28 Basic Diffusion Tensor Imaging

28.1 Introduction

[1]

Let us pause here for a moment and review what has been learned so far. In Chapter 14 we saw how an NMR signal can be generated and detected from a sample within a coil. We then elaborated on this idea in two different, and to this point, unconnected, ways. The first was to demonstrate (Chapter 16) that localization of the signal from different spatial regions of the sample could be achieved by the application of magnetic field gradients and reconstruction by Fourier transformation. This is the process of magnetic resonance imaging (MRI). We extended these ideas to ultrafast (single shot) imaging in Chapter 18, such as echo planar imaging (EPI), that allowed us to acquire images rapidly enough to obviate bulk subject motion, at the expense of greater sensitivity to a number of physical effects that create artifacts, that we nonetheless could mitigate. But we also reconsidered the simple case of a sample in a coil in the presence of a bipolar gradient pulse and discovered (Chapter 25) that diffusing spins produced a signal loss due to their motion through the gradients. Moreover, with an appropriate model for the diffusion process, the diffusion tensor could be estimated, from which we could derived scalar measures related to the mean diffusivity and the diffusion anisotropy. The bipolar pulse had the additional feature that it refocused stationary spins at its completion. In the present chapter we are going to tie this all together for the following purpose: To combine the diffusion weighting and subsequent analysis we examined for a small sample with the spatial localization (MRI) on a full human brain in order to produce images with diffusion weighting from which can be performed localized (e.g. in every voxel) estimation of the diffusion tensor. This entire process is called *diffusion tensor imaging*, or *DTI*.

As mentioned in the Introduction and in Chapter 25, the order I've chosen to present the aforementioned aspects of this process is somewhat non-standard, and at first glance perhaps even disconnected. But that is precisely the point. The payoff comes here in this chapter where we see that the process of DTI is really just the combination of several experimental procedures and physical processes that stand on their own. These distinctions are important to grasp because one of the biggest hurdles I've encountered in teaching DTI is the confusion over what aspects are related to diffusion, which are related to imaging, and which are related to reconstruction and estimation. This motivated the presentation in terms of the separated components to be combined only after each is understood in its own right. The central fact that allows this simplification is that bipolar gradients affect diffusing spins but leave stationary spins unchanged. As we shall see, this allows us to really just combine the results of these aforementioned chapters into one

[1] Make it clear here that you're just looking at the single fiber voxel model, etc, etc.

procedure. And the pleasing result is that this central chapter is relatively short: you have all the background you need to understand how DTI works.

Now, having said that, we still must proceed with caution, and state from the outset that our goal in this chapter is to describe the most basic form of DTI. By "basic" we mean specifically:[2]

- The diffusion is assumed to be Gaussian
- There is no subject movement
- Diffusion gradients cause no distortions
- Diffusion through imaging gradients has no effect

In later chapter we will relax these restrictions and consider important extensions to DTI that are necessary to more accurately characterize complex neural tissues.

28.2 Diffusion in a bipolar gradient: Review

In the previous chapter we investigated the effect of diffusion in a bipolar gradient, the results of which we review here. We found that a bipolar gradient rephased all of the stationary spins, causing a gradient echo. But this required that the spins be stationary because the rephasing required that the spins were subject to an equal and opposite gradient induced field. Since this field is spatially dependent, by definition, any movement of a spin would cause it to have a mismatch between the magnitude of its starting and ending phases, and thus we would expect that the total signal, no longer being from spins completely in phase, would be reduced. This is precisely the effect produced by diffusing spins: since their final location is different from their initial location, their phases are different, and thus all the spins no longer come back into phase and the final signal is reduced. How much the signal is reduced depends on how much they go out of phase, which depends on the particulars of their motion. Since it is impractical (ney, impossible) to follow the trajectory of individual spins, we must resort to probabilistic arguments.

So, in order to determine the signal in the presence of moving spins, we need to know where they start and where they end up, and then figure out the phase they've accrued by changing locations both in the presence of the gradients (i.e., during the time δ) as well as what happens in the time Δ between the centers of the gradients (i.e., the area/2). This problem sounds complicated. But what is remarkable is that, at least for the most basic case, the answer to these questions is surprisingly simple. We'll give a sketch of the answer here, and then go into the details in the rest of the chapter about how we arrived at this, what assumption were needed, and then touch on where those assumptions fail. The details of that will have to wait for a later chapter.

Let us return to our simple model of Gaussian diffusion. Recall the meaning of the Gaussian distribution that characterizes diffusion: it is the probability of finding a spin at a particular location after a certain time. This is just what we require in order to determine what location changes, and hence what phase changes, are accrued during the bipolar pulse. The answer, as we shall see, is that the signal decays in a very simple way as a function of the applied gradients, whose amplitudes and parameters are contained within a parameter b, called the *b-factor*. For an anisotropic tissue characterized by Gaussian diffusion with a diffusion tensor $\boldsymbol{D}$, we found that the signal in the presense of a bipolar gradient characterized by the b-factor is (Eqn **??**)

$$\mathfrak{s}(q,\tau) = e^{-b\tilde{D}} \tag{28.1}$$

[2] Have we listed all our implicit assumptions of this chapter?

where $\mathfrak{s} \equiv s(b)/s_o$ is the ratio of the diffusion weighted signal to the unweighted (i.e., $b = 0$) signal $s_o \equiv s(b = 0)$. The b-factor is given by

$$b = q^2\tau \tag{28.2}$$

where $q \equiv |\boldsymbol{q}|$ is just the area of the gradient that is pointing along the $\hat{\boldsymbol{q}}$ direction (Eqn **??**) and

$$\tilde{D} \equiv \hat{\boldsymbol{q}}^t \boldsymbol{D} \hat{\boldsymbol{q}} \tag{28.3}$$

is the projection of the diffusion tensor $\boldsymbol{D}$ along the direction $\hat{\boldsymbol{q}}$.

So the sensitivity of the bipolar gradient to diffusion depends upon the area squared: doubling the gradient amplitude *quadruples* the diffusion sensitivity. Note also that if the gradient is turned off ($G = 0$), then $b = 0$ and $s(b) = s_o$. So turning off the gradient pulse gives us the constant s_o. The decay is also proportional to the time τ, often called the *diffusion time*, which is given by

$$\tau = \Delta - \delta/3 \tag{28.4}$$

and so is proportional to the time Δ, minus a term proportional to the gradient width. Thus putting the gradients farther apart increases the sensitivity. This makes sense, since the longer spins are allowed to diffuse, the greater the spread in their locations, thus the greater the variations in the fields they subject to, and thus the wider the spread of phases, and thus the more signal loss when they're all added up.

And, because the diffusion tensor represents variations in the diffusion along different spatial directions, this decaying exponential characteristic of the signal depends upon which direction we measure the diffusion along. That is, which direction the diffusion weighting is along. We found, using the vectorial nature of gradients, that it was a simple matter to point the diffusion weighting along any desired direction by suitable combination of bipolar gradients along the three principal axes of the scanner. Performing measurements along different directions and using our model for the diffusion process, it was then possible to estimate the diffusion tensor.

Let us now turn to extending this to the imaging process. The key new feature is the incorporation of the diffusion weighting gradient into an imaging sequence.

28.3 Creating diffusion contrast in images: Diffusion Weighting Gradients

We now come to an important juncture in the book. We saw in Chapter 16 how to create MR images, and in Chapter 25 how the bipolar gradient causes a signal intensity that is proportional to diffusion. How can we then proceed to the real work of this book, which is to combine the two to get diffusion weighted images? Given the complicated sequence of RF and gradient pulses necessary to form a complete MR image, this might seem like a daunting task. But, surprisingly, this is not so! And all because of the nice qualities of the bipolar gradient pulse (recall that we use this term as a shorthand to refer specifically to two gradients of *equal* areas and opposite sign). Recall the two essential properties of the bipolar gradient pulse: 1) Stationary spins are refocussed at the end of the second pulse; 2) Diffusing spins suffer a signal loss proportional to the amplitude and timing of the two pulses. From these two facts alone we can develop the basic method of DTI. Because a bipolar gradient refocussing stationary spins, it can be inserted into a pulse sequence, as long as it doesn't interfere (i.e., overlap) with any of the imaging gradients, and be *invisible* to the imaging portion of the pulse sequence! And so, from fact (2), we can

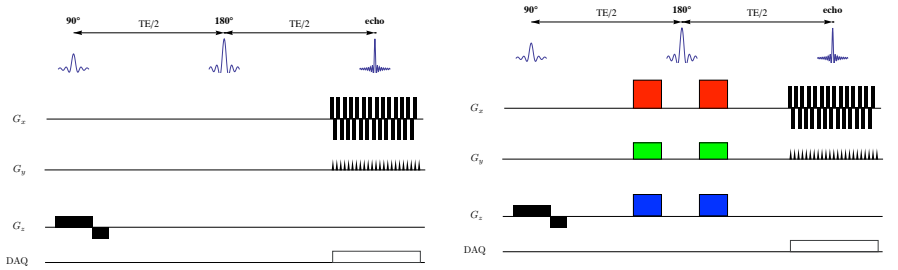

(a) Spin echo EPI acquisition.

(b) Diffusion weighted spin echo EPI acquisition.

Figure 28.1 Spin echo EPI diffusion weighted pulse sequence. Bipolar diffusion weighting gradient pulses (gray lobes) are inserted into a basic spin echo EPI pulse sequence without interfering with the imaging process because a bipolar pulse refocusses stationary spins and is thus *invisible* to the pulse sequence as long as it does not overlap with the imaging gradients.

add bipolar *diffusion weighting gradients* whose amplitudes and timings we can manipulate to probe the diffusion characteristics of our tissue, but create MR images simultaneously. The result is MR images that have diffusion weighting everywhere, ie, in each voxel, thus allowing us to investigate the spatial variations of whatever diffusion related quantities we can measure within a voxel. Such images are called *diffusion weighted images*.

The most straightforward, and most common, method of combining diffusion weighting with and imaging sequences is to insert a bipolar pair of diffusion weighting gradients into a standard spin echo pulse sequence, as shown in Figure 28.1.

Note the important fact that the second lobe of the diffusion gradients in Figure 28.1 is the same sign as the first lobe because of the presence of the 180° refocussing pulse This has been implicit in our previous description of "effective gradients" (Section **??**), in which we replace two gradient on either side of the 180° pulse with a bipolar pair. This allowed us to discuss diffusion effects without having to mention and RF pulses. Throughout what follows we will ignore the effects of the refocussing pulse(s). In practice, we have to be careful about where these diffusion weighting gradients are placed in the pulse sequence, and what system imperfections they are subject to. These will be discussed in detail in later chapters.

Notice the important fact that the diffusion weighting gradients are applied independently on each axis in Figure 28.1. That is, the second lobe of the bipolar gradient on the x-axis refocusses spins brought out of phase by the first lobe on the x-axis, regardless of what's happening on the y and z axes. So the bipolar gradients on each of the axes refocus stationary spins on their respsective axes, and produce diffusion weighting along these axes as well, independently of what happens on the other axes. But now recall the important and very practical fact that *gradients add like vectors*. Therefore, we see that *the combined effect of simultaneously applied diffusion weighting gradients is the rotation of the direction of diffusion sensitivity in the direction defined by the resulting combined gradient vector.* This was shown in Figure 29.1. Since we, the system operator, controls these gradients, we can point them in any directions we want. For example, in Figure 28.5 is shown diffusion encoding along the three principal scanner axes and an arbitrary direction.

So now we combine these ideas - bipolar diffusion weighting gradient put into a spin echo

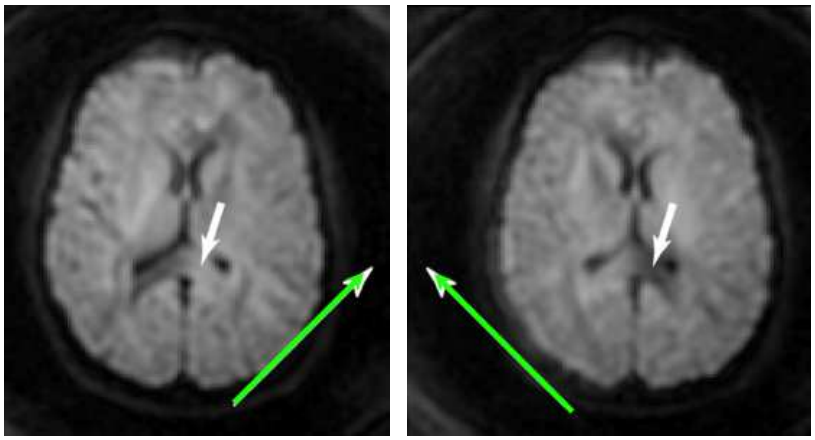

(a) Diffusion weighting parallel to the right splenium of the corpus callosum.

(b) Diffusion weighting parallel to the left splenium of the corpus callosum.

Figure 28.2 (**Use newer images!**) Two images with diffusion weighting in different directions (shown by the green arrows) approximately parallel and perpendicular to the left and right splenia of the corpus callosum. This example was specifically chosen because the two structure are close to perpendicular, so the two chosen gradient directions applied along these structures are close to perpendicular to one another. In both cases, there is signal loss in the portion of the corpus callosum parallel to the applied gradient, but not that perpendicular to it (white arrows).

sequence with our ability to diffusion encode in different directions, and see what happens. In Figure 28.2 is shown two images with diffusion sensitivity along two different directions chosen to be along the right and left splenia of the corpus callosum. This example was specifically chosen because the two structure are close to perpendicular, so the two chosen gradient directions applied along these structures are close to perpendicular to one another. In both cases, there is signal loss in the portion of the corpus callosum parallel to the applied gradient, but not that perpendicular to it (white arrows).

28.4 Interlude: What is the effect of the imaging gradients?

If diffusion through gradients results in a directionally dependent signal loss, don't the imaging gradients have some effect? In fact, they do, but in most clinical applications this effect is negligible. That is not the case is high field imaging using high powered gradients, however. To understand the effect of imaging gradients, it is useful to look at the diffusion weighting process in a different way.

In the simple case of Gaussian diffusion, the spread of spins in the spatial domain is equivalent to a convolution in the image domain by a Gaussian kernel. This is shown schematically in Figure 28.3 for the 1D case. As we saw in Section 11.5, convolution of the two functions $I(x)$ and $H(x)$ is one domain (x) is equivalent to multiplication of their Fourier transforms $\tilde{I}(k)$ and $\tilde{H}(k)$ in the conjugate domain:

$$I(x) \star H(x) = \tilde{I}(k)\tilde{H}(k) \tag{28.5}$$

(a) A group of spins at time $t = 0$ spreads into a Gaussian distribution at time $t = \tau$ with a standard deviation $\sigma = \sqrt{2D\tau}$.

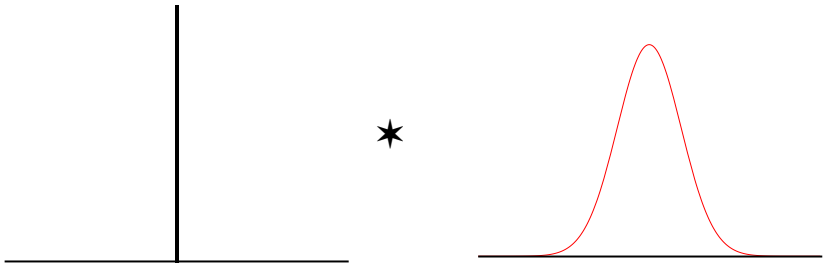

(b) The process in (a) is equivalent to convolution of the initial distribution with a Gaussian with standard deviation $\sigma = \sqrt{2D\tau}$.

Figure 28.3 Gaussian diffusion is equivalent to convolution in the image domain.

where "" denotes the Fourier transform. In this case x is the spatial (image) domain and k is the fourier (data acquisition) domain. Let $I(x)$ denotes the image and $H(x)$ the convolving kernel due to diffusion. Because the width of Gaussian convolution kernel $H(x)$ is very small compared to the voxel dimensions ($\tau = 100\,\text{ms} \rightarrow \sigma \approx 15\,\mu\text{m}$), its Fourier transform $\tilde{H}(k)$ is very broad with respect to the k-space data $\tilde{I}(k)$. This is illustrated in Figure 28.4 Because the width of $H(x)$ is very small, its Fourier tranform $\tilde{H}(x)$ (the red curve in Figure **??**) is very broad. The result is that small gradients (e.g., imaging gradients on a typical 3T clinical system) produce little diffusion effect and much larger diffusion weighting gradients that shift the center of k-space toward the edge of $\tilde{H}(x)$ are needed to produced a measureable effect. This is *not* necessarily true on high field high performance systems where the imaging gradients can be quite large.[3]

28.5 The spatial variations of the diffusion attenuation

Now that we have seen that the imaging process can be integrated with the diffusion weighting process (which consists of bipolar diffusion weighting gradients on each axis), the question, finally, is "What is the signal attenuation in each voxel?". Well, for our simple model of diffusion as unrestricted (free) Gaussian, we already know the answer to this, and with a slight addition

[3] Give examples of gradient values at 3T and on, say, and 11.7T system.

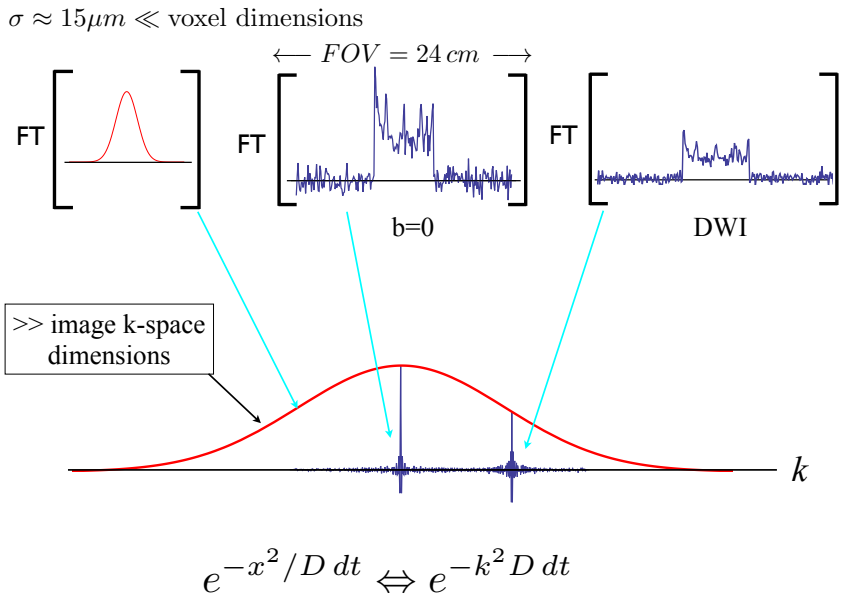

Figure 28.4 Convolution of the image by the diffusion kernel $H(x)$ is equivalent to multiplication of the data by the Fourier transform of the diffusion kernel. Because the width of $H(x)$ is very small, its Fourier tranform $\tilde{H}(x)$ (red) is very broad. The result is that small (e.g., imaging) gradients produce little diffusion effect and much larger diffusion weighting gradients that shift the center of k-space toward the edge of $\tilde{H}(x)$ are needed to produced a measureable effect.

to our notation we can make it obvious. The signal attenuation in 3-dimensions from a bipolar diffusion weighting gradient is given by Eqn 27.3. This is the signal model for a voxel so is easily extended to express the spatial variation of the diffusion in the entire image, the distribution of intensities in three dimensional space $\boldsymbol{x} = \{x, y, z\}$ is

$$s(\boldsymbol{x}, b) = s_0(\boldsymbol{x})\, e^{-b\tilde{D}(\boldsymbol{x})} + \eta(\boldsymbol{x}, b) \tag{28.6}$$

where $\tilde{D} \equiv \hat{\boldsymbol{u}}^T \cdot \boldsymbol{D} \cdot \hat{\boldsymbol{u}}$ is the projection of the diffusion tensor along the applied diffusion weighting gradient direction $\hat{\boldsymbol{u}}$, $b = q^2\tau = g^2\delta^2\,(\Delta - \delta/3)$ is the b-factor, and where $\eta(\boldsymbol{x}, b)$ is the noise and $s_0(\boldsymbol{x})$ is image acquired without diffusion weighting but all other timing parameters the same as in the diffusion weighted images. This is important because the relaxation contrast, which depends on the timing parameters, must be the same in $s(\boldsymbol{x}, 0)$ as in the diffusion weighted images in order that they are normalized correctly so that signal loss due to T_2 relaxation is not confounded with that from diffusion weighting (see the discussion surrounding m_0 in Eqn **??**). The image $s(\boldsymbol{x}, 0)$ is often referred to as the "b equals zero" image. Now, diffusion weighted sequences typically have long echo times TE in order to accommodate the diffusion weighting gradients. An image $s(\boldsymbol{x}, 0)$ acquired with the same timing parameters but without diffusion weighting gradients will thus tend to be $T_2{-}weighted$ (see Chapter 17). The image $s_0(\boldsymbol{x})$ is thus also often referred to as the "T_2" image. The noise $\eta(\boldsymbol{x}, b)$ is generally spatially varying, although we will simplify it by assuming that it is not: $\eta(b) = \eta(\boldsymbol{x}, b)$.

The application of diffusion weighting gradients therefore produces a spatial pattern of signal attenuation throughout the image concomitant with the spatial distribution of diffusion tensors $\boldsymbol{D}(\boldsymbol{x})$. In this chapter we will treat each of the voxels as independent, so that we reconstruct the diffusion tensor in each voxel without any reference to its neighbors. Thus in this section we will figure out what we need to do to reconstruct $\boldsymbol{D}$ and then just repeat the procedure for each voxel. This is again the procedure in standard DTI. Later, in Chapter 40 we will be concerned with the spatial relationship of the diffusion in the voxels and then we'll have to consider the fact that the local diffusion tensor is really a component of a more spatially extended diffusion field.

The noise $\eta(\boldsymbol{x}, b)$ is assumed to be Gaussian with zero mean and variance σ_η^2: that is $\eta(\boldsymbol{x}, b) \sim N(0, \sigma_\eta^2)$, because, as we found in Chapter 19, this is a good model for the background noise in MR images. This, of course, is predicated on using the complex MR images, rather than the magnitude images, which have noise that is characterized by a Rayleigh, (**?**), rather than a Gaussian, distribution (Chapter 19) [4]. It is common practice, however, to use the magnitude images, which means that the assumption of Gaussian noise is *incorrect*. This will become important for low SNR experiments, such as those using high b-values, where the noise becomes comparable to the signal.

Another implicit assumption in the modeling of $\eta(\boldsymbol{x}, b)$ as Gaussian is that all other noise sources that are not thermal have been eliminated. In practice, to perform an experiment precisely enough to make this statement even approximately true is a non-trivial affair. Suffice it to say for now that the largest sources of error in DTI or *not* thermal noise, but eddy currents and field distortions. We will return to this issue in Chapter 30 where we demonstrate the practical necessity of reducing these artifacts and show some methods by which this can be accomplished to a reasonable degree.

It is important that the diffusion weighting gradients in Figure 28.1 along each imaging axis do not overlap the imaging gradients or they would interfere with the imaging process[5]. But the diffusion gradients on each axis *can* be applied simultaneously, and in fact this is of critical importance. For, just as in imaging, gradient fields add as vectors so that combinations of simultaneously applied diffusion weighting gradients along the magnet's $\{x, y, z\}$ coordinates can be used to make the net diffusion weighting gradient occur in any direction. An example of $s_0(\boldsymbol{x})$ and some diffusion weighted images along different directions is shown Figure 28.5. Notice that the gradient directions are defined in terms of the unit vectors $\hat{\boldsymbol{u}}$. For the same timing parameters and gradient strength, the $b - value$ remains the same, although the direction of the diffusion sensitivity changes. The family of the endpoints of all arrow that are of constant length but arbitrary orientation form the surface of a sphere. Therefore, the surface of possible diffusion gradient vector endpoints for a constant b-value but arbitrary direction is a sphere. This is called the *diffusion sampling sphere*, and depicted in Figure 28.5(k). From this week see also that the length of the vector, and thus the sphere's radius, is the b-value. Thus sampling spheres for different $b - values$ are thus concentric spheres. Here again is another system for which is naturally described by a spherical coordinate system (Section 2.3).

[4] need discussion of Rayleigh (or is it Riccian?) noise in MRI!

[5] For gradient/imaging cross term interactions, see Mattielo refs in the Basser, Jones NMR in Biomedicine review.

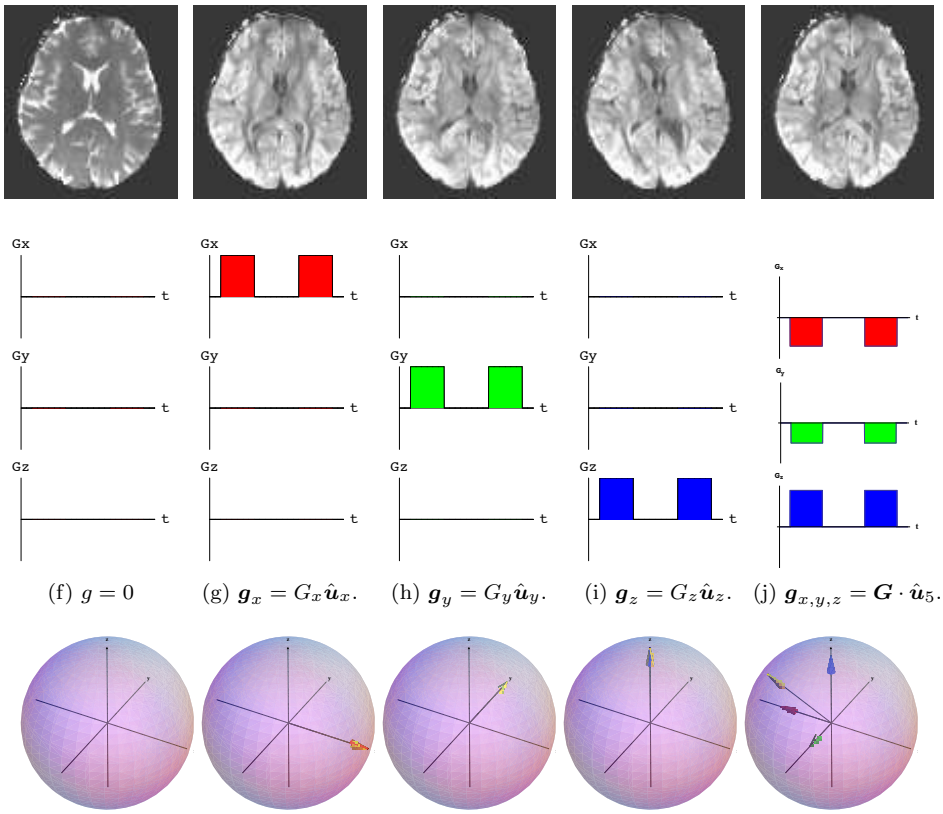

(f) $g = 0$ (g) $\boldsymbol{g}_x = G_x\hat{\boldsymbol{u}}_x$. (h) $\boldsymbol{g}_y = G_y\hat{\boldsymbol{u}}_y$. (i) $\boldsymbol{g}_z = G_z\hat{\boldsymbol{u}}_z$. (j) $\boldsymbol{g}_{x,y,z} = \boldsymbol{G}\cdot\hat{\boldsymbol{u}}_5$.

(k) Gradient directions and the sampling sphere.

Figure 28.5 Directional diffusion encoding. (Top row) Diffusion weighted images EPI images at 3T with $b = 2000s/mm^2$; (Second row) Diffusion gradients used in (a-e); (Bottom) Corresponding diffusion gradient directions in (f-j). (**Collect correct images!**). G_x, G_y, and G_z are the amplitudes of the x, y, and z gradients and $\boldsymbol{G} = \{G_x, G_y, G_z\}$.

28.6 Displaying the estimated tensor

There are several display methods that are particularly useful in visualizing the local structure of diffusion. As we saw in Section 27.2, estimation of the diffusion tensor is tantamount to reconstructing the 3D Gaussian pdf, from which can be constructed the contours of the probability distribution of particle positions: the diffusion ellipsoids. This can be calculated in each voxel, since the diffusion tensor is estimated in each voxel, and displayed as an image, as in Figure 28.6. This method of presentation is useful not only in allowing us to quickly assess the diffusion characteristics of individual voxels, but also to give a more global picture of the patterns of diffusion *amongst* voxels. We'll follow this line of thought in the next section.

(**This next figure was in adv-hard.tex for some reason. I moved it here but need text.**)

An example in human data is shown in Figure 28.7.

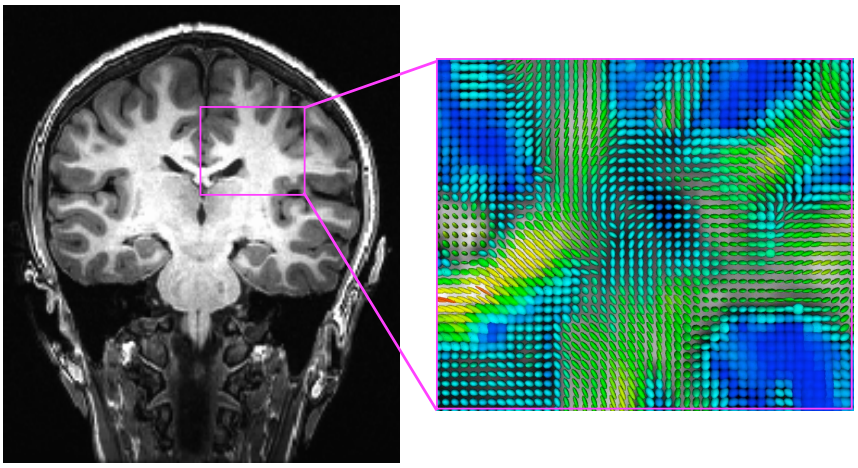

Figure 28.6 The reconstructed diffusion ellipsoids in a selected region of the brain.

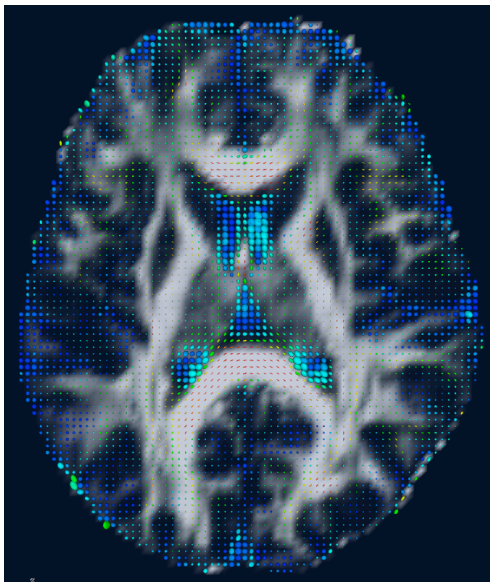

Figure 28.7 The estimated diffusion ellipsoids in a normal human brain. Note that the largest are in regions of CSF, which has the highest mean diffusivity. The anisotropy, on the other hand (shown in underlay in the form of the FA) accentuates the white matter, where it is largest.

28.7 Parameter maps

We have seen that the basic DTI method combines the results of Chapter 25 with the imaging methods of Chapter 18 to acquire data that is diffusion weighted along multiple directions. In Chapter 29 we saw how to estimate the diffusion tensor from this data. The beauty of MRI process is that the result of combining these operations is that we can now perform the *same* analysis we did in Chapter 29 for a single sample *in each voxel*. That is, we can essentially treat

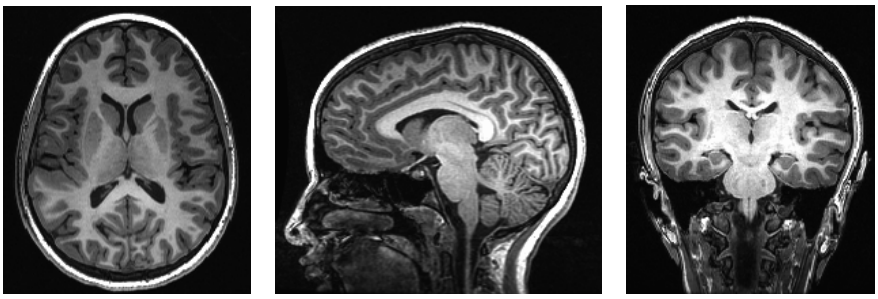

(a) Axial (b) Sagittal (c) Coronal

Figure 28.8 Anatomical images from a normal human subject shown at three orthogonal orientations. inversion recovert $T1$-weighted 3D fast spoiled gradient recalled echo pulse sequence with parameters: flip angle $\alpha = 12^\circ$, echo time $TE = 3\,ms$, repetition time $TR = 8\,ms$, matrix size = $(RL, AP, IS) = (172 \times 256 \times 256)$, field of view $FOV = (170 \times 240 \times 240)\,mm$ for a resolution of $(1 \times .938 \times .938)mm$. (Data courtesy of Dr Susan Tapert, UCSD.)

each voxel as an individual sample, and follow that analysis procedure outlined in Chapter 29. Specifically, once the gradient directions are specified, and the b-matrix constructed, the data in each voxel are fit to a model. For the Gaussian model, we simply plug in the data in each voxel to a routine such as AFNI's 3dDWItoDT which does a non-linear fit for the diffusion tensor under the assumption of additive, Gaussian noise, returning the eigenvalues and eigenvectors. The eigenvalues will tell us about the magnitude of the diffusion along the principal directions of the diffusion ellipsoid, as defined by the eigenvectors. Mathematically,

$$\underset{\text{measured}}{\boldsymbol{D}} = \underset{\text{eigenvectors}}{\boldsymbol{R}^t} \underset{\text{eigenvalues}}{\boldsymbol{D}_\Lambda} \underset{\text{eigenvectors}}{\boldsymbol{R}} \tag{28.7}$$

Having determined the eigensystem (i.e., the eigenvalues and eigenvectors) in each voxel from the measured diffusion tensor, we can then do several important things. First, we can create the maps of the diffusion magnitude parameters, such as the mean diffusivity and the fractional anisotropy, from the eigenvalues. We can also look at the direction of maximum diffusivity from the largest eigenvector, which should tell us something about the underlying structure. And from both the eigenvalues and the eigenvectors we can reconstruct the diffusion ellipsoid or, equivalently, and estimate of the measured apparent angular diffusion coefficient (i.e., the "shape" of diffusion we discussed in Section 27.4). We'll consider each of these in order in the next sections.

We can now apply to each voxel the procedures for estimation of the diffusion tensor discussion in the last chapter. From the eigenvalues in each voxel we can calculate the mean diffusivity $\overline{D}$ and the fractional anisotropy and make images of those parameters. These are shown for a set of data collected on a normal human brain at 3T whose anatomical images are shown in Eqn 28.8. The mean diffusivity and the fractional anisotropy in each voxel of a diffusion tensor image is shown in Figure 28.9. A very useful way to display this information is to overlay the parameters in color over the high resolution anatomical images, displayed in a grayscale. The mean diffusivity in each voxel of a diffusion tensor image (Section 29.6) is shown overlayed on the anatomical imaging in Figure 28.10. An example of the FA calculated from the diffusion tensor in each voxel (Section 29.7) of a diffusion tensor image is shown overlayed on the anatomical imaging in Figure 28.11. Notice that the area of high mean diffusivity are in the regions of cerebro-spinal

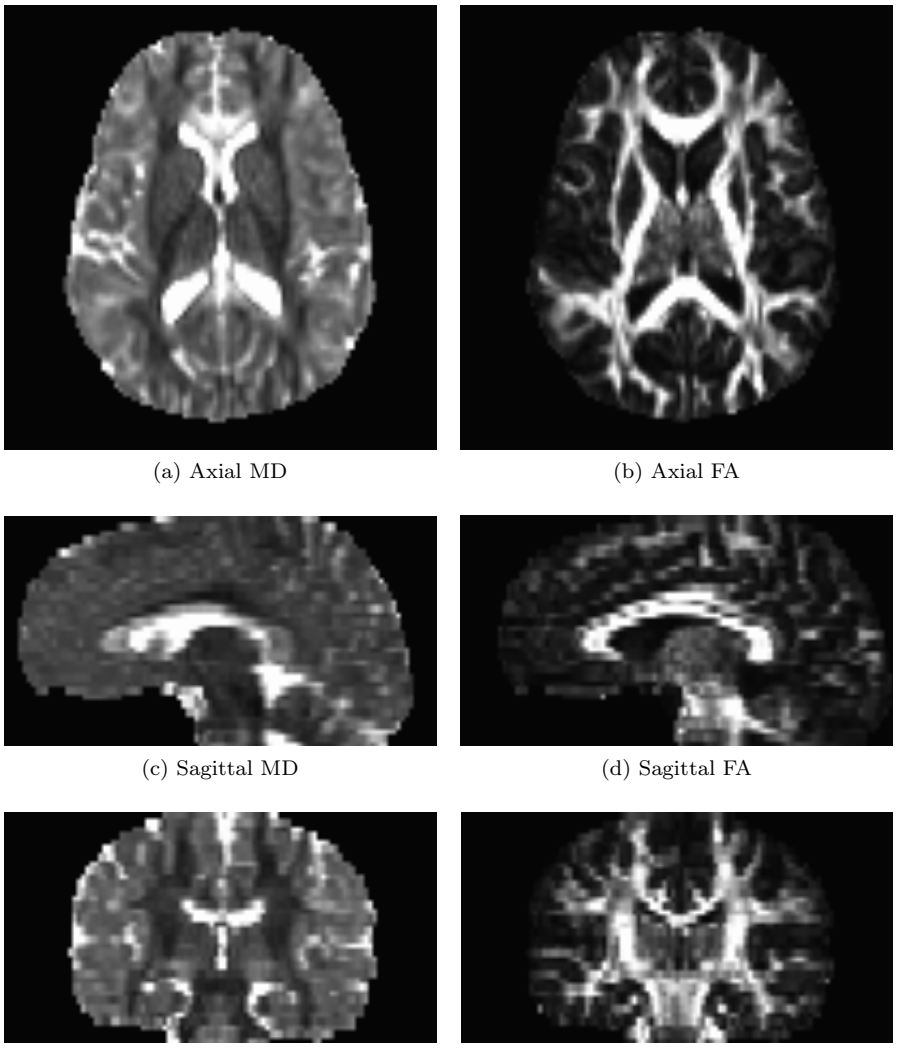

(a) Axial MD (b) Axial FA

(c) Sagittal MD (d) Sagittal FA

(e) Coronal MD (f) Coronal FA

Figure 28.9 Mean diffusivity (left) and fractional anisotropy (right) images from a normal human subject shown at three orthogonal orientations. Data were acquired using a diffusion weighted spin echo EPI pulse sequence with the following parameters: $TE/TR = 93/10,900ms$, $FOV = 240mm$, matrix = 128128, 34 contiguous slices, slice thickness = $3mm$, b-value = $1500s/mm^2 2$, one average. (Data courtesy of Dr Susan Tapert, UCSD).

fluid, or *csf*, which is a freely diffusing liquid that fills the ventricals and the sulci of the cortex[6]. The fractional anisotropy images, in contrast, show larger values in regions containing white matter, since the diffusion is higher along the fiber axis than in the direction perpendicular to that axis. While mean diffusivity can be an important clinical indicator, in stroke for example (e.g., (**?**, **?**, **?**, **?**, **?**)), in many white matter diseases it is the fractional anisotropy that is of greatest interest, since it preferentially distinguished white matter from gray matter. Also, in

[6] better have a neuroanatomist check this sentence!

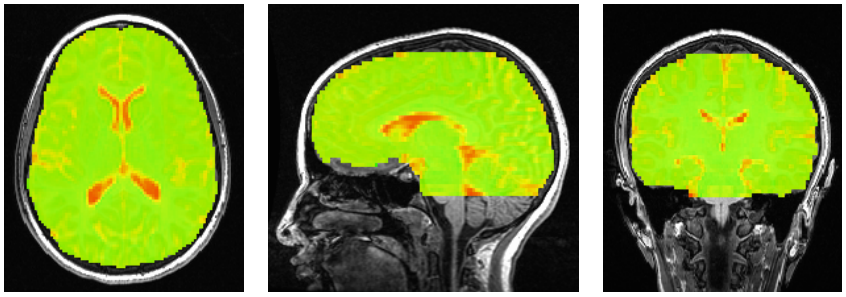

(a) Axial (b) Sagittal (c) Coronal

Figure 28.10 Mean diffusivity images from a normal human subject shown at three orthogonal orientations (top) and overlayed in color on the anatomical images in Figure 28.8 (bottom). (Data courtesy of Dr Susan Tapert, UCSD).

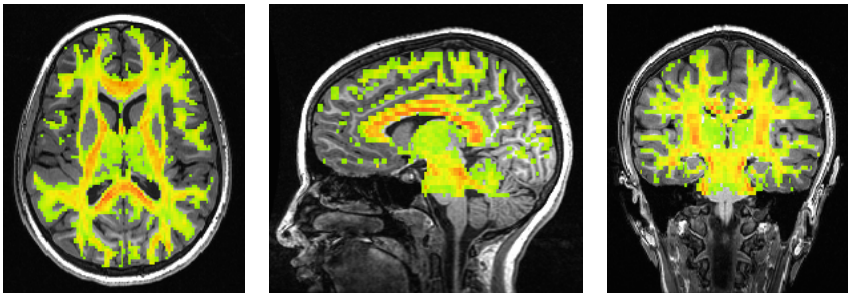

(a) Axial (b) Sagittal (c) Coronal

Figure 28.11 Fractional anisotropy images from a normal human subject shown at three orthogonal orientations (top) and overlayed in color on the anatomical images in Figure 28.8 (bottom). The FA threshold is .3. (Data courtesy of Dr Susan Tapert, UCSD).

many neuroscience application it is the neural connections that are of greatest interest, and thus the fractional anisotropy, as a measure of fiber orientation, is the parameter of interest. With this in mind, we show an additional example of the fractional anisotropy in several axial slices in a normal human brain is shown in Figure 28.12 with a greater number of diffusion directions and a higher signal-to-noise.

28.8 The direction of diffusion: The principal eigenvector

In the previous section we found that displaying the diffusion ellipsoid gave us a nice visual representation of the patterns of diffusion. One can take this notion one step further by considering what our simple Gaussian model of diffusion is suggesting. The longest eigenvector is associated with the highest diffusion, the principal eigenvector is assumed to be in the direction of the fiber. Therefore, if the 3D Gaussian model is correct, *the principal eigenvector can be used as a proxy for the fiber direction.* To effectively visualize the vectors, the following color scheme is typically employed: The principal axes of the scanner ($\{x, y, z\}$) are assigned, respectively, the colors red,

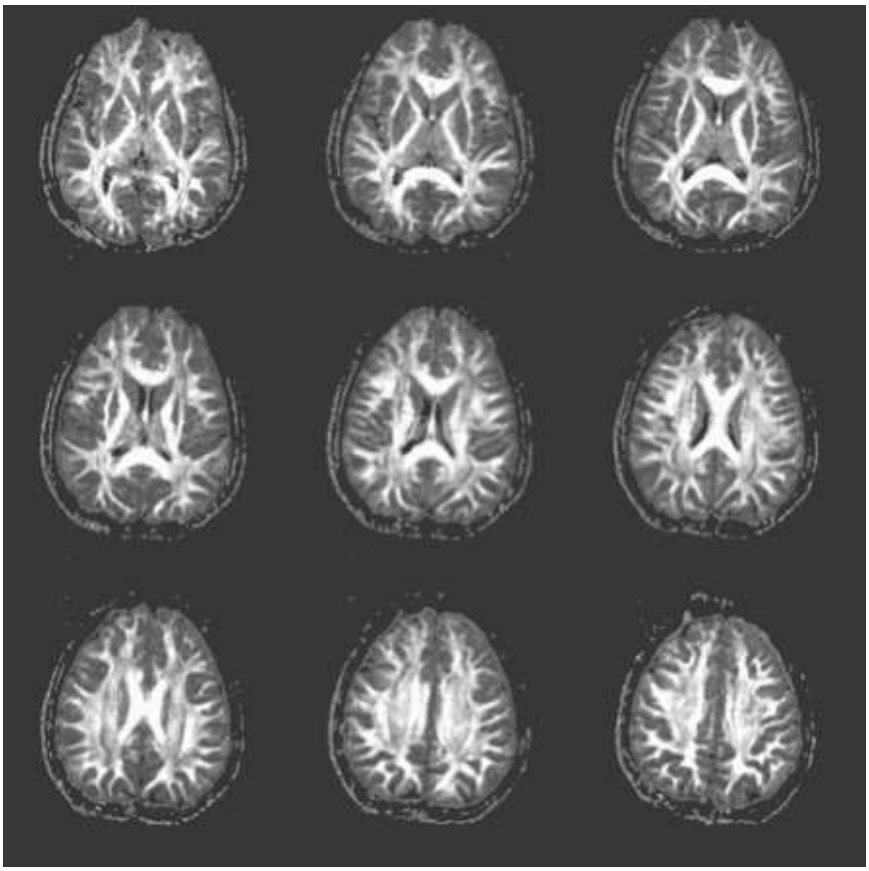

Figure 28.12 Fractional anisotropy in several axial slices in a normal human brain. (**Put imaging parameters here! This is 61 directions!**)

green, and blue. An arbitrary vector is assigned a color that is a mixture of these three colors, with the amount of each color determined by the projection of the vector along the principal axes. For example, a vector at a 45° angle in the $x - y$ is projected equally along the x and y axes, and thus is assigned half red and half green. The color schemes for each principal plane are shown in Figures 28.13a- 28.13c. An example is shown in Figure 28.13. One difficulty with representing the direction with arrows as in Figure 28.13 is that it if often quite hard to see these arrows in a full image. For this reason, an popular alternative scheme is to do away with the arrows and instead represent the direction of the principal eigenvector in a voxel by the directional color. Voxels with principal eigenvectors pointing along these directions are given these colors. This visualization technique allows one to quickly and easily assess the estimated fiber directions over an entire image, as is evident from the high resolution rat brain DTI images Figure 28.14. Combining the FA along with the principal eigenvector is also a useful way to combine parameter maps, as shown in Figure 28.15.

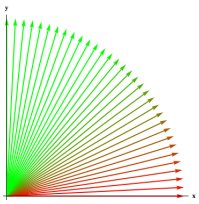

(a) $x - y$ plane

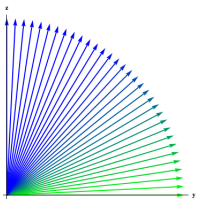

(b) $y - z$ plane

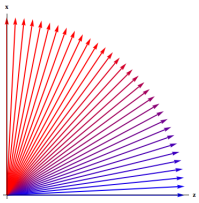

(c) $z - x$ plane

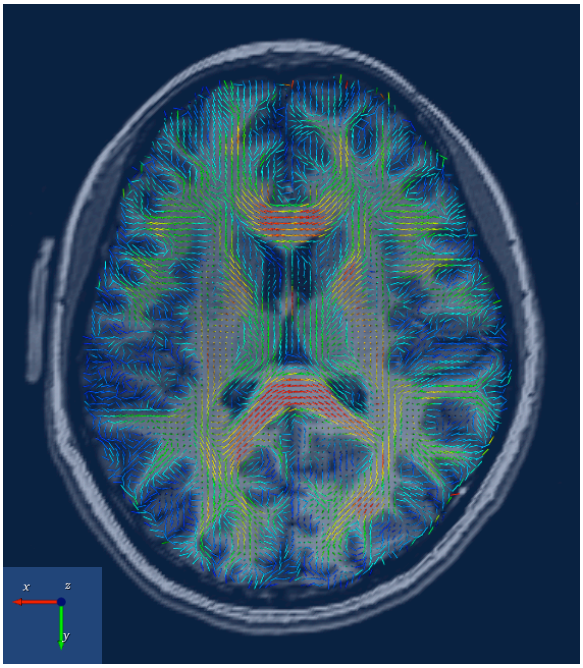

(d) The principal eigenvectors overlayed on anatomical scan.

Figure 28.13 Color encoding scheme for principal eigenvectors.

Suggested Reading

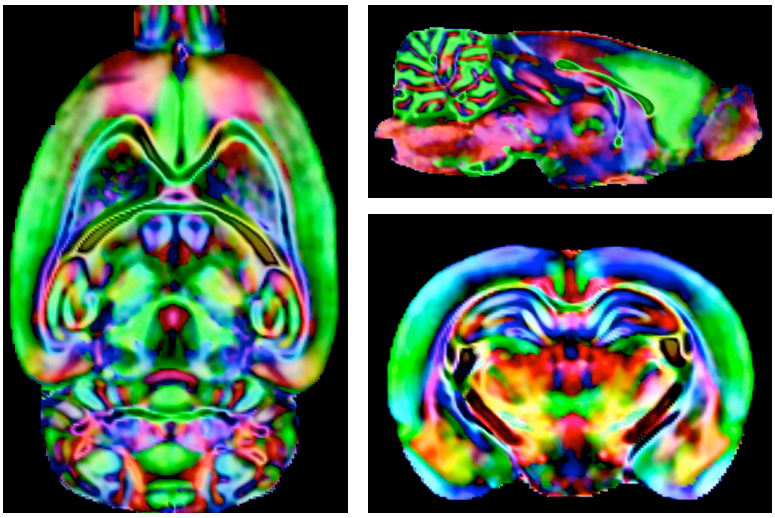

Figure 28.14 Color encoding the of the principal eigenvectors in the principal slices in a high field (11.7T) DTI data of a rat brain (data courtesy of Dr. J.M. Tyszka, CalTech).

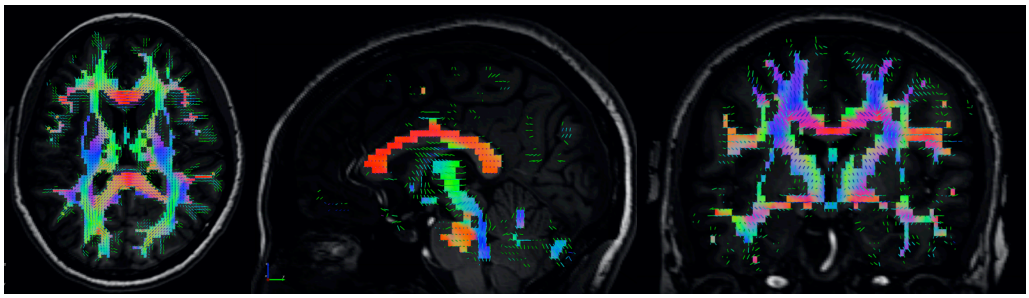

(a) Axial (b) Sagittal (c) Coronal

Figure 28.15 Fractional anisotropy and the principal eigenvectors images from a normal human subject shown at three orthogonal orientations (top) and overlayed in color on the anatomical images in Figure 28.8 (bottom). The FA threshold is .3. (Data courtesy of Dr Susan Tapert, UCSD.)