the white arrows in Figure 19.4.

From the DWI for gradient direction **g**, the diffusion rate (d) can be computed using the Stejskal–Tanner equation:

$$S = S_0 \exp(-bd) \tag{19.2}$$

where
S is the DWI intensity
S_0 is the standard
T_2 is the image intensity
b is the diffusion weighting [95]

The diffusion weighting b is in turn proportional to the strength and duration of the gradient pulse. The T_2 image used in (19.2) is typically referred to in this context as a B_0 image as it is acquired without the application of the gradient pulses (i.e., $b = 0$). The scalar d is commonly referred to as the apparent diffusion coefficient (ADC).

19.1.3 Correction of Image Artifacts

To obtain a full understanding of dMRI, we must acknowledge how the quality of the DWIs is affected or limited by the image acquisition process. All further analysis is going to depend on the accuracy of these diffusion measurements, and as such, we must address the presence of noise and imaging artifacts within these DWIs.

dMRI is susceptible to various artifacts, the three most common being eddy currents, subject motion, and Rician noise [13]. Let us consider each in turn.

19.1.3.1 Eddy Currents

As seen in the diffusion imaging sequence in Figure 19.3, multiple magnetic gradient pulses are applied in rapid succession. Switching between these gradients can result in fluctuations in the scanner's magnetic field. These fluctuations induce what are known as eddy currents in the coil of the MRI scanner. The eddy currents interfere with the currents induced by the scanned subject, thereby distorting the resulting DWIs [13].

Much is known of eddy currents, namely, that they are dependent on the magnitude of the gradient pulse, independent of the subject being scanned, and that they result in related geometric and intensity distortions in DWIs [51]. The geometric distortion produced from eddy currents has been shown to consist of a translation, scale, and shear of the resulting image and is commonly rectified using affine registration [27,51,70]. The DWIs are registered to a T_2-weighted image with the mutual information similarity measure showing the best results [70]. As the T_2 image is acquired

without gradient pulses that produce eddy currents, it is assumed to be free of geometric distortion, thereby making it an appropriate template to which we can register the DWIs. Intensity corrections are then calculated directly from the magnitudes of the shear, scaling, and translations of the affine warp [51,70].

One benefit of eddy currents being independent of the subject scanned is that the affine warp used in the correction can be obtained by imaging a physical phantom with known ground truth [34]. This warp can then be applied to later subject scans.

19.1.3.2 Subject Motion

Depending on the number of DWIs being acquired, the length of a dMRI scan can range from a couple of minutes [74] to a few hours [103]. During that time, the subject may move both voluntarily involuntarily (e.g., breathing). As a result, the same voxel location in two DWIs is not guaranteed to correspond to the same anatomical location in the subject.

While correcting for subject motion in a single image has been well studied (see [98] for a survey), the problem of correcting motion between separate DWIs has yet to receive a strong theoretical treatment [13]. Even so, two main approaches have been proposed to correct for subject motion between DWIs, both involving image registration. First, we can, as with eddy current correction, align the DWIs to a T_2-weighted image with the mutual information similarity measure [70,90]. The alternative approach is to model the diffusion at each voxel (as discussed further in Section 19.2) and to align images so as to minimize the residual of the model fit [7,10]. A recent quantitative comparison of these approaches suggests that both methods are equally capable of correcting for subject motion [91].

Note however that if the rotational motion of the subject is large, the directions of the applied gradient pulses would need to be corrected as well [66,90].

19.1.3.3 Rician Noise

Any environment is going to contain a certain amount of background noise. In the case of dMRI, this noise has been well modeled using a Rician distribution [16] given as

$$p(x \mid \mu, \sigma) = \frac{x}{\sigma^2} \exp\left(-\frac{x^2 + \mu^2}{2\sigma^2}\right) I_0\left(\frac{x\mu}{\sigma^2}\right) \tag{19.3}$$

where

x is the observed image intensity
μ is the noise-free signal
σ is the standard deviation of the noise
I_0 is the zeroth-order Bessel function of the first kind

At high signal to noise ratios, the Rician distribution is occasionally approximated using a Gaussian distribution [43]. This additive noise can have an adverse effect on the diffusion rates calculated using the Stejskal–Tanner equation, particularly for images taken at a high diffusion weighting [60].

Historically, variational methods have been applied to remove this Rician noise, with anisotropic filtering [82] and total variation regularization [16,43], both showing success. Weighted-mean filtering approaches have also been used [97,107]. The main conceptual difference in noise removal in DWIs is whether to denoise one image at a time (the scalar approach) or all images at once (the vector approach) [43]. Recent results suggest that the vector-based algorithms improve signal to noise to a greater extent [97].

19.2 Modeling Local Diffusion Patterns

Since the introduction of dMRI, two key advancements have propelled the field to where it is today: first, the introduction of the diffusion tensor by Basser et al. [14] and second, the introduction of higher angular resolution diffusion imaging (HARDI) [101]. The introduction of the diffusion tensor brought forth the concept of modeling the diffusion rates from the DWIs as a 3D function within each voxel. HARDI, on the other hand, allowed us to increase the complexity of these models to better represent the local diffusion properties. This section will show how these two contributions underlie the ability to perform brain connectivity analysis.

19.2.1 Diffusion Tensor Model

While we have shown that the Stejskal–Tanner equation (19.2) relates diffusion rates to the DWI intensities, we can consider a more general formulation of the diffusion properties at a voxel. Since water molecules undergo random Brownian motion, we can consider a probability density function (PDF) $p_t(\mathbf{x})$ describing the probability of a water molecule experiencing a displacement $\mathbf{x}$ over the observation time t. It has been shown that the distribution p_t is related to the DWI intensities via the Fourier transform [5].

$$\frac{S(\mathbf{g})}{S_0} = \int_{\mathbf{x}\in\mathbb{R}^3} p_t(\mathbf{x}) \exp\left(-ib\mathbf{g}\mathbf{x}\right) d\mathbf{x} \tag{19.4}$$

As, $S(\mathbf{g})$ represents the diffusion-weighted signal for the gradient direction $\mathbf{g}$, S_0 is the unweighted B_0 image signal, and b is the diffusion weighting. With enough DWIs $S(\mathbf{g})$, the Fourier transform can be inverted to obtain p_t. This is known as *q-space* imaging [103]. In practice however, the number of DWIs

required to accurately perform the inversion leads to scanning times on the order of hours [54] that is generally not available in a clinical setting. As a result, it has become common to assume a model for p_t, the simplest model being a zero-mean Gaussian:

$$p_t(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^3 \mid 2t\mathbf{D} \mid}} \exp\left(\frac{-\mathbf{x}^T\mathbf{D}^{-1}\mathbf{x}}{4t}\right) \tag{19.5}$$

where the covariance is $2t\mathbf{D}$.

Plugging the Fourier transform of (19.5) into (19.4) results in a more general case of the Stejskal–Tanner equation:

$$S(\mathbf{g}) = S_0 \exp\left(-b\mathbf{g}^T\mathbf{D}\mathbf{g}\right) \tag{19.6}$$

The 3×3 second-order positive-definite symmetric matrix $\mathbf{D}$ is referred to as the diffusion tensor [14]. It contains six unique elements and therefore six DWIs are required, along with the B_0 image, to estimate the tensor. The DWIs are obtained from uniform, noncollinear gradient directions so as to not favor a given direction in the tensor fitting process. These seven images can be obtained with an MRI scan on the order of 1–2 min [74], thereby making it a clinically feasible imaging protocol.

Many factors affect the quality of the diffusion tensors. As mentioned earlier, noise, motion, and distortions in the DWIs will result in poor tensor estimates. Aside from post-processing the DWIs, it is also common to obtain DWIs from more than six gradient directions in order to overfit the tensor, thereby reducing the effect of having some corrupted DWI signals [60]. The fitting procedure also affects the quality of the resulting tensors. The simplest approach is to take the logarithm of (19.6) and fit the tensor using least squares [14]. This approach, however, does not ensure that the resulting tensor be positive definite (i.e., have positive eigenvalues). Nonlinear fitting allows for this constraint and generally results in a less noisy tensor field [60], especially if spatial regularization is incorporated into the fitting procedure [102]. Additional approaches include using a weighted least squares fitting of the log-transformed equation (19.6) that is used to detect and remove outlier DWI signals prior to the final tensor fit [33].

19.2.2 Tensor Image Visualization

The power of the diffusion tensor lies in its ability to measure and visualize more detailed properties of the diffusion than the scalar DWIs provide. For example, we can look at how the rates of diffusion vary with direction or calculate the average diffusion rate at each voxel. In fact, significant diagnostic information can be obtained from the diffusion tensor by analyzing

its principal components obtained through the tensor's eigendecomposition [50]. Given a diffusion tensor **D**, we can obtain the eigendecomposition:

$$\mathbf{D} = \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 & \mathbf{e}_3 \end{bmatrix} \begin{bmatrix} \lambda_1 & & \\ & \lambda_2 & \\ & & \lambda_3 \end{bmatrix} \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 & \mathbf{e}_3 \end{bmatrix}^T \tag{19.7}$$

where the eigenvalues are positive and sorted in descending order (i.e., $\lambda_1 \geq \lambda_2 \geq \lambda_3$). The eigenvectors $\mathbf{e}_1$, $\mathbf{e}_2$, $\mathbf{e}_3$ are considered the main axes of diffusion, while the eigenvalues encode the rates of diffusion along each corresponding axis. Given this interpretation, we can visualize the diffusion tensor as an ellipsoid as shown in Figure 19.5. The axes of the ellipsoid are the eigenvectors of the tensor, while the tensor's eigenvalues describe the ellipsoid's stretch along each axis. Another interpretation of the ellipsoid is as an isoprobability surface of the Gaussian diffusion model given in (19.5).

The tensor eigendecomposition allows for the computation of two key diffusion properties: the mean diffusivity (MD) and the fractional anisotropy (FA) [58,105]. These two measures respectively capture the mean and variance of the diffusion rate with respect to direction. They are computed from the tensor's eigenvalues as

$$MD = \frac{1}{3}\left(\lambda_1 + \lambda_2 + \lambda_3\right) \tag{19.8}$$

$$FA = \sqrt{\frac{3}{2}} \frac{\sqrt{(\lambda_1 - MD)^2 + (\lambda_2 - MD)^2 + (\lambda_3 - MD)^2}}{\sqrt{\lambda_1^2 + \lambda_2^2 + \lambda_3^2}} \tag{19.9}$$

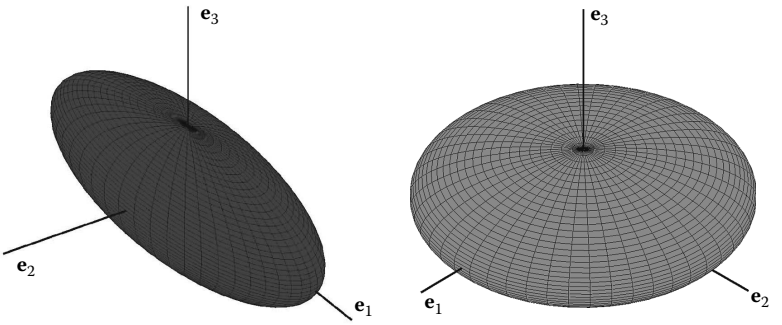

FIGURE 19.5
Examples of the ellipsoidal representation of prolate (a) and oblate (b) diffusion tensors. (Adapted from Jones, D.K., *Cortex*, 44(8), 936, September 2008.)

Examples of MD and FA on a slice of the brain are shown in Figure 19.6. From these images, we observe the brain microstructure described in Section 19.1.1. Note that the MD is higher in the ventricles than the rest of the brain due to the lack of tissue structure. Conversely, the FA is highest in the white matter of the brain due to coherent orientation of

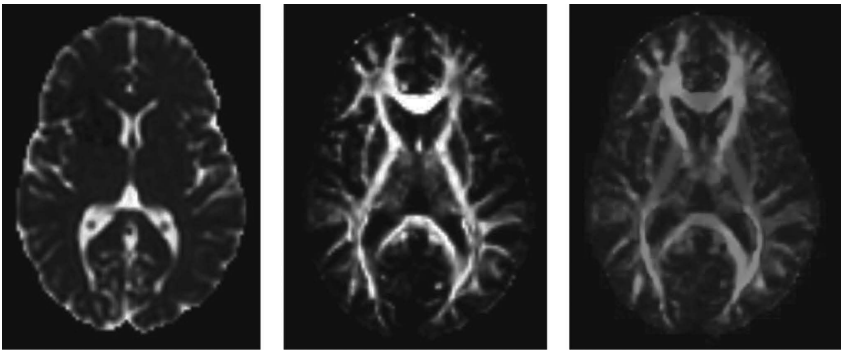

(a)

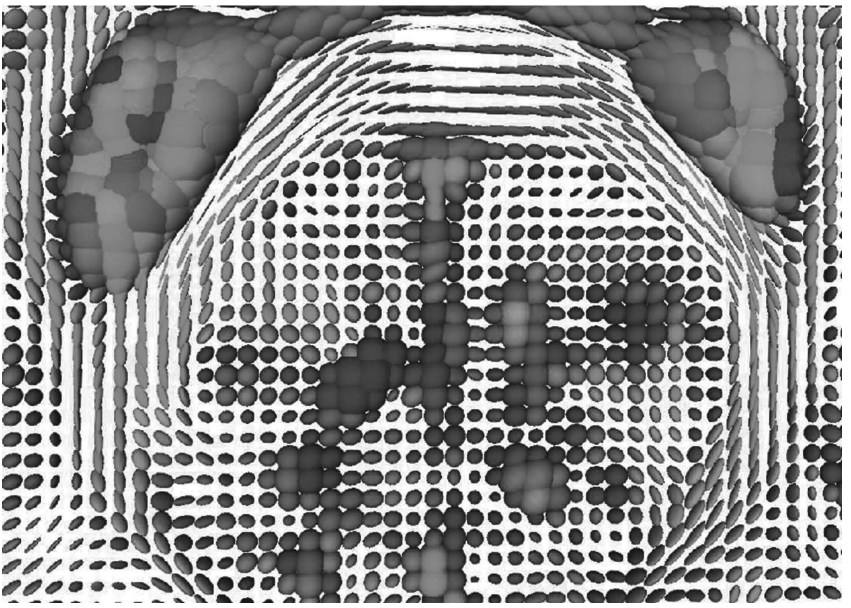

(b)

FIGURE 19.6
(See color insert.) Various methods of visualizing the information contained in a diffusion tensor field. (a) MD (left), FA (center), and color-coded orientation map (right) and (b) ellipsoidal visualization. (Images generated using MedINRIA http://www-sop.inria.fr/asclepios/software/MedINRIA/ on data obtained from Mori, S., John Hopkins Medical Institute: Laboratory of Brain Anatomical MRI, in vivo human database, http://lbam.med.jhmi.edu/, accessed February 2009.)

the tissue microstructure. We can further estimate the orientation of this microstructure as being equivalent to $\mathbf{e}_1$. Of course, the quality of this estimate will depend in part on the FA. Low FA would imply a less coherent orientation in the tissue microstructure, making the estimation of this orientation not well founded.

Other scalar measures have been generated to characterize both the shape and anisotropy of diffusion tensors, but FA and MD are most commonly used in practice. A review of other scalar tensor measures can be found in [105,114].

Aside from visualizing the scalar FA and MD maps, approaches have been developed to display the tensor's orientation information as well. The two most common approaches are shown in Figure 19.6. First, the primary diffusion direction (PDD) $\mathbf{e}_1$ can be visualized as a color image, where the RGB values are $R = FA|\mathbf{e}_1 \cdot [1,0,0]|$, $G = FA|\mathbf{e}_1 \cdot [0,1,0]|$, and $B = FA|\mathbf{e}_1 \cdot [0,0,1]|$ [80]. Such a scheme allows for an intuitive visualization of orientation weighted by the orientation strength, yet color assignments are not unique. For example, the color yellow would be assigned to voxels with $\mathbf{e}_1 = [1, 1, 0]$ and $\mathbf{e}_1 = [-1, 1, 0]$, leading to ambiguity of the underlying diffusion direction. As a result, it is occasionally necessary to visualize the tensor ellipsoids themselves as seen in Figure 19.6b. Generating a less ambiguous color representation of tensor data remains an area of open research [115].

Regardless of visualization strategy, the value of the orientation information in dMRI is significant as it allows us to infer the orientation of neural pathways in the white matter of the brain. If we take, for example, the FA in Figure 19.6, we can observe four major neural pathways. The forceps minor can be seen in the upper portion of the image arching upward in a U shape. A similar looking pathway, the forceps major, can be seen in the bottom half of the image as an inverted U shape. Flanking the forceps major on either side are the optic radiations. These pathways, as seen in dMRI, agree with histological studies [31], thereby making dMRI a powerful tool for mapping out these neural pathways noninvasively.

19.2.3 High Angular Resolution Diffusion Models

While the diffusion tensor model provides a powerful tool for visualizing and assessing the microstructure of brain tissue, it suffers from a significant limitation: the assumption that diffusion follows a Gaussian model. While this model may hold for simple examples such as those in Figure 19.1, there exist many situations where the local diffusion is more complex.

Take, for example, this situation shown in Figure 19.7a. This example shows a mixture of fibrous tissues oriented along the positive and negative diagonal directions, resulting in a diffusion profile in bold. Ideally, we would like to model this diffusion as shown in Figure 19.7b. Unfortunately, the diffusion tensor model assumes ellipsoidal Gaussian diffusion. As a result, we would obtain for this example the tensor shown in Figure 19.7c. This tensor would misleadingly suggest that diffusion is equal for all directions in the

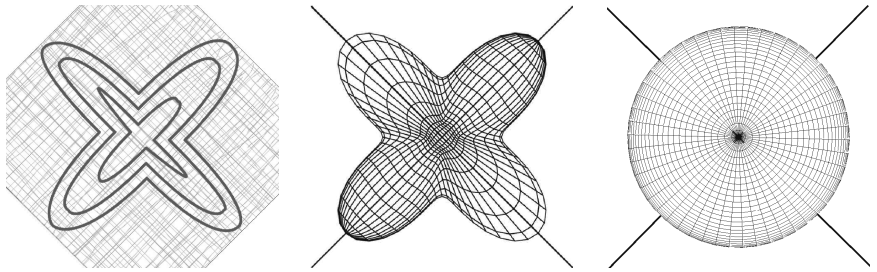

FIGURE 19.7
Example of crossing fibers and how they are modeled using dMRI. (a) Crossing fibers, (b) corresponding diffusion ODF, and (c) corresponding diffusion tensor. (Adapted from Alexander, D.C., *Visualization and Processing of Tensor Fields*, Chapter 5: an introduction to computational diffusion MRI: The diffusion tensor and beyond, Springer, Berlin Heidelberg, Germany, pp. 83–106, 2006; Descoteaux, M. et al. *IEEE Trans. Med. Imaging*, 28(2), 269, 2009; Westin, C.-F. et al., *Med. Image Anal*., 6, 93, 2002, respectively.)

plane of the crossing. Further, the primary eigenvector of the tensor is not guaranteed to align with either fiber direction.

Such an example is common in the white matter of the brain. The neural pathways are made up of aligned tissue fibers whose diameter is on the order of microns [3]. In contrast, the resolution of DWIs is typically on the order of millimeters cubed. As a result, this type of averaging of diffusion from multiple pathways is unavoidable. In fact, it has been estimated that at least one-third [21] to two-thirds [39] of voxels in the brain may exhibit this crossing fiber property. As a result, attempts have been made to come up with better diffusion models.

Tuch et al. first proposed the use of more descriptive diffusion models by showing that there are regions in the brain where fibers cross [101]. They noted that in order to detect these crossing fibers, DWIs from a greater number of diffusion directions, and at a higher gradient weighting, were required. Thus was born the concept of higher angular resolution diffusion imaging (HARDI).

Initial attempts to model more complicated diffusion profiles revolved around fitting multiple tensors to the DWI signals [101]. A mixture of Gaussian's model was assumed and the Stejskal–Tanner equation was updated to incorporate the mixture:

$$\frac{S(\mathbf{g})}{S_0} = \sum_i f_i \exp\left(-b\mathbf{g}^T\mathbf{D}_i\mathbf{g}\right) \tag{19.10}$$

Each tensor $\mathbf{D}_i$ has a corresponding volume fraction f_i representing the fraction of the local diffusion the tensor represents. Later work using the CHARMED [8,9] and FORECAST [6] methods assumed a particular shape

for each fitted tensor. The former approach attempts to model intra-fiber and extra-fiber diffusion using prolate (cigar-like) and spherical tensors, respectively. The latter approach models prolate tensors with an equal and known MD. More recent work has instead assumed a mixture of Wishart (MOW) distributions—effectively a distribution over tensors—as the choice of diffusion model [56,57].

While these mixture model approaches allow for the same intuitive representation as the single tensor model, they also have their limitations. These include the following:

- The number of tensors being fitted to the DWI signal has to be specified ahead of time. While there has been work on estimating this number from the data [5,101], there is no ground truth specification for the number of tensors to fit at a voxel.
- There is no guarantee that the assumed shape of the fitted tensors is appropriate for the underlying diffusion. If the shape assumption is poor, the volume fractions can be poorly estimated [6].
- The mixture model, and not the underlying mixture components, is fit to the DWIs. While the peaks of the mixture model will align with the directions of maximal diffusion, there is no guarantee that the peaks of the underlying distributions will align with these directions as well [101].

Some of these limitations have been addressed in recent work. For example, instead of fitting a fixed number of tensors to the data, volume fractions can be calculated for a set of basis tensors [56,57,69]. Those tensors with a volume fraction above a given threshold are maintained to model the diffusion. Also, instead of using the mixture components for further analysis, the mixture model itself is used to analyze the diffusion [6,56,57].

On the other end of the spectrum, model-free approaches have also been proposed to capture local diffusion properties. Again, Tuch instituted this approach through the introduction of *q-ball* imaging [100]. Based on the earlier *q-space* approach described by the Fourier transform in (19.4), Tuch noticed that the directional dependence of the diffusion rate is the most commonly used information for dMRI analysis and that the radial distance component of the diffusion does not play a significant role. As a result, instead of modeling the diffusion as a PDF $p(\mathbf{x})$, where $\mathbf{x}$ is a vector of any length, *q-ball* imaging models the diffusion orientation distribution function (ODF) $\psi(\theta, \varphi)$, where θ, φ are spherical angles. As such, the ODF captures the probability of diffusion along different angular directions (θ, φ) but without a radial distance parameter. An estimation of the ODF can be more efficiently obtained through the use of the Funk–Radon transform [100]. By ignoring the radial component, the *q-ball* ODF can be estimated with fewer DWI samples than the original PDF from *q-space* imaging, leading to more reasonable scanning times.

Two other model-free approaches have also gained traction in the dMRI community. First, the diffusion orientation transform (DOT) shares similarities with *q-ball* imaging as both are based on the earlier *q-space* approach. In contrast, DOT assumes diffusion decays exponentially along the radial direction and uses this assumption to perform the Fourier transform in (19.4) using fewer DWI samples [79]. The DOT diffusion ODF is then obtained by analytically integrating the resulting PDF along the radial direction.

An alternative model-free approach is Jansons and Alexander's persistent angular structure (PAS) approach [54]. The goal behind PAS is to find a diffusion PDF $p(\mathbf{x})$ from (19.4) that is smooth yet captures the key angular structure of the diffusion. This goal is achieved through optimization by finding a PDF $p(\mathbf{x})$ that maximizes entropy while minimizing the error in fitting to the DWI samples. A Lagrange multiplier is used to weight the two competing terms [54].

One of the key limitations of the model-free HARDI approaches is precisely that a model is not assumed. In areas of low anisotropy, both PAS and *q-ball* imaging can overestimate the directional dependence of the diffusion as a result of image noise [54,100]. This overestimation can result in spurious maxima in the diffusion ODFs. While work has been done in reviewing and comparing different HARDI approaches [1,3,4,87], there is generally no consensus as to which HARDI model is best suited to represent diffusion MR characteristics.

As with the diffusion tensor, the quality of the diffusion ODF is going to depend in part on the algorithm used to fit the ODF to the diffusion data. Recently, Jian and Vemuri pointed out that many of these diffusion model fitting algorithms can be unified using a spherical deconvolution framework [55]. Given DWI signals $S(\mathbf{g})$ and the B_0 image S_0, the diffusion ODF ψ can be considered as a deconvolved version of the signal:

$$\frac{S(\mathbf{g})}{S_0} = \int_{\mathcal{M}} R(\mathbf{g}, \theta, \phi)\psi(\theta, \phi) \tag{19.11}$$

where

R is the convolution kernel

$\mathcal{M}$ is a manifold, typically $\mathbb{R}^3$ in the case of a PDF and $\mathbb{S}^2$ for the ODF

With a discretization of $\mathcal{M}$, (19.11) can be converted into a linear least squares problem. However, the kernel matrix representing R is typically ill-conditioned and highly sensitive to noise in the sampled $S(\mathbf{g})$ [55]. As a result, nonlinear fitting approaches are recommended [2,55].

19.2.4 HARDI Representation and Visualization

While a diffusion PDF, such as the Gaussian PDF used with the diffusion tensor, provides information on the radial aspect of the measured diffusion,

the concept of using a diffusion ODF has become commonplace with HARDI data [55]. As such, an efficient representation of the ODF would be useful for further computation. Since the diffusion ODF is a spherical function, the most popular choice for its representation is a real spherical harmonic expansion [6,40,55]. The diffusion ODF ψ can be represented as

$$\psi(\theta,\phi) = \sum_{\ell=0}^{K} \sum_{m=-\ell}^{\ell} F_\ell^m Y_\ell^m(\theta,\phi) \tag{19.12}$$

where integers ℓ and m are the degree and order of the harmonics, respectively.

The basis harmonics Y_ℓ^m are given as

$$Y_\ell^m = \begin{cases} \sqrt{\dfrac{2\ell+1}{4\pi}} P_\ell^0(\cos(\phi)), & \text{if } m = 0 \\ \sqrt{2}\sqrt{\dfrac{2\ell+1}{4\pi}\dfrac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos(\phi))\cos(m\theta), & \text{if } m > 0 \\ \sqrt{2}\sqrt{\dfrac{2\ell+1}{4\pi}\dfrac{(\ell-m)!}{(\ell+m)!}} P_\ell^{-m}(\cos(\phi))\sin(m\theta), & \text{if } m < 0 \end{cases} \tag{19.13}$$

where P_ℓ^m is the associated Legendre function of degree ℓ and order m.

As the ODF is antipodally symmetric, only the even-degree basis harmonics are used [40]. Typically, the expansion is limited to degree $\ell \leq 8$ to suppress noise artifacts in the resulting ODF [39]. Other ODF representations have also seen limited use, including von Mises–Fisher distributions [88] and fourth-order tensors [11].

The notions of MD and FA have also been extended to HARDI diffusion ODFs, with the latter being referred to in this context as generalized anisotropy (GA). As before, the two measures correspond to the mean and variance of the diffusion ODF ψ:

$$MD = \frac{1}{4\pi}\int \psi(\theta,\phi)dS \qquad GA = \frac{1}{4\pi}\int \left(\psi(\theta,\phi) - MD\right)^2 dS \tag{19.14}$$

Unlike the diffusion tensor model, MD and GA generally do not have an elegant solution. Analytical solutions have been proposed for both measures [11,78] but involve ad hoc scaling and normalization weights. Examples of MD and GA images are shown in Figure 19.8.

Finally, we note that the visualization of the orientation information in the diffusion ODF typically involves visualizing the spherical ODFs themselves as seen in Figure 19.8.

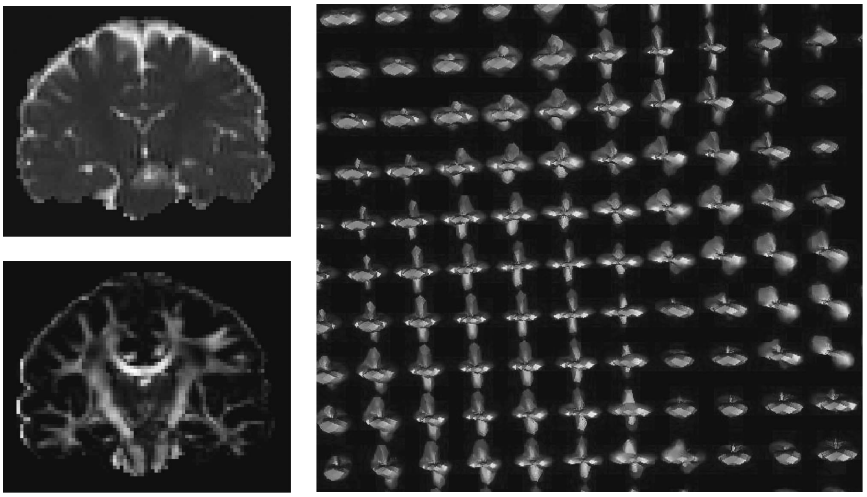

(a) (b)

FIGURE 19.8
Sample visualization techniques for diffusion ODFs obtained from HARDI. (a) MD (top) and GA (bottom) and (b) *q-ball* diffusion ODFs.

19.2.5 HARDI versus the Diffusion Tensor

Despite the presence of these HARDI models that better represent the underlying diffusion properties, the use of the diffusion tensor model still persists in a clinical setting [74–76]. There are various reasons for the use of what is perceived to be an inferior model and these reasons highlight some of the limitations of HARDI:

- The number of gradient directions, and in turn DWIs, required for the reconstruction of HARDI models is still significantly larger than for diffusion tensor imaging. With scanning time as a bottleneck, the opportunity to obtain enough DWIs for a HARDI reconstruction remains in many cases a luxury.
- To observe non-Gaussian diffusion, the strength of the magnetic gradients used in the scan is increased [101]. Increasing the gradient strength increases diffusion rates, which in turn are inversely proportional to relaxation time and DWI intensity. If we increase the gradient strength enough, the DWI intensities can fall below the noise floor, an effect seen with HARDI imaging settings [60].
- Recent research suggests that limitations of the tensor model with regard to crossing fibers can be overcome by taking into account neighborhood information [12,92]. With such advancements, it remains unclear if HARDI can provide information that cannot be recovered from a diffusion tensor image.

Due to these aforementioned reasons and the wealth of diffusion tensor medical research [76], the tensor model cannot be ignored.

19.2.6 Diffusion MR Image Regularization

While modeling local diffusion through tensors or ODFs provides useful information with regard to orientation dependence and anisotropy, we must not lose sight of the fact that the resulting diffusion profiles are estimates of the underlying diffusion and are open to error. While noise reduction can be performed on DWIs as seen in Section 19.1.3, any remaining noise can be amplified due to a poor model fit. Further, the choice of model, particularly the tensor model, can result in a poor fit to the diffusion measurements. These potential errors will affect any further analysis of the diffusion data and so it is important that they be addressed.

Methods for addressing model error generally fall into two categories: correcting for noise and correcting for model choice. We look at these in turn.

19.2.6.1 Noise Reduction

Total variational regularization has been a popular choice for noise removal in the diffusion tensor images [35,37,43,85,99,108]. Generally, total variation approaches involve the minimization of an energy functional:

$$E_{TV}(\mathcal{I}) = E_{data}(I_{orig}, \mathcal{I}) + E_{diff}(\mathcal{I}) \tag{19.15}$$

where E_{data} is a term that measures the distance from the noise reduced image $\mathcal{I}$ from the original version of the image I_{orig}. The term E_{diff} is the Perona and Malik anisotropic diffusion term that controls the smoothing in the image [104]. The variations between the proposed methods relate to what is regularized and how they maintain the constraints of the diffusion tensor. In some cases, only the principal diffusion direction (i.e., the tensor's primary eigenvector) is regularized [108]. In other cases, the tensor's eigenvectors and eigenvalues are regularized separately [37,99]. Still other approaches regularize all tensor elements at once, using the Cholesky factorization [35] or tensor distance metrics [43,85] to ensure positive definiteness of the diffusion tensors.

Other regularization approaches include graph-based anisotropic diffusion filtering of the tensor elements [110], Markov random field minimization of the principal diffusion directions [86], and bilateral [49] and nonlocal means [107] filtering using tensor distance metrics.

19.2.6.2 Model Correction

While noise in the diffusion measurements is one source of inaccuracy, the chosen diffusion model may also be inaccurate. As seen earlier, the diffusion tensor model is unable to accurately represent diffusion profiles in areas of

crossing fibers. Recent work has looked at determining crossing fiber locations from neighborhood information in tensor images. Barmpoutis et al. estimate an ODF called a *tractosema* by performing a neighborhood kernel integration with a specialized kernel [12]. Meanwhile, Savadjiev et al. produce a similar ODF by measuring the probability of diffusion along 3D curves through the surrounding neighborhood [92]. These approaches are able to resolve certain crossing fiber situations, yet it remains unclear whether all crossing fibers can be discovered using some similar post-processing on diffusion tensor images.

Despite attempts to fully address noise and inconsistencies within the data, it is unlikely that we will ever be sure that all imperfections will be removed. As such, recent work has looked at quantifying this error using statistical methods [52,106].

19.3 Brain Connectivity Mapping from dMRI

The orientation information in dMRI is incredibly valuable in mapping out structure in the white matter of the brain. As diffusion is strongest along the fiber tracts that make up neural pathways, the directions of maximal diffusion at each voxel location can be used to help reconstruct the fiber tracts, thereby mapping out connectivity in the brain. The problem of mapping out these connections is known as tractography and is complicated by many factors. We have already mentioned two: poor diffusion model fitting and noisy diffusion measurements. Here, we look at how these complications, and others, affect existing algorithms as well as how the tractography results are used.

19.3.1 Streamline Tractography

The earliest approaches to the tractography problem surrounded tracing out 3D curves that followed the direction of strongest diffusion [15,72]. These 3D curves, known as streamlines, evolve using the following Euler equation:

$$\mathbf{r}(s_1) = \mathbf{r}(s_0) + \alpha\varepsilon_1(\mathbf{r}(s_0)) \tag{19.16}$$

where $\mathbf{r}$ is the streamline curve parametrized by its length from a given seed point and s_i are points along the curve. ε_1 is the PDD at the given location on the curve and the choice of notation comes from the use of the diffusion tensor's primary eigenvector as the PDD. The PDD acts as the tangent to the streamline as it evolves with stepsize α, where α is sufficiently smaller than the voxel size to limit discretization effects on the evolving curve. These aspects are further shown in Figure 19.9. While the aforementioned Euler equation most easily describes the streamline evolution, a higher-order Runge–Kutta method is commonly used to improve numerical stability in

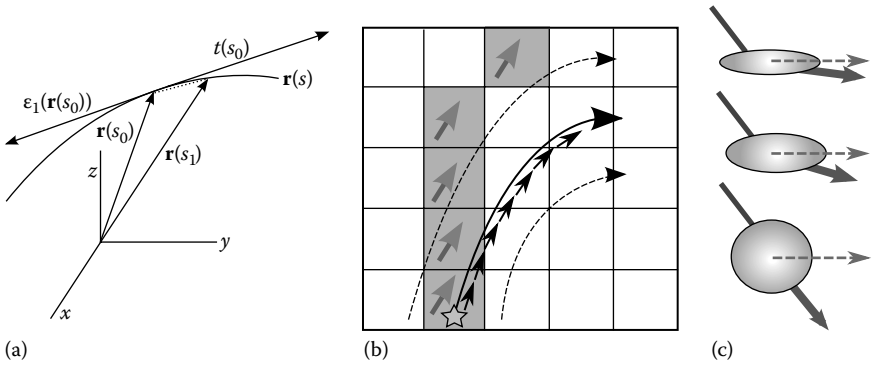

FIGURE 19.9
Examples of streamline evolution. Streamlines evolve in the direction tangent to the local PDD. Stepsizes in the evolution equation are chosen sufficiently small so as to avoid poor tracking due to discretization. (a) Streamline tangent to PDD, (b) effect of stepsize on streamline tractography, and (c) tensor deflection. (Adapted from Basser, P.J. et al., *Magn. Reson. Med.*, 44, 625, 2000; Lazar, M. et al., *Hum. Brain Map.*, 18, 306, 2003; Mori, S. et al., *Ann. Neurol.*, 45(2), 265, 1999, respectively.)

the streamline evolution [15]. The streamline evolution continues until the PDD becomes unreliable. Typically, the reliability of the PDD is captured through either FA [15] or via neighborhood PDD coherence [72]. This initial tractography approach is referred to in the literature as the *fiber assignment by continuous tracking* (FACT) method.

One of the concerns with the FACT approach is that it follows the PDD regardless of whether the PDD at a voxel is an accurate estimation of fiber tract orientation. In areas of lower anisotropy (but still above the termination threshold of FACT), the PDD may become more unreliable. In these situations, we may wish to regulate the effect of the local PDD on the direction of the evolving streamline. This is the idea behind the tensor deflection (TEND) approach [65]. In this algorithm, the local diffusion tensor $\mathbf{D}$ is used to deflect the incoming streamline curve as given by the evolution equation:

$$\mathbf{r}(s_1) = \mathbf{r}(s_0) + \alpha \mathbf{D}(\mathbf{r}(s_0)) \cdot \mathbf{r}(s_0) \tag{19.17}$$

The greater the anisotropy of the tensor, the more reliable the PDD and therefore the stronger the deflection of the streamline fiber. An example of this evolution is shown in Figure 19.9c.

By deflecting the incoming streamline curve, TEND implicitly creates a curvature constraint on the evolving streamline. The streamline can only bend as much as a diffusion tensor will allow. In some situations, this curvature constraint can cause the TEND algorithm to deviate from a high-curvature fiber tract, thereby generating a poor result [38]. To compensate for this effect, the tensorline approach was proposed [65] that evolves the streamline curve based on a weighted combination of (19.16) and (19.17). As a result, the curvature can be turned on and off based on local anisotropy

or prior knowledge. Streamline tractography has also been extended to multitensor models [23] and HARDI using extracted ODF maxima [26]. In these cases, streamlines follow the PDD of minimal angle with the incoming curve. A further review can be found in [112].

From a computational standpoint, these streamline approaches have many limitations, namely:

- Streamline approaches only follow one tract at a time. The algorithm cannot naturally handle situations where tracts branch or cross. One approach to address this concern is to perform a brute force implementation of the algorithm where every point of the brain is, in turn, used as a seed. The tracts that are kept are ones that flow through one or more regions of interest [73]. Even so, this brute force approach doesn't guarantee that crossing or kissing (i.e., barely touching) fibers are appropriately handled.
- These algorithms, particularly the FACT algorithm, assume that the principal diffusion direction is an accurate and error-free estimate of the fiber direction. Any error in the PDD measurement will result in accumulation of error with each step taken. Error in the PDD, while small at each step, can accumulate to the point where the streamline can "jump" into a neighboring tract, thereby giving a false display of anatomical connectivity [73].
- Despite the aforementioned issues, these tractography algorithms present a binary result: a 3D space curve. There is no representation of the confidence or accuracy of the resulting streamline tract. Attempts are being made to quantify that confidence from streamlines [116], but addressing this problem is still in the early stages.

Even so, streamline tractography has been successful in detecting major fiber tracts like the forceps major shown in Figure 19.10a.

19.3.2 Probabilistic Tractography

One of the major concerns with tractography is the amount of confidence we can have in the accuracy of the generated tracts. As such, significant work has gone into performing tractography from a probabilistic point of view. Given points A and B in a diffusion MR image $\mathcal{I}$, we can define the probability of a tract connecting A and B as

$$p(A \to B \mid \mathcal{I}) = \sum_{n=1}^{\infty} \int_{\Omega_{AB}^{n}} p(n) p(\mathbf{v}_{1:n} \mid \mathcal{I}) d\Omega_{AB}^{n} \tag{19.18}$$

where n is the length of the tract, $\mathbf{v}_{1:n}$ is a random path of length n, and Ω_{AB}^{n} is the space of fiber tracts of length n that connect A to B [42]. Given the exponential number of paths in the space Ω_{AB}^{n}, this integration cannot be done

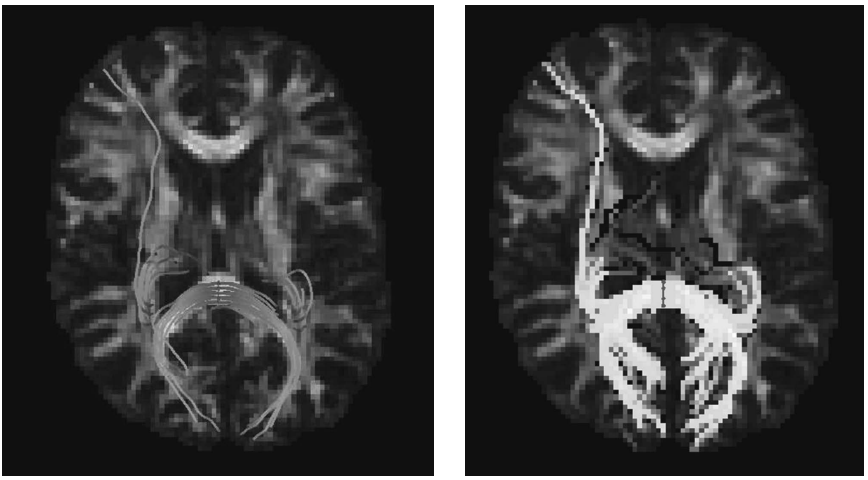

(a) (b)

FIGURE 19.10
Examples of streamline and probabilistic tractography applied to a seed region in the splenium of the corpus callosum. Note that since probabilistic tractography uses streamline tractography as an underlying mechanism, the results are similar. (a) Streamline tractography and (b) probabilistic tractography.

analytically. Instead, the probability $p(A \rightarrow B|\mathcal{I})$ is sampled through the use of Markov chain Monte Carlo (MCMC) [22,25,42,63].

Conceptually, MCMC-based probabilistic tractography shares many similarities with streamline tractography algorithms. Both trace out 3D streamlines by following a local tangent vector. The differences with probabilistic tractography approaches are that instead of exclusively using the principal diffusion direction as the local tangent to the curve, we sample each tangent vector $\mathbf{v}_i$ from a given distribution $p_i(\mathbf{v}_i|\mathbf{v}_{i-1}, \mathcal{I})$. Also, we repeat the streamline tractography many times from the same seed A. Each resulting streamline is considered a sample of $p(A \rightarrow B|\mathcal{I})$. With enough of these samples (K), we can obtain a reasonable approximation of the probability that regions A and B are connected:

$$p(A \rightarrow B \mid \mathcal{I}) = \sum_{n=1}^{\infty} \sum_{k=1}^{K} p(n) \frac{\vartheta(\mathbf{v}_{1:n}^k)}{K} \tag{19.19}$$

The function $\vartheta(\mathbf{v}_{1:n}^k)$ is equal to one if path k connects regions A and B and zero otherwise. The prior probability $p(n)$ is usually taken to be uniform, thereby being unbiased to path length. Effectively, the probability that A and B are connected is equal to the fraction of random paths that connect A and B [25]. Probability maps containing the values from (19.19) can then be displayed and analyzed. An example is shown in Figure 19.10b.

While many probabilistic tractography algorithms have been proposed, their key differences seem to lie in how the distribution $p_i(\mathbf{v}_i|\mathbf{v}_{i-1},\mathcal{I})$ models the local tangent vectors $\mathbf{v}_i$ that make up each random path. In the succeeding text are the most popular approaches:

Diffusion profile: We noted earlier that the diffusion tensor describes a Gaussian model of diffusion. As such, this Gaussian model is commonly used to describe the distribution $p_i(\mathbf{v}_i|\mathbf{D}_i)$ [48,63]. The incoming tract direction $\mathbf{v}_{i-1}$ is modeled as being independent of the local diffusion tensor $\mathbf{D}_i$, thereby splitting the tangent vector distribution into two terms:

$$p_i(\mathbf{v}_i \mid \mathbf{v}_{i-1}, \mathcal{I}) = p_i(\mathbf{v}_i, \mathcal{I}) p(\mathbf{v}_i \mid \mathbf{v}_{i-1}) \tag{19.20}$$

where the diffusion data is the local diffusion tensor $\mathcal{I} = \mathbf{D}_i$. The conditional probability $p(\mathbf{v}_i|\mathbf{v}_{i-1})$ is typically chosen as a binary distribution with $p(\mathbf{v}_i|\mathbf{v}_{i-1}) = 1$ if the angle between $\mathbf{v}_i$ and $\mathbf{v}_{i-1}$ is less than ninety degrees (and zero otherwise).

These Gaussian approaches typically replace the diffusion tensor $\mathbf{D}$ with a scaled version $\mathbf{D}^\alpha$. As expected, the diffusion data contain diffuse information and α is used to decrease probabilities perpendicular to the PDD. Sample values for α range from 2 [48] to 7 [63].

HARDI versions have also followed a similar approach with a sharpened version of the diffusion ODF used for $p_i(\mathbf{v}_i, \mathcal{I})$ [39,62].

Heuristic approaches: Parker et al. proposed the probabilistic index of connectivity (PICo) approach where the tangent vector distribution also takes the form in (19.20). The key difference is how the diffusion data are used. The distribution $p_i(\mathbf{v}_i, \mathcal{I})$ is not Gaussian in this case and instead is replaced by a heuristic distribution based on local anisotropy [81]. The PDD is taken as the mean tangent direction with a cone of uncertainty whose apex angle is a function of a local anisotropy measure.

Bayesian formulations: The tangent vector distribution can also be described using Bayes' rule with respect to the diffusion data:

$$p_i(\mathbf{v}_i \mid \mathbf{v}_{i-1}, \mathcal{I}) = \frac{p(\mathcal{I} \mid \mathbf{v}_i, \mathbf{v}_{i-1}) p(\mathbf{v}_i \mid \mathbf{v}_{i-1})}{p(\mathcal{I})} \tag{19.21}$$

The posterior distribution $p(\mathcal{I} \mid \mathbf{v}_i, \mathbf{v}_{i-1})$ captures how well the diffusion model fits the DWI samples and is typically approximated using a Gaussian distribution on the model's residual fit [21,22,36,42]. Friman et al. update the curvature prior $p(\mathbf{v}_i|\mathbf{v}_{i-1})$ to be the dot product between the two tangent vectors [42].

Statistical bootstrap: Instead of assuming some distribution for the tangent vector, some have used bootstrap techniques to approximate the distribution from multiple samples [59,64,106]. During the image acquisition process, multiple DWIs can be obtained for each gradient direction. When it comes to fitting a diffusion

model, we can do so by fitting to a randomly selected subset of the DWIs. This fitting process can be repeated for many image subsets, thereby generating multiple diffusion MR images. The distribution of the PDDs generated from this set of dMRIs can then be used as a model-free approximation to $p_i(\mathbf{v}_i, \mathcal{I})$.

In the absence of multiple DWI acquisitions, wild bootstrap can be performed [59,106]. In this situation, noise is added to the DWIs by using random perturbations of the residual of the model fit. The fitting is then reperformed for each set of noise-simulated DWIs to obtain multiple diffusion MR images from which the distribution $p_i(\mathbf{v}_i, \mathcal{I})$ can be generated.

Probabilistic tractography approaches have the advantage of characterizing uncertainty in the tractography algorithm. Even so, these methods too have their limitations:

- As each step taken along a tract contains some uncertainty, the connection probabilities we obtain using this approach are inevitably linked to the length of the tract. As such, we cannot interpret these probabilities as a measure of tract quality since they are not invariant to length [58].
- As with streamline tractography, noise can still cause the maxima of the tangent vector distribution to be off. There exists no mechanism in the tractography algorithm to correct for this error.
- The number of path samples required to approximate (19.18) is commonly on the order of thousands [22,42,59]. This results in significant computational cost and running times on the order of an hour or more for a given seed point [111]. Some recent work has tried to address this issue through, for example, the use of particle filters [111].

19.3.3 Front Propagation Tractography

A third set of tractography algorithms can be described as front propagation approaches where some form of information propagates outward from a given seed region at a speed proportional to the amount of fiber tract evidence. The information propagated by the front can then be used to reconstruct the fiber tracts. These algorithms can be divided into three main groups based on their computational aspects.

19.3.3.1 Fast Marching Tractography

Conceptually, the fast marching tractography approaches are distinguished by the calculation of a time of arrival of the propagating front for each voxel [83]. This arrival time T is related to the speed of the front F via the Eikonal equation

$$|\nabla T|F = 1 \quad \text{or} \quad T(\mathbf{r}_i) = T(\mathbf{r}_{i-1}) + \frac{|\mathbf{r}_i - \mathbf{r}_{i-1}|}{F(\mathbf{r}_i)} \tag{19.22}$$

where $\mathbf{r}_i$ and $\mathbf{r}_{i-1}$ are neighboring voxels on opposite sides of the propagating front.

With the arrival times calculated for all voxels, fiber tracts can be delineated by performing gradient descent on the arrival time map. By generating tracts in this fashion, situations of branching and merging fibers are handled naturally through the propagation of the front.

The speed F of the front is set based on the presence or absence of a fiber tract. A common choice is to use the diffusion profile as the speed function [67,68,77,93], thereby ensuring faster speed along directions of faster diffusion. Another choice is neighborhood PDD coherence [83]. By making the front speed an indicator of tract presence, we can characterize the tract's "quality" as some function of the speed. One approach is to characterize the confidence of a tract γ by its weakest link τ [83]:

$$\zeta(\gamma) = \min_{\tau} F(\gamma(\tau)) = \min_{\tau} \frac{1}{|\nabla T(\gamma(\tau))|} \tag{19.23}$$

While measuring tract quality in this fashion is a heuristic approach, it does provide us with a measure of confidence that is invariant to path length.

19.3.3.2 Tractography via Flow Simulation

Instead of using an arrival time map for tract reconstruction, we can interleave the two operations, thereby recovering the tract as we propagate the front. Noting that diffusion is fastest along a fiber tract, some researchers [17,18,32,45,61,77,96,113,117] have proposed that we simply simulate the diffusion and reconstruct candidate tracts through the analysis of the diffusion front.

The diffusion is simulated using Fick's second law [77], given as

$$\frac{\partial u}{\partial t} = \nabla \cdot (\psi \nabla u) \tag{19.24}$$

where

u is the local molecular concentration

ψ is the diffusion function (either the tensor $\boldsymbol{D}$ [18,45,61,77,113] or a diffusion ODF in the case of HARDI [32])

Given a seed point, the diffusion process is simulated for a fixed time t. The resulting concentration map u is then thresholded to obtain the hard diffusion front shown in Figure 19.11a. The voxels along the diffusion front are then scored based on a set of criteria to determine the likelihood that they are on a fiber tract [32,61,113]. Sample criteria include distance from the seed point, FA, and path curvature [61]. The diffusion is then simulated at each candidate point and the process repeats itself.

While this approach has generally gone out of favor due to the ad hoc criteria used to select fiber tract points, the ideas generated by this tractography approach have been applied elsewhere. One example is the work of O'Donnell

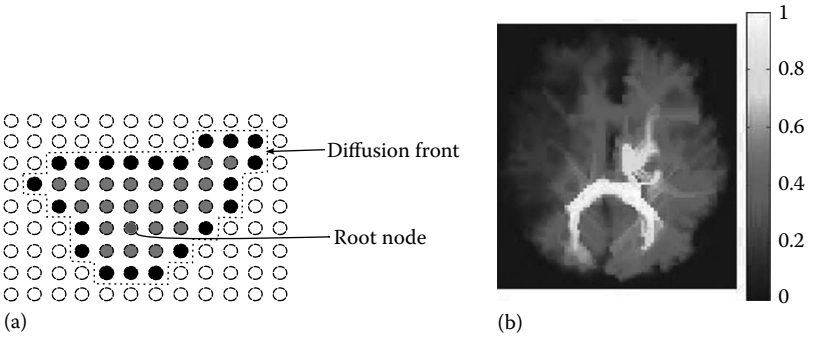

FIGURE 19.11
Examples of front propagation tractography, specifically the representation of the diffusion front and the connection strengths generated using minimal path tractography on a seed in the splenium of the corpus callosum. (a) An example diffusion front and (b) example of minimum path tractography. (Adapted from Kang, N. et al., *IEEE Transactions on Medical Imaging*, 24(9), 1127, September 2005, Booth, B.G. and Hamarneh, G. Exact integration of diffusion orientation distribution functions for graph-based diffusion MRI analysis, *Proceedings of ISBI*, pp. 935–938, 2011; respectively.)

et al. where the steady state flux (i.e., $\partial u/\partial t = 0$) is solved for and tracts are generated that maximize the resulting flux [77]. A similar approach is used by Hageman et al. where instead of modeling diffusion, they model fluid flow using the Navier–Stokes equation [47]. By using the fluid flow model, Hageman et al. are capable of modeling further concepts of the flow (e.g., viscosity).

19.3.3.3 Minimal Path Tractography Algorithms

A third set of tractography algorithms also display this concept of front propagation: graph-based minimal path algorithms [53,94,109]. These tractography algorithms discretize the image space into a graph formulation and use Dijkstra's algorithm to obtain the path of strongest diffusion. In this case, the front being propagated is the boundary between the visited and unvisited nodes.

To ensure the shortest path is the path of strongest diffusion, the edge weights in the graph are set to $w(e_{ij}) = -log(P_{diff}(i,j))$, where the pseudo-probability P_{diff} is given as

$$P_{diff}(i,j) = \frac{1}{Z}\left(\int_{(\theta,\phi)\in\beta_i} \psi_i(\theta,\phi)dS + \int_{(\theta,\phi)\in\beta_j} \psi_j(\theta,\phi)dS \right) \tag{19.25}$$

where

ψ_i is the diffusion ODF at voxel i
Z is a normalizing constant
β_i is the solid angle around the graph edge between i and j [53]

As with the fast marching algorithm, we can consider a "weakest link" connection strength here as well by selecting the largest edge weight along the tract.

One must be concerned when using this method to ensure that the angular discretization provided by the edge connectivity is fine enough to avoid diverging effects similar to those in Figure 19.9. An example of this form of tractography is shown in Figure 19.11b.

19.4 Conclusions

dMRI provides us with the ability to analyze brain connectivity noninvasively. By measuring the diffusion of water molecules along various directions in 3D and knowing that cell structure restricts molecular diffusion, we are able to infer the directional organization and integrity of fibrous tissue. Further modeling of these diffusion measurements allows us to assess characteristics such as bulk diffusivity and anisotropy. The directional dependence of the diffusion can also be used to trace out the imaged axonal fibers.

Various computational aspects of dMRI have been presented in this chapter, from acquiring diffusion-weighted MRI to modeling the diffusion through the use of diffusion tensors or diffusion ODFs to uncovering neural pathways with tractography algorithms. These analysis techniques have become even more established with their incorporation into software packages like FSL, MedINRIA, Camino, and TrackVis. These aspects of image analysis merely scratch the surface of what may be possible with this relatively new imaging technique. Already, there is work being done in the areas of segmentation [118], registration [119], and statistics of dMRI data [84]. It is hoped that continued work in dMRI will culminate in the ability to generate a human connectome: a detailed connectivity map of the human brain [30].

References

1. D. Alexander. A comparison of q-ball and PASMRI on sparse diffusion MRI data. In *Proceedings of International Society of Magnetic Resonance in Medicine (ISMRM)*, Kyoto, Japan, p. 90, 2004.
2. D.C. Alexander. Maximum entropy spherical deconvolution for diffusion MRI. In *Proceedings of Information Processing for Medical Imaging—IPMI*, Glenwood Springs, Co, Vol. 19, pp. 76–87, 2005.
3. D.C. Alexander. Multiple-fiber reconstruction algorithms for diffusion MRI. *Annals of the New York Academy of Sciences*, 1046:113–133, 2005.
4. D.C. Alexander. *Visualization and Processing of Tensor Fields*, Chapter 5: An introduction to computational diffusion MRI: The diffusion tensor and beyond, pp. 83–106. Springer Berlin Heidelberg, Heidelberg, Germany, 2006.

5. D.C. Alexander, G.J. Barker, and S.R. Arridge. Detection and modeling of non-Gaussian apparent diffusion coefficient profiles in human brain data. *Magnetic Resonance in Medicine*, 48:331–340, 2002.
6. A.W. Anderson. Measurement of fiber orientation distributions using high angular resolution diffusion imaging. *Magnetic Resonance in Medicine*, 54:1194–1206, 2005.
7. J.L.R. Andersson and S. Skare. A model-based method for retrospective correction of geometric distortions in diffusion-weighted EPI. *NeuroImage*, 16:177–199, 2002.
8. Y. Assaf and P. Basser. Composite hindered and restricted model of diffusion (CHARMED) MR imaging of the human brain. *NeuroImage*, 27(1):48–58, 2005.
9. Y. Assaf, R.Z. Freidlin, G.K. Rohde, and P.J. Basser. New modeling and experimental framework to characterize hindered and restricted water diffusion in brain white matter. *Magnetic Resonance in Medicine*, 52:965–978, 2004.
10. Y. Bai and D.C. Alexander. Model-based registration to correct for motion between acquisitions in diffusion MR imaging. In *Proceedings of International Symposium on Biomedical Imaging—ISBI*, Paris, France, pp. 947–950, 2008.
11. A. Barmpoutis, M.S. Hwang, D. Howland, J.R. Forder, and B.C. Vemuri. Regularized positive-definite fourth order tensor field estimation from DW-MRI. *NeuroImage*, 45:S153–S162, 2009.
12. A. Barmpoutis, B.C. Vemuri, D. Howland, and J.R. Forder. Extracting tractosemas from a displacement probability field for tractography for DW-MRI. In D. Metaxas, L. Axel, G. Fichtinger, and G. Székely, eds., *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008*, Vol. 5241 of *LNCS*, pp. 9–16. Springer, Heidelberg, Germany, 2008.
13. P.J. Basser and D.K. Jones. Diffusion-tensor MRI: Theory, experimental design and data analysis—A technical overview. *NMR in Biomedicine*, 15:456–467, 2002.
14. P.J. Basser, J. Mattiello, and D. LeBihan. MR diffusion tensor spectroscopy and imaging. *Biophysical Journal*, 66:259–267, 1994.
15. P.J. Basser, S. Pajevic, C. Pierpaoli, J. Duda, and A. Aldroubi. In vivo fiber tractography using DT-MRI data. *Magnetic Resonance in Medicine*, 44:625–632, 2000.
16. S. Basu, T. Fletcher, and R. Whitaker. Rician noise removal in diffusion tensor MRI. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, Vol. 4190 of *LNCS*, pp. 117–125. Springer, Heidelberg, Germany, 2006.
17. P.G. Batchelor, D.L.G. Hill, D. Atkinson, F. Calamanten, and A. Connelly. Fibre-tracking by solving the diffusion-convection equation. In *Proceedings of International Society of Magnetic Resonance in Medicine (ISMRM)*, Honolulu, Vol. 10, 2002.
18. P.G. Batchelor, D.L.G. Hill, F. Calamante, and D. Atkinson. Study of connectivity in the brain using the full diffusion tensor from MRI. In M.F. Insana and R.M. Leahy, eds., *Proceedings of Information Processing in Medical Imaging: IMPI 2001*, Davis, CA, pp. 121–133, 2001.
19. C. Beaulieu. The basis of anisotropic water diffusion in the nervous system—A technical review. *NMR in Biomedicine*, 15:435–455, 2002.
20. C. Beaulieu. The biological basis of diffusion tractography. In *Proceedings of International Symposium on Biomedical Imaging (ISBI)*, Arlington, VA, pp. 347–350, 2006.
21. T.E.J. Behrens, H.J. Berg, S. Jbabdi, M.F.S. Rushworth, and M.W. Woolrich. Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage*, 34:144–155, 2007.

22. T.E.J. Behrens, M.W. Woolrich, M. Jenkinson, H. Johansen-Berg, R.G. Nunes, S. Clare, P.M. Matthews, J.M. Brady, and S.M. Smith. Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance in Medicine*, 50:1077–1088, 2003.
23. O. Bergmann, G. Kindlmann, S. Peled, and C.-F. Westin. Two-tensor fiber tractography. In *Proceedings of International Symposium on Biomedical Imaging (ISBI)*, Washington, DC, pp. 796–799, 2007.
24. D.L. Bihan and E. Breton. Imagerie de diffusion in vivo par résonance magnétique nucléaire. *Compte Rendus de l'Académie de Sciences Paris*, 301:1109–1112, 1985.
25. M. Björnemo, A. Brun, R. Kikinis, and C.-F. Westin. Regularized stochastic white matter tractography using diffusion tensor MRI. In T. Dohi and R. Kikinis, eds., *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, Tokoyo, Japan, Vol. 2488, pp. 435–442, 2002.
26. L. Bloy and R. Verma. On computing the underlying fiber directions from the diffusion orientation distribution function. In D. Metaxas, L. Axel, G. Fichtinger, and G. Székely, eds., *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008*, Vol. 5241 of *LNCS*, pp. 1–8. Springer, Heidelberg, Germany, 2008.
27. N. Bodammer, J. Kaufmann, M. Kanowski, and C. Tempelmann. Eddy current correction in diffusion-weighted imaging using pairs of images acquired with opposite diffusion gradient polarity. *Magnetic Resonance in Medicine*, 51:188–193, 2004.
28. B.G. Booth and G. Hamarneh. Exact integration of diffusion orientation distribution functions for graph-based diffusion MRI analysis. In *Proceedings of ISBI*, Chicago, pp. 935–938, 2011.
29. R. Brown. A brief account of microscopical observations made in the months of June, July and August 1827 on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. *Philosophical Magazine*, 4:161, 1828.
30. E. Bullmore and O. Sporns. Complex brain networks: Graph theoretical analysis of structural and functional systems. *Nature Review Neuroscience*, 10(3):186–198, March 2009.
31. U. Bürgel, K. Amunts, L. Hoemke, H. Mohlberg, J.M. Gilsbach, and K. Zilles. White matter fiber tracts of the human brain: Three-dimensional mapping at microscopic resolution, topography and intersubject variability. *NeuroImage*, 29:1092–1105, 2006.
32. J.S.W. Campbell, K. Siddiqi, V.V. Rymar, A.F. Sadikot, and G.B. Pike. Flow-based fiber tracking with diffusion tensor and q-ball data: Validation and comparison to principal diffusion direction techniques. *NeuroImage*, 27:725–736, 2005.
33. L.-C. Chang, D.K. Jones, and C. Pierpaoli. RESTORE: Robust estimation of tensors by outlier rejection. *Magnetic Resonance in Medicine*, 53:1088–1095, 2005.
34. B. Chen, H. Guo, and A.W. Song. Correction for direction-dependent distortions in diffusion tensor imaging using matched magnetic field maps. *NeuroImage*, 30:121–129, 2006.
35. O. Christiansen, T.-M. Lee, J. Lie, U. Sinha, and T.F. Chan. Total variation regularization of matrix-valued images. *International Journal of Biomedical Imaging*, 2007(27432):11, 2007.
36. P.A. Cook, H. Zhang, S.P. Awate, and J.C. Gee. Atlas-guided probabilistic diffusion-tensor fiber tractography. In *Proceeding of International Symposium on Biomedical Imaging (ISBI)*, Paris, France, pp. 951–954, 2008.

37. O. Coulon, D.C. Alexander, and S. Arridge. Diffusion tensor magnetic image regularization. *Medical Image Analysis*, 8:47–67, 2004.
38. S. Crettenand, S.D. Meredith, M.J. Hoptman, and R.B. Reilly. Quantitative analysis and comparison of diffusion tensor imaging tractography algorithms. In *Proceedings of Irish Signals and Systems Conference—ISSC*, Dublin, Ireland, pp. 105–110, 2006.
39. M. Descoteaux, R. Deriche, T.R. Knösche, and A. Anwander. Deterministic and probabilistic tractography based on complex fibre orientation distributions. *IEEE Transaction Medical Imaging*, 28(2):269–286, 2009.
40. M. Descoteaux, E. Angelino, S. Fitzgibbons, and R. Deriche. Regularized, fast, and robust analytical q-ball imaging. *Magnetic Resonance in Medicine*, 58:497–510, 2007.
41. A. Einstein. Über die von der molekularkinetischen theorie der Wärme geforderte bewegung von in ruhenden flussigkeiten suspendierten teilchen. *Annalen der Physik*, 4:549–590, 1905.
42. O. Friman, G. Farneback, and C.-F. Westin. A Bayesian approach for stochastic white matter tractography. *IEEE Transactions on Medical Imaging*, 25(8):965–978, 2006.
43. C. Frindel, M. Robini, P. Croisille, and Y.-M. Zhu. Comparison of regularization methods for human cardiac diffusion tensor MRI. *Medical Image Analysis*, 13:405–418, 2009.
44. J. Gee and D. Alexander. *Welckert and Hagen: Visualization and Image Processing of Tensor Fields*, Chapter 20—Diffusion tensor image registration. Springer, Berlin, Germany, 2005.
45. D. Gembris, H. Schumacher, and D. Suter. Solving the diffusion equation for fiber tracking in the living human brain. In *Proceedings of International Society of Magnetic Resonance in Medicine (ISMRM)*, Glasgow, UK, p. 1529, 2001.
46. A.C. Guyton. *Textbook of Medical Physiology,* 8th edn. W.B. Saunders, Philadelphia, PA, 1991.
47. N.S. Hageman, A.W. Toga, K. Narr, and D.W. Shattuck. A diffusion tensor imaging tractography algorithm based on navier-stokes fluid mechanics. *IEEE Transactions on Medical Imaging*, 28:348–360, 2009.
48. P. Hagmann, J.-P. Thiran, L. Jonasson, P. Vandergheynst, S. Clarke, P. Maeder, and R. Meuli. DTI mapping of human brain connectivity: Statistical fibre tracking and virtual dissection. *NeuroImage*, 19:545–554, 2003.
49. G. Hamarneh and J. Hradsky. Bilateral filtering of diffusion tensor magnetic resonance images. *IEEE Transactions on Image Processing*, 16(10):2463–2475, October 2007.
50. K.M. Hasan, P.J. Basser, D.L. Parker, and A.L. Alexander. Analytical computation of the eigenvalues and eigenvectors in DT-MRI. *Journal of Magnetic Resonance*, 152:41–47, 2001.
51. J.C. Haselgrove and J.R. Moore. Correction for distortion of echo-planar images used to calculate the apparent diffusion coefficient. *Magnetic Resonance in Medicine*, 36:960–964, 1996.
52. M.O. Irfanoglu, C.G. Koay, S. Pajevic, R. Machiraju, and P.J. Basser. Diffusion tensor field registration in the presence of uncertainty. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, Vol. 5761 of *LNCS*, pp. 181–189. Springer, Heidelberg, Germany, 2009.

53. Y. Iturria-Medina, E.J. Canales-Rodríguez, L. Melie-García, P.A. Valdés-Hernández, E. Martínez-Montes, Y. Alemán-Gómez, and J.M. Sánchez-Bornot. Characterizing brain anatomical connections using diffusion weighted MRI and graph theory. *NeuroImage*, 36:645–660, 2007.
54. K.M. Jansons and D.C. Alexander. Persistant angular structure: New insights from diffusion magnetic resonance imaging data. *Inverse Problems in Physics*, 19:1031–1046, 2003.
55. B. Jian and B.C. Vemuri. A unified computational framework for deconvolution to reconstruct multiple fibers from diffusion weighted MRI. *IEEE Transaction Medical Imaging*, 26(11):1464–1471, 2007.
56. B. Jian and B.C. Vemuri. Multi-fiber reconstruction from diffusion MRI using mixture of Wisharts and sparse deconvolution. In N. Karssemeijer and B. Lelieveldt, eds., *Proceedings of Image Processing for Medical Imaging (IPMI)*, Vol. 4584 of *LNCS*, pp. 384–395. Springer, Heidelberg, Germany, 2007.
57. B. Jian, B.C. Vemuri, E. Özarslan, P.R. Carney, and T.H. Mareci. A novel tensor distribution model for the diffusion-weighted MR signal. *NeuroImage*, 37:164–176, 2007.
58. D.K. Jones. Studying connections in the living human brain with diffusion MRI. *Cortex*, 44(8):936–952, September 2008.
59. D.K. Jones. Tractography gone wild: Probabilistic fiber tracking using the wild bootstrap with diffusion MRI. *IEEE Transactions on Medical Imaging*, 27(9):1268–1274, September 2008.
60. D.K. Jones and P.J. Basser. Squashing peanuts and smashing pumpkins: How noise distorts diffusion-weighted MR data. *Magnetic Resonance in Medicine*, 52:979–993, 2004.
61. N. Kang, J. Zhang, E.S. Carlson, and D. Gembris. White matter fiber tractography via anisotropic diffusion simulation in the human brain. *IEEE Transactions on Medical Imaging*, 24(9):1127–1137, September 2005.
62. I. Kezele, M. Descoteaux, C. Poupon, F. Poupon, and J.-F. Mangin. Spherical wavelet transform for ODF sharpening. *Medical Image Analysis*, 14:332–342, 2010.
63. M.A. Koch, D.G. Norris, and M. Hund-Georgiadis. An investigation of functional and anatomical connectivity using magnetic resonance imaging. *NeuroImage*, 16:241–250, 2002.
64. M. Lazar and A.L. Alexander. Bootstrap white matter tractography (BOOT-TRAC). *NeuroImage*, 24:524–532, 2005.
65. M. Lazar, D.M. Weinstein, J.S. Tsuruda, K.M. Hasan, K. Arfanakis, M.E. Meyerand, B. Badie, H.A. Rowley, V. Haughton, A. Field, and A.L. Alexander. White matter tractography using diffusion tensor deflection. *Human Brain Mapping*, 18:306–321, 2003.
66. A. Leemans and D.K. Jones. The B-matrix must be rotated when correcting for subject motion in DTI data. *Magnetic Resonance in Medicine*, 61:1336–1349, 2009.
67. C. Lenglet, R. Deriche, and O. Faugeras. Diffusion tensor magnetic resonance imaging: Brain connectivity mapping. Technical Report 4983, INRIA, Sophia Antipolis, France, October 2003.
68. C. Lenglet, R. Deriche, and O. Faugeras. Inferring white matter geometry from diffusion tensor MRI: Application to connectivity mapping. In *Proceedings of European Conference on Computer Vision (ECCV)*, Prague, CZE, pp. 127–140, 2004.
69. A.D. Leow, S. Zhu, K. McMahon, G.I. de Zubicaray, M. Meredith, M. Wright, and P.M. Thompson. The tensor distribution model. In *Proceedings of the International Symposium on Biomedical Imaging (ISBI)*, Paris, France, pp. 863–866, May 2008.

70. J.-F. Mangin, C. Poupon, C. Clark, D. Le Bihan, and I. Bloch. Distortion correction and robust tensor estimation for MR diffusion imaging. *Medical Image Analysis*, 6:191–198, 2002.
71. S. Mori. John Hopkins Medical Institute: Laboratory of Brain Anatomical MRI, in vivo human database. http://lbam.med.jhmi.edu/ (accessed February 2009).
72. S. Mori, Barbara J. Crain, V.P. Chacko, and Peter C.M. van Zijl. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Annals of Neurology*, 45(2):265–269, 1999.
73. S. Mori and P.C.M. van Zijl. Fiber tracking: Principles and strategies—A technical review. *Nuclear Magnetic Resonance in Biomedicine*, 15:468–480, 2002.
74. P. Mukherjee, J.I. Berman, S.W. Chung, C.P. Hess, and R.G. Henry. Diffusion tensor MR imaging and fiber tractography: Theoretic underpinnings. *American Journal of Neuroradiology*, 29:632–641, April 2008.
75. P. Mukherjee, S.W. Chung, J.I. Berman, C.P. Hess, and R.G. Henry. Diffusion tensor MR imaging and fiber tractography: Technical considerations. *American Journal of Neuroradiology*, 29:843–852, May 2008.
76. P.G.P. Nucifora, R. Verma, S.-K. Lee, and E.R. Melhem. Diffusion-tensor MR imaging and tractography: Exploring brain microstructure and connectivity. *Radiology*, 245(2):367–384, November 2007.
77. L. O'Donnell, S. Haker, and C.-F. Westin. New approaches to estimation of white matter connectivity in diffusion tensor MRI: Elliptic PDEs and geodesics in a tensor-warped space. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2002*, pp. 459–466, Tokyo, Japan, 2002.
78. E. Ozarslan, B.C. Vemuri, and T.H. Mareci. Generalized scalar measures for diffusion MRI using trace, variance, and entropy. *Magnetic Resonance in Medicine*, 53:866–876, 2005.
79. E. Özerslan, T. Shepherd, B. Vemuri, S. Blackband, and T. Mareci. Resolution of complex tissue microarchitecture using the diffusion orientation transform (DOT). *NeuroImage*, 31(3):1086–1103, 2006.
80. S. Pajevic and C. Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: Application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine*, 42:526–540, 1999.
81. G.J.M. Parker, H.A. Haroon, and C.A.M. Wheeler-Kingshott. A framework for streamline-based probabilistic index of connectivity (PICo) using a structural interpretation of MRI diffusion measurements. *Journal of Magnetic Resonance Imaging*, 18:242–254, 2003.
82. G.J.M. Parker, J.A. Schnabel, M.R. Symms, D.J. Werring, and G.J. Barker. Nonlinear smoothing for reduction of systematic and random errors in diffusion tensor imaging. *Journal of Magnetic Resonance Imaging*, 11:702–710, 2000.
83. G.J.M. Parker, C.A.M. Wheeler-Kingshott, and G.J. Barker. Estimating distributed anatomical connectivity using fast marching methods and diffusion tensor imaging. *IEEE Transactions on Medical Imaging*, 21(5):505–512, 2002.
84. O. Pasternak, N. Sochen, and P.J. Basser. The effect of metric selection on the analysis of diffusion tensor MRI data. *NeuroImage*, 49:2190–2204, 2010.
85. X. Pennec, P. Fillard, and N. Ayache. A Riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006.
86. C. Poupon, C.A. Clark, V. Frouin, J. Régis, I. Bloch, D. Le Bihan, and J.-F. Mangin. Regularization of diffusion-based direction maps for the tracking of brain white matter fascicles. *NeuroImage*, 12:184–195, 2000.

87. A. Ramirez-Manzanares, P.A. Cook, and J.C. Gee. A comparison of methods for recovering intra-voxel white matter fiber architecture from clinical diffusion imaging scans. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, Vol. 5241 of *LNCS*, pp. 305–312. Springer, Heidelberg, Germany, 2008.
88. Y. Rathi, O. Michailovich, M.E. Shenton, and S. Bouix. Directional functions for orientation distribution estimation. *Medical Image Analysis*, 13:432–444, 2009.
89. C.A.-L. Rodrigo de Luis-Garcia and C.-F. Westin. *Tensors in Image Processing and Computer Vision*, Chapter Segmentation of Tensor Fields: Recent Advances and Perspectives. Springer, London, U.K., 2009.
90. G.K. Rohde, A.S. Barnett, P.J. Basser, S. Marenco, and C. Pierpaoli. Comprehensive approach for correction of motion and distortion in diffusion-weighted MRI. *Magnetic Resonance in Medicine*, 51:103–114, 2004.
91. K.E. Sakaie and M.J. Lowe. Quantitative assessment of motion correction for high angular resolution diffusion imaging. *Magnetic Resonance Imaging*, 28:290–296, 2010.
92. P. Savadjiev, J.S.W. Campbell, G.B. Pike, and K. Siddiqi. 3D curve inference for diffusion MRI regularization and fibre tractography. *Medical Image Analysis*, 10:799–813, 2006.
93. G. Sebastiani, F. de Pasquale, and P. Barone. Quantifying human brain connectivity from diffusion tensor MRI. *Journal of Mathematical Imaging and Vision*, 25:227–244, 2006.
94. S.N. Sotiropoulos, L. Bai, P.S. Morgan, C.S. Constantinescu, and C.R. Tench. Brain tractography using Q-ball imaging and graph theory: Improved connectivities through fibre crossings via a model-based approach. *NeuroImage*, 49:2444–2456, 2010.
95. E.O. Stejskal and J.E. Tanner. Spin diffusion measurements: Spin echoes in the presence of a time-dependent field gradient. *The Journal of Chemical Physics*, 42(1):288–292, January 1965.
96. J.-D. Tournier, F. Calamante, D.G. Gadian, and A. Connelly. Diffusion-weighted magnetic resonance imaging fibre tracking using a front evolution algorithm. *NeuroImage*, 20:276–288, 2003.
97. A. Tristan-Vega and S. Aja-Fernandez. DWI filtering using joint information for DTI and HARDI. *Medical Image Analysis*, 14:205–218, 2010.
98. T.P. Trouard, Y. Sabharwal, M.I. Altbach, and A.F. Gmitro. Analysis and comparison of motion correction techniques in diffusion-weighted imaging. *Journal of Magnetic Resonance Imaging*, 6(6):925–935, 1996.
99. D. Tschumperle and R. Deriche. Diffusion tensor regularization with constraints preservation. In *Proceedings of Computer Vision and Pattern Recognition—CVPR*, Kawai, HI, Vol. 1, pp. 948–953, 2001.
100. D.S. Tuch. Q-ball imaging. *Magnetic Resonance in Medicine*, 52:1358–1372, 2004.
101. D.S. Tuch, T.G. Reese, M.R. Wiegell, N. Makris, J.W. Belliveau, and V.J. Wedeen. High angular resolution diffusion imaging reveals intravoxel white matter fiber heterogeneity. *Magnetic Resonance in Medicine*, 48:577–582, 2002.
102. Z. Wang, B.C. Vemuri, Y. Chen, and T.H. Mareci. A constrained variational principle for direct estimation and smoothing of the diffusion tensor field from complex DWI. *IEEE Transactions on Medical Imaging*, 23(8):930–939, 2004.
103. V. Wedeen, T. Reese, D. Tuch, M. Wiegel, J.-G. Dou, R. Weiskoff, and D. Chessler. Mapping fiber orientation spectra in cerebral white matter with fourier-transform diffusion MRI. In *Proceedings of the International Society of Magnetic Resonance in Medicine*, Denver, CO, p. 82, 2000.

104. J. Weickert and T. Brox. Diffusion and regularization of vector- and matrix-valued images. Technical Report 58, Universität des Saarlandes, Saarbrücken, 2002.
105. C.-F. Westin, S.E. Maier, H. Mamata, A. Nabavi, F.A. Jolesz, and R. Kikinis. Processing and visualization for diffusion tensor MRI. *Medical Image Analysis*, 6:93–108, 2002.
106. B. Whitcher, D.S. Tuch, J.J. Wisco, A.G. Sorensen, and L. Wang. Using the wild bootstrap to quantify uncertainty in diffusion tensor imaging. *Human Brain Mapping*, 29:346–362, 2008.
107. N. Wiest-Daessle, S. Prima, P. Coupe, S.P. Morrissey, and C. Barillot. Non-local means variants for denoising of diffusion-weighted and diffusion tensor MRI. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI*, Vol. 4792 of *LNCS*, pp. 344–351. Springer, Heidelberg, Germany, 2007.
108. E. Yoruk and B. Acar. Structure preserving regularization of DT-MRI vector fields by nonlinear anisotropic diffusion filtering. In *Proceedings of European Signal Processing Conference (EUSIPCO)*, Antalya, Turkey, p. 4, 2005.
109. A. Zalesky. DT-MRI fiber tracking: A shortest paths approach. *IEEE Transaction Medical Imaging*, 27(10):1458–1571, 2008.
110. Fan Zhang and E.R. Hancock. Tensor MRI regularization via graph diffusion. In *Proceedings of British Machine Vision Conference (BMVC)*, Edinburgh, UK, pp. 578–589, 2006.
111. F. Zhang, E.R. Handcock, C. Goodlett, and G. Gerig. Probabilistic white matter fiber tracking using particle filtering and von Mises-Fisher sampling. *Medical Image Analysis*, 13:5–18, 2009.
112. J. Zhang, H. Ji, N. Kang, and N. Cao. Fiber tractography in diffusion tensor magnetic resonance imaging: A survey and beyond. Technical Report 437–05, University of Kentucky, Lexington, KY, April 2005.
113. J. Zhang, N. Kang, and S.E. Rose. Approximating anatomical brain connectivity with diffusion tensor MRI using kernel-based diffusion simulations. In *Proceedings of Information Processing for Medical Imaging—IPMI*, Vol. 3565 of *LNCS*, pp. 64–75. Springer, Heidelberg, Germany, 2005.
114. K. Nand, R. Abugharbieh, B. Booth, and G. Hamarneh. Detecting structure in diffusion tensor MR images. In Lecture Notes in Computer Science, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Toronto, CA, 6892:90–97, 2011.
115. G. Hamarneh, C. McIntosh, and M. Drew. Perception-based visualization of manifold-valued medical images using distance-preserving dimensionality reduction. *IEEE Transactions on Medical Imaging (IEEE TMI)*, 30(7):1314–1327, 2011.
116. C. Brown, B.G. Booth, and G. Hamarneh. K-Confidence: Assessing uncertainty in tractography using k optimal paths. In *IEEE International Symposium on Biomedical Imaging (IEEE ISBI)*, San Francisco, 250–253, 2013.
117. B.G. Booth and G. Hamarneh. Global multi-region competitive tractography. In *IEEE workshop on Mathematical Methods for Biomedical Image Analysis (IEEE MMBIA)*, Breckenridge, CO, 73–78, 2012.
118. B.G. Booth and G. Hamarneh. A cross-sectional piecewise constant model for segmenting highly curved fiber tracts in diffusion MR images. In Lecture Notes in Computer Science, *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Nagoya, Japan, to appear, 2013.
119. B.G. Booth and G. Hamarneh. Consistent information content estimation for diffusion tensor MR images. *In IEEE Conference on Healthcare Informatics, Imaging and Systems Biology (IEEE HISB)*, San Jose, 166–173, 2011.