29 Estimation of the diffusion tensor I: The Basic Approach

[1]

29.1 Introduction

In the previous chapter we found that spins diffusing through a bipolar gradient result in a diminution of the signal from that which would have been produced by an equal number of stationary spins because the diffusing spins acquire a distribution of phases concommitant with their distribution of spatial locations, and thus the integration of the signal over all spins leads to incoherent averaging of the signal. Stationary spins, on the other hand, are refocussed by the second lobe in the bipolar pair, forming a gradient echo. The signal loss of the diffusing spin depends upon their phase distribution which, in turn, depends on the spatial distribution of spins along the gradient direction, and thus our model for the signal depends upon our model for the diffusion. For the simple model of Gaussian diffusion[2], the resulting signal took a very simple form of an exponential decay as a function of both the imaging parameters, encapsulated in the b-value, and the diffusion tensor $\boldsymbol{D}$.

In this chapter we turn to the question that is faced at the end of all such experiments: Given the data and our knowledge about the data, how do we estimate the physical quantities of interest? The first question is, of course, "What are the physical quantities of interest?" In the case of Gaussian diffusion, the answer to this question appears to be easy, since there is a single quantity, the diffusion tensor $\boldsymbol{D}$, that characterizes the diffusion. And, indeed, give this model, that's about all we can do. But, as we shall see in later chapters, for more complex models this question is not so easy to answer. Even in the present case the simplicity is somewhat illusory. For example, if we are looking at diffusion in white matter in a patient with multiple scleroris, what is the relationship between the diffusion tensor and, say, the breakdown of the myelin sheath? This central question about the relationship of the estimated parameters of our *model* to the *true* underlying tissue structure is worth keeping in mind as a reality check on the information content of the images that we acquire. Given the complexity of a block of neural tissue within a voxel, it is clear that we must strike a balance between a model that is sufficiently descriptive to provide useful information on some aspect of tissue structure, while not so complex that we cannot efficiently acquire the data necessary to estimate the parameters.

The focus of this chapter is just that of the previous chapter: a single, diffusion weighted voxel

[1] This chapter has several sections commented out. They may have useful stuff, but they were, for the most part, confusing

[2] Shouldn't there be a law of large number argument for why gaussian diffusion is a good model for lots of diffusing spins?

(Figure 25.1). And I will present the essential first step in DTI analyses - the estimation of the diffusion tensor $\boldsymbol{D}$, within the much simpler discussion of diffusion sensitivity in a single voxel presented in the previous chapter. For that is all we need in order to discuss the estimation of $\boldsymbol{D}$. Once we have estimated $\boldsymbol{D}$, we will see that we can derive some interesting quantities that will tell us about the tissue. Later, when we are concerned with more global aspects such as fiber connections, it will be important to have considered the entire imaging process. But for now we will consider the signal from a single voxel without any imaging. The analysis will be *the same* when we expand to include image acquisition. But it is important to understand the centeral feature that makes this separation possible (at least in the ideal case): The bipolar diffusion weighting gradient *are invisible* to stationary spins. Therefore the diffusion weighting gradients and the imaging procedure can be decoupled in our discussion. In reality, things are not so simple, as the imaging gradients produce diffusion weighting that must be accounted for. We'll leave this complication for a later chapter.

29.2 What's the problem?

So here is the problem that we are faced with. We have a set of measurements of the diffusion weighted voxel and from these measurements we want to estimate the diffusion tensor $\boldsymbol{D}$. As we have seen, the different measurements can be either at different diffusion sensitivities (i.e., different b-values) or along different directions, or both. Typically we will be interested in the variations of the diffusion with direction, since this will tell us something about the underlying geometry (i.e., tissue architecture) of the environment in which spins are diffusion. But we have also seen that estimating the diffusion coefficient (the 1D problem) from multiple b-value measurements involves the non-trivial problem of estimating a parameter from an exponentially decaying signal. Surely the more complex diffusion tensor must also involve this aspect as well, and be even more non-trivial. In other words, shouldn't the estimation of $\boldsymbol{D}$ involved measurements along different directions *and* with multiple b-values? The answer is *yes* but it should be noted that historically, and even in the majority of present practices, this is *not* what is typically done. Rather, a single b-value is chosen and $\boldsymbol{D}$ is estimated from multiple angular measurements taken at this single b-values. There is another obvious but very important thing to keep in mind here: The measurements required to estimate $\boldsymbol{D}$ will *depend on the structure of* $\boldsymbol{D}$, or rather our model for its structure. The simple model of unrestricted Gaussian diffusion where $\boldsymbol{D}$ is a real, symmetric 3×3 matrix requires less measurements than a more complicated $\boldsymbol{D}$ that arises in voxels that have a mixture of different tissues (which we address in Chapter 33). But this is just the point that we made in Chapter 12: Both the measurements and estimates depend on the model.

29.3 Constructing the b-matrix

Our data model, and thus our estimation of $\boldsymbol{D}$, depends upon the tensor product $\boldsymbol{bD}$. Before considering specifically how to estimate $\boldsymbol{D}$, we're going to take a slight detour and first consider the structure of $\boldsymbol{bD}$. In the simple problem of estimating the diffusion coefficient we made multiple measurements at different b-values in order to map out the exponential decay.

Now let's consider the problem of estimating the diffusion tensor $\boldsymbol{D}$ from a set of angular measurements. Our data model is

$$s(q,\tau) = s_o e^{-\boldsymbol{bD}} + \eta(q) \tag{29.1}$$

with

$$\boldsymbol{bD} = \sum_i \sum_j b_{ij} D_{ij} \tag{29.2}$$

where the factor that multiplies D_{ij}, the i,j'th element of the diffusion tensor, is

$$b_{ij} = \boldsymbol{q}_i \boldsymbol{q}_j \tau \tag{29.3}$$

where $q = g\delta$ and $\tau = \Delta - \delta/3$. The i'th gradient points in the direction of the unit vector $\hat{\boldsymbol{u}}_i$, so denoting $g_i \equiv |\boldsymbol{g}_i|$, the q-vectors can be written $\boldsymbol{q}_i = q_i \hat{\boldsymbol{u}}_i$ where $q_i = g_i \delta$ and thus Eqn 29.2 can be written

$$\boldsymbol{bD} = b_{ij} \tilde{D}_{ij} \tag{29.4}$$

where

$$b_{ij} = q_i q_j \tau \tag{29.5a}$$
$$\tilde{D}_{ij} = \hat{\boldsymbol{u}}_i D_{ij} \hat{\boldsymbol{u}}_j \tag{29.5b}$$

The set of n measurement at different angles is of the form Eqn 29.1, where b_{ij} is given by Eqn 29.5, in full matrix form:

$$\boldsymbol{b} = \begin{pmatrix} q_x^2 & q_x q_y & q_x q_z \\ q_y q_x & q_y^2 & q_y q_z \\ q_z q_z & q_z q_y & q_z^2 \end{pmatrix} \tau \tag{29.6}$$

and the diffusion tensor $\boldsymbol{D}$, in its most general form, is

$$\boldsymbol{D} = \begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix} \tag{29.7}$$

Since this matrix is symmetric, i.e., $\boldsymbol{D} = \boldsymbol{D}^t$, we've colored the equivalent off-diagonal components the same color. (In addition, the components are real, i.e., $D_{ij} \in \Re$, although we won't use that fact just quite yet.) Let's write out explicitly what the components of $\boldsymbol{bD}$ in Eqn 29.2 for a *single* gradient direction:

$$\begin{aligned} \frac{1}{\tau} \sum_i \sum_j b_{ij} D_{ij} &= q_x^2 D_{xx} + q_x q_y D_{xy} + q_x q_z D_{xz} \\ &+ q_y q_x D_{yx} + q_y^2 D_{yy} + q_y q_z D_{yz} \\ &+ q_z q_x D_{zx} + q_z q_y D_{zy} + q_z^2 D_{zz} \\ &= q_x^2 D_{xx} + 2 q_x q_y D_{xy} + 2 q_x q_z D_{xz} \\ &+ \qquad q_y^2 D_{yy} + 2 q_y q_z D_{yz} \\ &+ \qquad\qquad + q_z^2 D_{zz} \end{aligned}$$

where we've used the fact that $\boldsymbol{D}$ is symmetric so we can collect off-diagonal terms. Recall that the symmetry of $\boldsymbol{D}$ in Eqn 29.7 derives from the physical property of our measurements that

the diffusion related signal loss is the same whether the net motion is along the positive or negative gradient direction[3]. If we visualize the diffusion gradient direction as a line that goes from the origin to a point on the surface of the sampling sphere, then the signal loss due to this gradient is the same if the gradient is applied in the opposite direction, which is represented by a line from the origin to a point directly on opposite sides of the sampling sphere, the so-called *antipodal* points[4]. This symmetric is therefore called *antipodal symmetry*. The result is that the off-diagonal terms are equal.

Now, a trick is usually pulled here in practice (**?**) is to note that

$$s_o e^{-\boldsymbol{bD}} = e^{\log(s_o)} e^{-\boldsymbol{bD}} = e^{-\boldsymbol{bD}+\log(s_o)} = e^{-\boldsymbol{Bd}+\log(s_o)} \tag{29.8}$$

Therefore, if we define a vector

$$\boldsymbol{\mathfrak{B}}^t = \left(q_x^2\,,\, q_y^2\,,\, q_z^2\,,\, q_x q_y\,,\, q_x q_z\,,\, q_y q_z\,,\, 1\right)\tau \tag{29.9}$$

and

$$\boldsymbol{d}^t = (D_{xx}\,,\, D_{yy}\,,\, D_{zz}\,,\, D_{xy}\,,\, D_{xz}\,,\, D_{yz}\,,\, -\log s_o) \tag{29.10}$$

and so the equation for our data looks like

$$\boldsymbol{s}(b) = \exp\left(-\boldsymbol{\mathfrak{B}}^t \boldsymbol{d}\right) + \boldsymbol{\eta} \tag{29.11}$$

This form is nice because it involves only two quantities, $\boldsymbol{\mathfrak{B}}$ which are the known gradient directions, and $\boldsymbol{d}$ which are the elements of the diffusion tensor we want to estimate. However, there are 7 unknowns in $\boldsymbol{d}$, so we need 7 equations to solve for them all. To do this, we need to acquire 7 different measurements. If all the measurements were pointing in the same directions, this would add no new information to the problem. So we need to acquire measurements *at different directions*. If we gather n directions, each produces a product of the form $\boldsymbol{\mathfrak{B}}_k^t \boldsymbol{d} \quad k = 1, \ldots, n$, so we can describe the whole problem by constructing a matrix where each $\boldsymbol{\mathfrak{B}}_k$ of the form Eqn 29.9 are the different rows

$$\underbrace{\begin{pmatrix} s(b_1) \\ s(b_2) \\ \vdots \\ s(b_n) \end{pmatrix}}_{\boldsymbol{y}} = \exp\left[- \underbrace{\begin{pmatrix} \boldsymbol{\mathfrak{B}}_1^t \\ \boldsymbol{\mathfrak{B}}_2^t \\ \vdots \\ \boldsymbol{\mathfrak{B}}_n^t \end{pmatrix}}_{\boldsymbol{\mathfrak{B}}} \boldsymbol{d} \right] + \underbrace{\begin{pmatrix} \eta(b_1) \\ \eta(b_2) \\ \vdots \\ \eta(b_n) \end{pmatrix}}_{\boldsymbol{\eta}} \tag{29.12}$$

where $\boldsymbol{\mathfrak{B}}$ is called the *b-matrix* (**?**). Writing it out in full, it is

$$\boldsymbol{\mathfrak{B}} = \begin{pmatrix} \hat{q}_{1,x}^2 & \hat{q}_{1,y}^2 & \hat{q}_{1,z}^2 & \hat{q}_{1,x}\hat{q}_{1,y} & \hat{q}_{1,x}\hat{q}_{1,z} & \hat{q}_{1,y}\hat{q}_{1,z} & 1 \\ \hat{q}_{2,x}^2 & \hat{q}_{2,y}^2 & \hat{q}_{2,z}^2 & \hat{q}_{2,x}\hat{q}_{2,y} & \hat{q}_{2,x}\hat{q}_{2,z} & \hat{q}_{2,y}\hat{q}_{2,z} & 1 \\ \vdots & & & & & & \\ \hat{q}_{n,x}^2 & \hat{q}_{n,y}^2 & \hat{q}_{n,z}^2 & \hat{q}_{n,x}\hat{q}_{n,y} & \hat{q}_{n,x}\hat{q}_{n,z} & \hat{q}_{n,y}\hat{q}_{n,z} & 1 \end{pmatrix} \tau \tag{29.13}$$

[3] True?!
[4] put earlier in the discussion of gradient directions! in dti-sensit. Show picture!

Alternatively, under ideal conditions in which the gradients are perfect, Eqn 29.13 can be rewritten in terms of the directional vector $\boldsymbol{q} = q\hat{\boldsymbol{u}}$, where $q = |q|$.

$$\mathfrak{B} = b \begin{pmatrix} \hat{u}_{1,x}^2 & \hat{u}_{1,y}^2 & \hat{u}_{1,z}^2 & \hat{u}_{1,x}\hat{u}_{1,y} & \hat{u}_{1,x}\hat{u}_{1,z} & \hat{u}_{1,y}\hat{u}_{1,z} & 1/b \\ \hat{u}_{2,x}^2 & \hat{u}_{2,y}^2 & \hat{u}_{2,z}^2 & \hat{u}_{2,x}\hat{u}_{2,y} & \hat{u}_{2,x}\hat{u}_{2,z} & \hat{u}_{2,y}\hat{u}_{2,z} & 1/b \\ \vdots & & & & & & \\ \hat{u}_{n,x}^2 & \hat{u}_{n,y}^2 & \hat{u}_{n,z}^2 & \hat{u}_{n,x}\hat{u}_{n,y} & \hat{u}_{n,x}\hat{u}_{n,z} & \hat{u}_{n,y}\hat{u}_{n,z} & 1/b \end{pmatrix} \tag{29.14}$$

where $b = q^2\tau$ and $q = g\delta$ and $\tau = (\Delta - \delta/3)$.

In practice (i.e., in an actual scanner), the full matrix $\boldsymbol{B}$ is calculated by incorporating the actual true waveforms on the scan (which might have slight variations from ideality) and, if one is being careful, the diffusion weighting effects of the imaging gradients. But Eqn 29.14 is useful in understanding the influence of gradient direction choices.

It is useful to recall at this point the relationship between the matrix $\mathfrak{B}$ and what is going on in the scanner. Recall that a single diffusion weighted direction is created by a combination of bipolar gradients along the three axes, as shown in Figure 29.1. The multiple measurements are made by suitable combinations of the gradients for producing diffusion weighting directions that are isotropically distributed on a spherical surface (by various methods, as discussed in Section 29.9 below). From these are determined the matrix $\mathfrak{B}$.

29.4 Estimation of the diffusion tensor D: Angular sampling

We have thus reduced the problem of estimating the diffusion tensor to finding the solution to the matrix equation

$$\boldsymbol{y} = \exp\left[-\mathfrak{B}\boldsymbol{d}\right] + \boldsymbol{\eta} \tag{29.15}$$

where

$$\begin{aligned} \boldsymbol{y} &= [n \times 1] \quad , \quad \text{data} \\ \mathfrak{B} &= [n \times m] \quad , \quad \text{B-matrix} \\ \boldsymbol{d} &= [m \times 1] \quad , \quad (s_o, \text{ and components of } \boldsymbol{D}) \end{aligned}$$

If there were no noise, we could just take the log of both sides of Eqn 29.15 to get $-\log(\boldsymbol{y}) = \mathfrak{B}\boldsymbol{d}$ and solving this by linear least squares (**?**). However, this approach is incorrect because the model (no noise) is incorrect (**?**): least squares assumes that the noise is additive and each sample has the same variance. Both of these assumptions are violated when the logarithm is taken. In addition, the $\boldsymbol{D}$ estimated from the log of the data is not guaranteed to have positive eigenvalues. [5] Rather, it is best to directly fit the data to the model. This can be accomplished using any number of fitting routines which are widely available, including in such standard platforms at MATLAB (**?**) and Mathematica (**?**). However, even fitting the model directly does not guarantee negative eigenvalues. Therefore, a fitting routine that also incorporates the constraint of positive eigenvalues is important in getting physically meaningful result. One such method (that we use in my lab) is that due to Cox (**?**), which is distributed in the AFNI package (**?**). This is a non-linear method that uses a gradient descent algorithm modified so that the estimated diffusion tensor

[5] More explanation here? Should have this discussion somewhere.

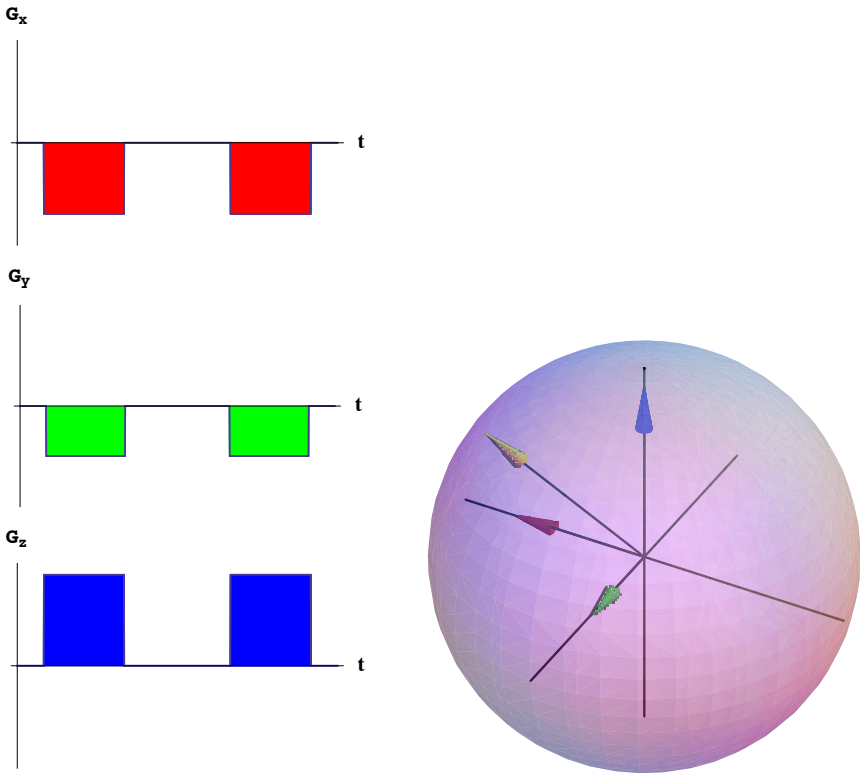

(a) Gradients (b) Diffusion weighting direction.

Figure 29.1 An example of gradients for a single diffusion encoding direction. The effects of gradients add like vectors so the direction and amplitude of the applied diffusion weighting gradient is defined by the vector sum of the gradients. The effect of gradients only along the x, y, or z axes is shown in (a-c), red, green, blue, respectively. With all gradients turned on simultaneously at different amplitudes produces and direction that is the vector sum (yellow).

remains positive definite (and thus has positive eigenvalues Section 8.12). Results of this method in comparison with the log-linear method are shown in Figure 29.2.[6]

29.5 Eigenvectors and eigenvalues of the diffusion tensor

As we saw in Chapter 24, the diffusion tensor $\boldsymbol{D}$ is symmetric and its elements are real. These two facts are expressed mathematically as:

$$\boldsymbol{D} = \boldsymbol{D}^t \quad \text{or} \quad D_{ij} = D_{ji} \quad , \quad \text{symmetric matrix} \tag{29.16}$$

$$D_{ij} \in \Re \quad , \quad \text{real matrix} \tag{29.17}$$

But from our discussion in Section 8.6 we can restate the spectral theorem in more practical terms: *The diffusion tensor $\boldsymbol{D}$ is a real, symmetric matrix and therefore can be rotated to make it diagonal.* That is, there exists some 3D rotation matrix $\boldsymbol{R}$ for which the matrix $\boldsymbol{D}$ (Eqn 29.7)

[6] This shows an image and so should be in dti-basic.tex where we first show images!

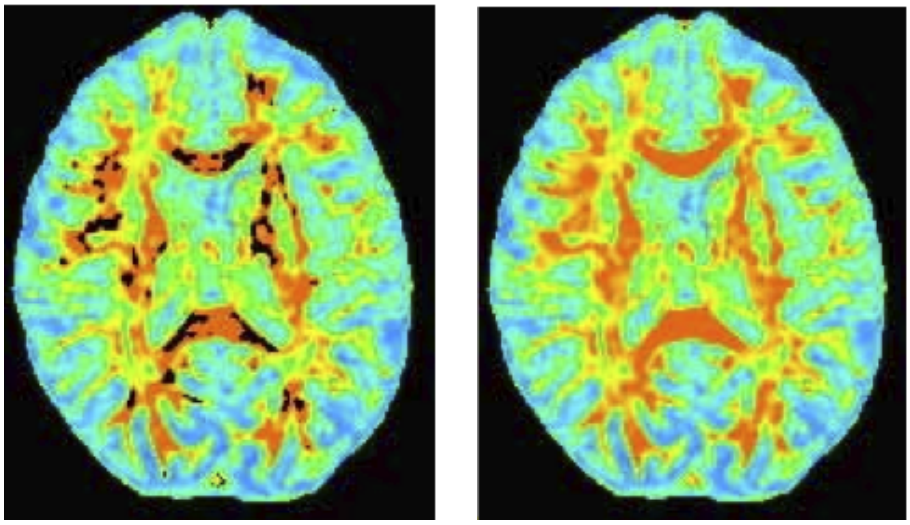

(a) Fitting the log-linear model produces spurious FA values (> 1) because $\boldsymbol{D}$ is no longer assured to be positive definite.

(b) Directly fitting the exponential form of the signal allows constraining $\boldsymbol{D}$ to be positive definite and consequently that $0 \leq FA \leq 1$.

Figure 29.2 Estimation of the diffusion tensor using the method of Cox (**?**) to directly fit the exponential decay. Shown is the fractional anisotropy.

can be transformed into a diagonal form $\boldsymbol{D}_\Lambda$:

$$\boldsymbol{D}_\Lambda = \boldsymbol{R}^t \boldsymbol{D} \boldsymbol{R} = \boldsymbol{R}^t \boldsymbol{R} = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \tag{29.18}$$

In practice, this procedure of *diagonalization* is done numerically by specialized algorithms designed specifically for this procedure.

We can now return to the results of Section 27.2 where we found that for free diffusion the probability distribution of the spins is given by Eqn 27.1

$$p(\bar{\boldsymbol{r}}, \tau) = \frac{1}{\sqrt{|\boldsymbol{D}|\,(4\pi\tau)^3}} \exp\left[-\bar{\boldsymbol{r}}^t \boldsymbol{D}^{-1} \bar{\boldsymbol{r}}/(4\tau)\right] \tag{29.19}$$

where $\bar{\boldsymbol{r}}$ is the mean displacement vector and $\tau = \Delta - \delta/3$. This distribution describes the spin displacement during the diffusion measurement period, $|.|$ is the determinant and $\boldsymbol{D}$ is the diffusion tensor. This is a multivariate Gaussian distribution with covariance matrix $2\tau\boldsymbol{D}$. The covariance matrix determines the width of the distribution and thus describes how much the spins have "spread out" during the measurement time. But we studied the multivariate Gaussian in-depth in Section 8.7 and we can now use those results.

So for a tissue where the diffusion varies as a function of direction, the distribution of described by the multivariate Gaussian of Eqn 8.48, with the variance $\sigma_i^2 = 2D_i\tau$ where $i = \{x, y, z\}$. In

other words, the covariance matrix is

$$\boldsymbol{M} = \begin{pmatrix} 2D_x\tau & 0 & 0 \\ 0 & 2D_y\tau & 0 \\ 0 & 0 & 2D_z\tau \end{pmatrix} \tag{29.20}$$

where the eigenvalues are

$$\lambda_1 = 2D_x\tau \qquad , \qquad \lambda_2 = 2D_y\tau \qquad , \qquad \lambda_3 = 2D_z\tau \tag{29.21}$$

Because of the common constant factor of 2τ in these equations, it is more efficient to write the covariance matrix as $\boldsymbol{M} = 2\tau\boldsymbol{D}$ where

$$\boldsymbol{D}_\Lambda = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \tag{29.22}$$

where $\boldsymbol{D}_\Lambda$ is called the *diffusion tensor* and has eigenvalues

$$\{\lambda_1, \lambda_2, \lambda_3\} = \{D_x, D_y, D_z\} \tag{29.23}$$

Note that, because the diffusion coefficients are related to the variances, which are inversely proportional to the eigenvalues, that they are thus linearly related to the eigenvalues. That is, *the eigenvalues are the diffusion coefficients!* Again, this is the covariance matrix structure if the variance directions - i.e., the diffusion coefficients directions - are aligned with the coordinate (spatial) axes $\{x, y, z\}$. If they are not, then the covariance matrix is not diagonal. But, again, we can make it so, and thus determine the eigenvalues and eigevectors, which in turn gives us the diffusion coefficients and the directions of diffusion. The orientation of the diffusion tensor relative to the magnet coordinates depends upon the orientation of the tissue relative to these coordinates, and in generate is rotated. The diffusion tensor in the frame in which it is diagonal (the tissue frame) is related to the magnet frame by the similarity transform we discussed in Section 8.5. The diffusion tensor is thus (Eqn 27.14)

$$\boldsymbol{D} \equiv \boldsymbol{R}^t\, \boldsymbol{D}_\Lambda\, \boldsymbol{R} \tag{29.24}$$

The situation is just as we discussed in Section 8.7: the contours of probability for the multivariate Gaussian diffusion are ellipsoids, as shown in Figure 29.3.

29.6 The average eigenvalue: the mean diffusivity

A natural question to ask is "What is the average diffusion $\overline{\boldsymbol{D}}$ in the voxel?" where $\boldsymbol{D}$ is given, generally, by Eqn 29.7. Well, how do you take the average of this? Let's first look at $\boldsymbol{D}$ in the coordinate system in which it is diagonal, that is, in the eigencoordinates of $\boldsymbol{D}$. Here the average is easy:

$$\langle D \rangle = \frac{1}{3}\left(D_{e_1,e_1} + D_{e_2,e_2} + D_{e_3,e_3}\right) = \frac{1}{3}\left(\lambda_1 + \lambda_2 + \lambda_3\right) = \frac{1}{3}\mathrm{Tr}(\boldsymbol{D}) \tag{29.25}$$

where we've used the fact that the trace Tr(.) of a matrix is (Section 5.6) is the sum of the diagonal elements of the matrix. Therefore $\mathrm{Tr}(\boldsymbol{D})$ is a measure of the *total* diffusion in a voxel and $\frac{1}{3}\mathrm{Tr}(\boldsymbol{D})$ is the *average* over the different directions of the eigenvectors, and is called the *mean diffusivity*. Since each measurement provides an estimate of the *apparent diffusion coefficient* or

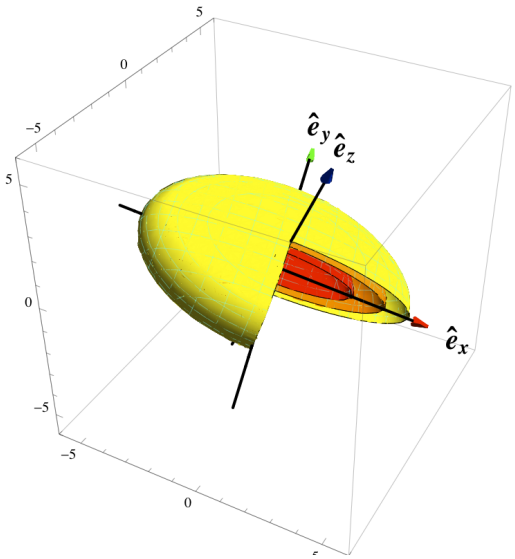

Figure 29.3 Estimation of the diffusion tensor eigenvalues and eigenvectors allows reconstruction of the Gaussian probability distribution whose contours of equal probability are ellipsoidal. The principle axes of the ellipsoid are the eigenvectors of the diffusion tensor. The lengths of the eigenvectors are proportional to the square-root of the diffusion tensor eigenvalues λ times the diffusion time τ.

ADC along that direction, the $\langle D \rangle$ is the average ADC. When we say "average", we really need to be more specific and say "average over directions". Charaterizing these directions as different angles $\Omega \equiv \{\vartheta, \varphi\}$, we should write $\langle D \rangle_\Omega$ which symbolizes that the average is over the values of Ω at which the measurements were taken.

Now let's return to our original problem of determining $\langle \boldsymbol{D} \rangle$ for $\boldsymbol{D}$ given by Eqn 29.7, which is related to $\boldsymbol{D}_\Lambda$ by Eqn 29.24. But now recall (Eqn 8.43) the invariance of the trace to cyclic permutation of the elements:

$$\mathrm{Tr}(\boldsymbol{D}) = \mathrm{Tr}(\boldsymbol{R}\boldsymbol{\Lambda}\boldsymbol{R}^t) = \mathrm{Tr}(\boldsymbol{R}^t\boldsymbol{R}\boldsymbol{\Lambda}) = \mathrm{Tr}(\boldsymbol{\Lambda}) \tag{29.26}$$

and the fact that rotation matrices are orthonormal (Section 5.16): $\boldsymbol{R}^t\boldsymbol{R} = 1$, so that

$$\mathrm{Tr}(\boldsymbol{D}) = \mathrm{Tr}(\boldsymbol{\Lambda}) = \sum_{i=1}^{3} \lambda_i = 3\overline{\lambda} \tag{29.27}$$

where $\overline{\lambda} = \frac{1}{3}\sum_{i=1}^{3} \lambda_i$ is the average eigenvalue. This tells us mathematically what you had probably already guessed - the average diffusion should not change just because we changed the orientation of the probability distribution. But it also tells us something that is useful computationally: The trace of $\boldsymbol{D}$ of the form Eqn 29.7 is the same as the trace of $\boldsymbol{D}_\Lambda$ in its diagonal form. So it is not necessary to diagonalize the matrix in order to determine the trace - you just sum the elements along the diagonal and it works out the same! This scalar quantity $\langle D \rangle$ is invariant to orientation and thus is called a *scalar invariant*.

An intuitive picture of what is happening physically is provided by a simple geometrical depiction. Denoting $\delta\lambda_i = \lambda_i - \overline{\lambda}$, the difference of the i'th eigenvalue from the mean, we can rewrite

Eqn 29.22 in the form

$$\boldsymbol{D}_\Lambda = \underbrace{\begin{pmatrix} \overline{\lambda} & 0 & 0 \\ 0 & \overline{\lambda} & 0 \\ 0 & 0 & \overline{\lambda} \end{pmatrix}}_{\overline{\boldsymbol{D}}} + \underbrace{\begin{pmatrix} \delta\lambda_1 & 0 & 0 \\ 0 & \delta\lambda_2 & 0 \\ 0 & 0 & \delta\lambda_3 \end{pmatrix}}_{\delta\boldsymbol{D}} \tag{29.28}$$

where we've used the fact (Section 5.3) that we can write a matrix as the sum of other matrices. We've now written $\boldsymbol{D}$ in terms of a diagonal matrix $\overline{\boldsymbol{D}}$ where all the diagonal elements are equal and another diagonal matrix $\delta\boldsymbol{D}$ where the elements are the deviations from the mean $\overline{\lambda}$. But a diagonal matrix with all the elements equal has the geometrical representation of a sphere - the eigenvectors are all of the same length. This means that the contour of constant probability are *spherical*, as shown in Figure 29.4. This represents the average diffusion over the different directions and is independent of orientation. Recall from our discussion in Section 8.5 that the trace is invariant to rotations of the matrix so that

$$\mathrm{Tr}(\boldsymbol{D}) = \underbrace{\begin{pmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{pmatrix}}_{D_{xx}+D_{yy}+D_{zz}} = \underbrace{\begin{pmatrix} \lambda_x & 0 & 0 \\ 0 & \lambda_y & 0 \\ 0 & 0 & \lambda_z \end{pmatrix}}_{\lambda_x+\lambda_y+\lambda_z} \tag{29.29}$$

29.7 The variance for the eigenvalues: diffusion anisotropy

The matrix $\delta\boldsymbol{D}$ contains the deviations of the eigenvalues from their average value, and thus is related to the differences between the diffusion along the principle axes. If there we no variations, i.e., $\delta\lambda_1 = \delta\lambda_2 = \delta\lambda_3 = 0$, then the diffusion is the same in all directions and is said to be *isotropic*. But if any of the $\delta\lambda_i$ are different, then the diffusion is not equal along all directions, and is said to be *anisotropic*. A natural way to characterize these differences is to look at the *variance* of the eigenvalues. We can rearrange Eqn 29.28 to write the matrix of deviations as

$$\delta\boldsymbol{D} = \begin{pmatrix} \delta\lambda_1 & 0 & 0 \\ 0 & \delta\lambda_2 & 0 \\ 0 & 0 & \delta\lambda_3 \end{pmatrix} = \begin{pmatrix} \lambda_1 - \overline{\lambda} & 0 & 0 \\ 0 & \lambda_2 - \overline{\lambda} & 0 \\ 0 & 0 & \lambda_3 - \overline{\lambda} \end{pmatrix} \tag{29.30}$$

Now these deviations $\delta\lambda_i$ can be either positive or negative, depending upon whether the particular eigenvalue is greater than or less than the mean. But if we are only interested in how much our shape deviates from a sphere, that is, how much our diffusion deviates from isotropic, we only care about the magnitude of the difference. So what the average magnitude of the deviation? Well, this is just

$$\sqrt{\overline{(\delta\lambda)^2}} = \sqrt{\frac{1}{3}\sum_{i=1}^{3}\delta\lambda_i^2} \tag{29.31}$$

But this is just

$$\sqrt{\overline{(\delta\lambda)^2}} = \sqrt{\frac{1}{3}\delta\boldsymbol{D}^t\delta\boldsymbol{D}} = \sqrt{\frac{1}{3}\left[(\lambda_1-\overline{\lambda})^2 + (\lambda_2-\overline{\lambda})^2 + (\lambda_3-\overline{\lambda})^2\right]} = \sqrt{\frac{1}{3}\sigma_\lambda^2} \tag{29.32}$$

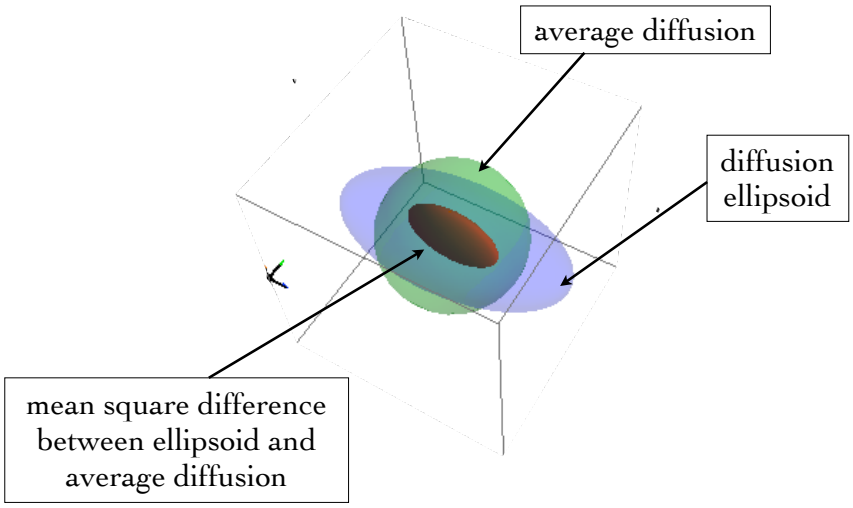

Figure 29.4 The diffusion ellipsoid, the average diffusion, and the diffusion anisotropy.

where σ_λ^2 is just the variance of the eigenvalues. Now in an experimental situation where we have, for instance, a voxel containing white matter, the tissue will are often interested in both the mean diffusion and the diffusion anisotropy. The diffusion ellipsoid, the average diffusion, and the diffusion anisotropy are depicted graphically in Figure 29.4.

But the variance in and of itself may not be as important as how much of the how large these variations are *relative to the mean squared diffusion in the voxel*, i.e., if all $\delta_i = 0$, then

$$\sqrt{\overline{\lambda^2}} = \sqrt{\frac{1}{3}\left(\lambda_1^2 + \lambda_2^2 + \lambda_3^2\right)} = \sqrt{\frac{1}{3}\sum_i \lambda_i^2} = \sqrt{\frac{1}{3}\boldsymbol{D}_\Lambda^t \boldsymbol{D}_\Lambda} \tag{29.33}$$

So the fraction that the variations make up of the mean squared diffusion is, from Eqn 29.33 and Eqn 29.32 is

$$\sqrt{\frac{\overline{(\delta\lambda)^2}}{\overline{\lambda^2}}} = \sqrt{\frac{\sigma_\lambda^2}{\boldsymbol{D}_\Lambda^t \boldsymbol{D}_\Lambda}} \tag{29.34}$$

This is a natural measure of anisotropy. Notice that if

$$(\lambda_1, \lambda_2, \lambda_3) = (1, 0, 0) \tag{29.35}$$

then

$$\begin{aligned}
\overline{\lambda^2} &= \frac{1}{3}(\lambda_1^2 + \lambda_2^2 + \lambda_3^2) = \frac{1}{3} \\
\overline{(\delta\lambda)^2} &= \frac{1}{3}\left[\left(\frac{2}{3}\right)^2 + \left(\frac{1}{3}\right)^2 + \left(\frac{1}{3}\right)^2\right] = \frac{2}{9} \\
\Rightarrow \sqrt{\frac{\overline{(\delta\lambda)^2}}{\overline{\lambda^2}}} &= \sqrt{\frac{2/9}{1/3}} = \sqrt{\frac{2}{3}}
\end{aligned}$$

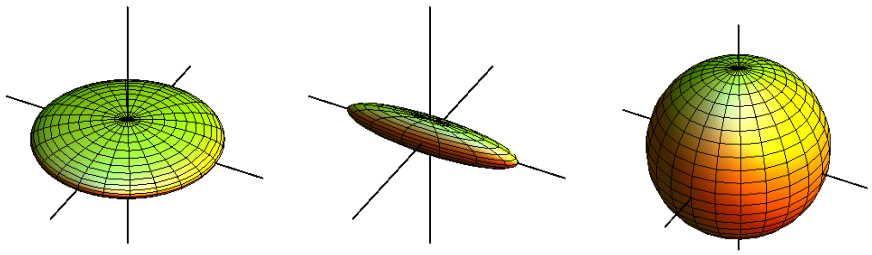

(a) Oblate. $\lambda_1 = \lambda_2 \gg \lambda_3$ (b) Prolate. $\lambda_1 \gg \lambda_2 \approx \lambda_3$ (c) Spherical. $\lambda_1 \approx \lambda_2 \approx \lambda_3$

Figure 29.5 Diffusion ellipsoid shape classifications for special cases of eigenvalue combinations.

So normalizing Eqn 29.34 so that the maximum value is 1 we can define

$$FA = \sqrt{\frac{3}{2}\frac{\overline{(\delta\lambda)^2}}{\overline{\lambda^2}}} \quad , \qquad \text{Fractional anisotropy} \tag{29.36}$$

Eqn 29.36 is called the *fractional anisotropy* or *FA*. It is the most common index of anisotropy used in practice. [7]. It is quite remarkable that this rather straightforward calculation gives what appears to be a remarkably good image of the white matter! [8]

29.8 Other Anisotropy Measures

There are some practical cases in which there are two (nearly) equal eigenvalues, in which case there is a nice shape description for the diffusion ellipsoid in terms of *oblate* and *prolate* ellipsoids, as shown in Figure 29.5. Sometimes these shapes are called *linear* and *planar* rather than prolate and oblate, respectively. These three shapes can be characterized by the following anisotropy indeces (**?**):

$$\text{linear}: \qquad c_l = \frac{\lambda_1 - \lambda_2}{3\bar{\lambda}} \tag{29.37a}$$

$$\text{planar}: \qquad c_p = \frac{2(\lambda_2 - \lambda_3)}{3\bar{\lambda}} \tag{29.37b}$$

$$\text{spherical}: \qquad c_s = \frac{\lambda_3}{\bar{\lambda}} \tag{29.37c}$$

where

$$\bar{\lambda} \equiv \frac{1}{3}(\lambda_1 + \lambda_2 + \lambda_3) \tag{29.38}$$

is the average eigenvalue.

[7] Show a single voxel example here!
[8] Check all the FA stuff for the constants!

29.9 How do we choose our angular sampling?

In Section 29.4 we noted that we need at least 7 measurements to estimate the six unique components of $\boldsymbol{D}$ plus the normalizing factor s_o. But what 7 measurements? And is 7 really enough? And what do we mean by "enough"? [9] For angular measurements where $\boldsymbol{b}$ is a matrix, it is required that $\boldsymbol{b}^1$ is *invertible*. As we discussed at length in Eqn 5.12, in order for a matrix to be invertible, there are certain restrictions on the relationship amongst the column vectors that compose it. In particular, they must be *linearly independent* (**CHECK!**). This means that none of the vectors can be colinear with any of the other vectors. In other words, *the vectors must be non-colinear*. That is, they must all point in different directions. But is that enough? Intuitively, you could imagine that sampling n non-colinear directions that lie on the surface of a cone with a very angle[10] is not taking measurements in directions that are very different, and so must have trouble detecting the ellipsoidal variations in the MR signal as a function of angle. How do we quantify this intuitive concept?

The answer to this question was found in our discussion of the amplitude estimates in our matrix formulation of the parameter estimation problem in probability theory (Section 12.9). The amplitude estimates depend on the *pseudo-inverse*, which we found in Eqn 12.97. For the 3-dimensional problem here, this means that there are 3 non-zero eigenvalues of the the pseudo-inverse of the b-matrix, $\boldsymbol{b}^{(-1)} = (\boldsymbol{b}^t\boldsymbol{b})^{-1}\boldsymbol{b}^t$. The angular spacing of the samples decides the relative weighting of these eigenvalues. In order for our estimates to be equally weighted, we want these eigenvalues to be approximately equal. This is equivalent to having the samples isotropically distributed about the sphere [11]. This is shown in Figure 34.7 where we compare the tesselation method (Figure 34.7a), which produces essentially an equi-angular distribution, with a random distribution of angles (Figure 34.7b). Note that the equi-angular distribution produces a strong diagonal component in the pseudo-inverse matrix. This, then, conforms with our intuition: we want samples that are equally spaced on a sphere.

Having decided that we want samples equally distributed on a sphere, how can this be achieved numerically? One common method, illustrated in Figure 29.7, is to use the vertices of a platonic solid. These vertices are well known and can be easily programmed into the pulse sequence software. One limitation of this method, however, is that the platonic solids do not have many vertices. As we shall see in later chapters, it is often necessary to collect many different directions (say, 64). However, there is a nice way to extend the platonic sampling method to more directions. Each "face" of the solid, defined by the flat surface between the vertices, is cut up into smaller versions of the same shape by cutting the lines between the vertices in half. This process is called *tessellation*, and is illustrated for triangular faces in Figure 29.8a. This process can be repeated *ad infinitum* on the new, smaller faces. An example of a tessellated icosahedron is shown in Figure **??**. However, this method still has a major limitation: For any given platonic solid, there are only a discrete set of directions available, since the tesselation procedure produces a fixed number of additional vertices. For example, in Figure 29.8a we see that each triangular face defined by 3 vertices becomes 4 faces with 6 unique vertices. Moreover, the number of vertices grows very rapidly, since this procedure takes place at every face, generated more faces, each which then gets tessellated again. For example, the icosahedron in Figure **??** has 12 vertices, and

[9] Haven't discussed total number of samples yet!
[10] what is that cone angle called?
[11] Why?!

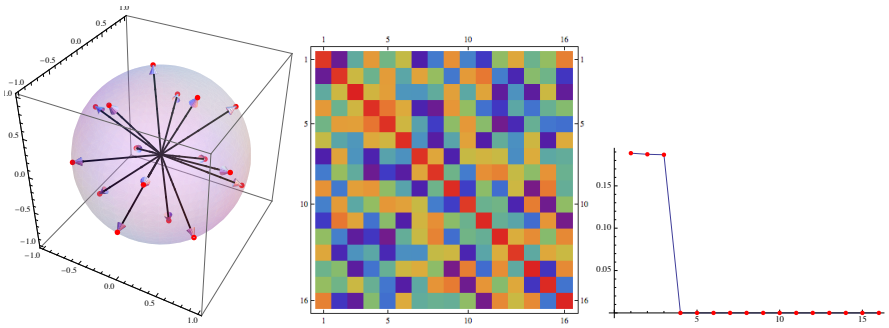

(a) Approximately equally spaced angularly through tesselation. (left) Sampling; (center) b-matrix; (right) eigenvalue spectrum of b-matrix.

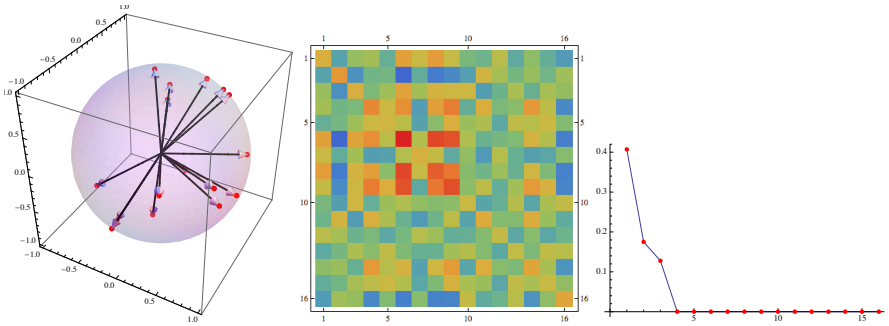

(b) Randomly spaced. (left) Sampling; (center) b-matrix; (right) eigenvalue spectrum of b-matrix.

Figure 29.6 Eigenvalues of $\boldsymbol{b}^{(-1)} = (\boldsymbol{b}^t\boldsymbol{b})^{-1}\boldsymbol{b}^t$, the pseudo-inverse of the b-matrix. (**Say something about eigenvalue spectra!**)

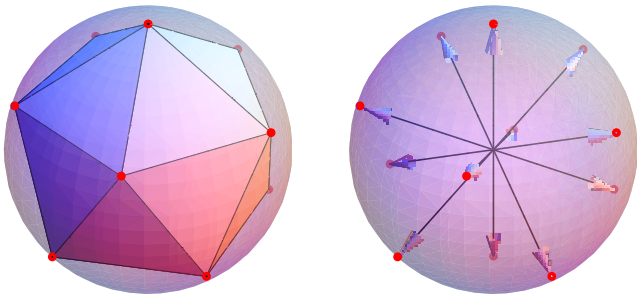

Figure 29.7 Isotropically distributed diffusion directions samples chosen as the vertices of a platonic solid (an icosahedron).

on subsequent levels of tessellation produces $\{42, 162, 642, 2562\}$ vertices. In a standard imaging protocol a "large" number of directions is around 60, which means that only the second level of tesselation (i.e., 42 vertices) is allowed. Thus the tesselation method is quite restrictive in what directions it produces.

An alternative strategy that allows an essentially arbitrary number of equally points on a sphere to be generated is to utilize what is a numerical method wherein a desired number of

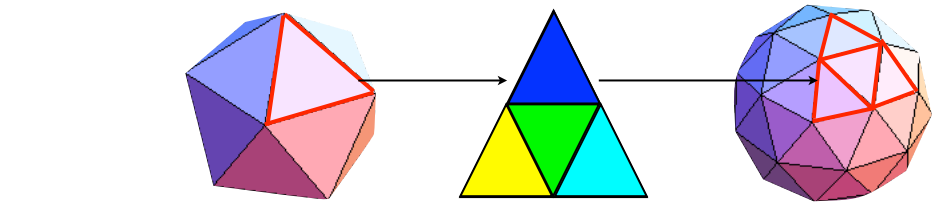

(a) Tessellation.

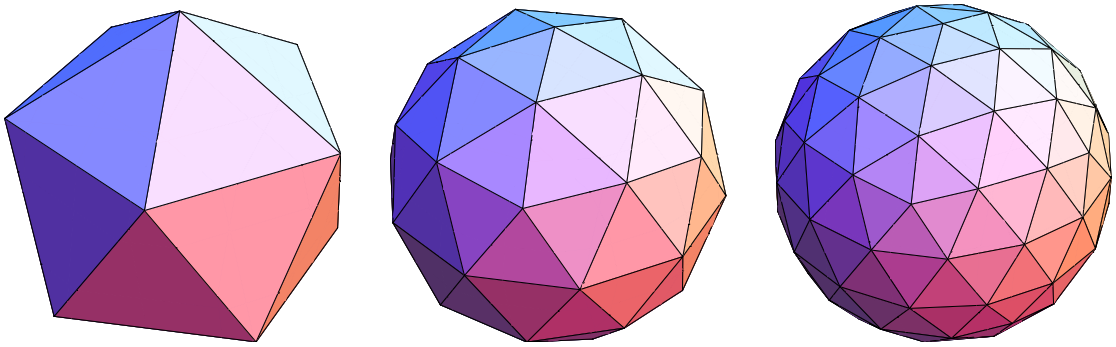

(b) $n_{tess} = 1$ (c) $n_{tess} = 2$ (d) $n_{tess} = 3$

Figure 29.8 Tessellation of a platonic solid (icosahedron) to generate (approximately) equi-angular diffusion encoding directions. n_{tess} is the tesselation level.

electrically charged particles (i.e., one for each diffusion directions) is put on the surface of a conducting sphere. These particles tend to repel one another and so their state of equilibrium is to distribute themselves isotropically around the sphere, thereby minimizing the potential energy of the system, and maximizing the mean distance between particles. This is called on *electrostatic repulse model* (**?**). Where these charges finally settle is taken to be the final points from which the diffusion directions are chosen (**?**). This method can be run on an increasing number of particles and the results stored in a table. The pulse sequence then just choses the results from the table entry for the requested number of points. An example of the electrostatic repulsion model is shown in Figure 29.9.

29.10 Non-linear sampling*

Of course, in most applications the noise will not be known, so sampling at only two points would be problematic. So, if a limited number of points of an exponential are to be collected, how should they be spaced? In most applications, such as the sine wave problem above, data are collected in evenly spaced increments in time. For sinusoids, the important criterion is that the samples be spaced close enough to satisfy the Nyquist criterion (Section 19.7.1). But for the decay problem we're fighting a different demon - the falloff of the signal into the noise. Would it not be better to get more samples earlier in the signal before it dies away into the noise? Intuitively, this makes sense, and it is true. An example is shown in Figure 29.10 where the sampling pattern, often called the *schedule*, is logarithmic, so that the density of sample points is highest at the start and decreases with b. An example of logarithmic sampling is shown in Figure 29.11.

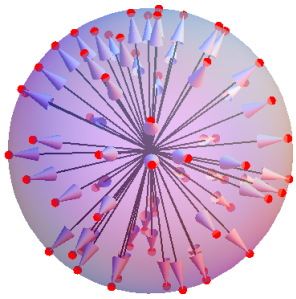

(a) Diffusion directions generated from the origin to the points in (a).

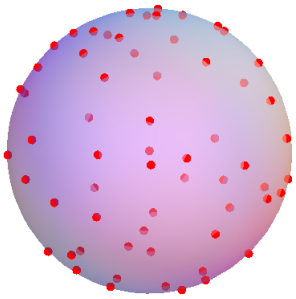

(b) Points on a sphere generated by the electrostatic repulsion model

Figure 29.9 The electrostatic repulse model. Points are chosen so as to minimize the energy of charged points on a sphere.

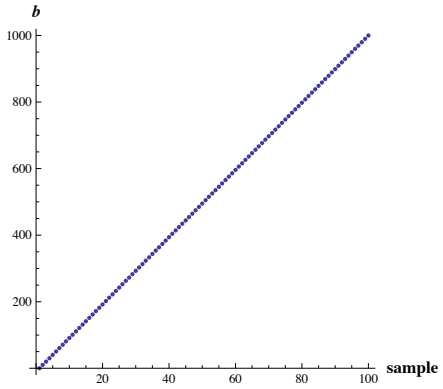

(a) Linear sampling of b.

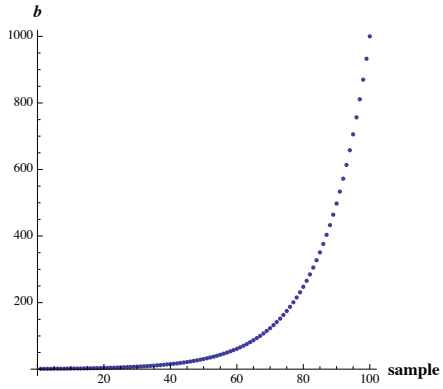

(b) Logarithmic sampling of b.

Figure 29.10 The default strategy for sampling is usually equally spaced samples, as in (a). For the decay problem, however, a higher sampling density during the initial part of the signal allows the acquisition of more data with high *SNR* than equally spaced samples.

Suggested reading

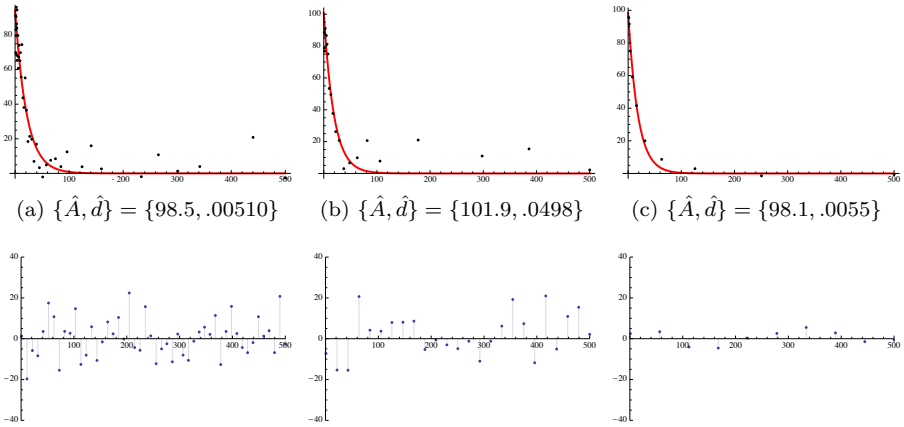

(a) $\{\hat{A}, \hat{d}\} = \{98.5, .00510\}$ (b) $\{\hat{A}, \hat{d}\} = \{101.9, .0498\}$ (c) $\{\hat{A}, \hat{d}\} = \{98.1, .0055\}$

(d) residual = 1.69 (e) residual = 2.03 (f) residual = 2.44

Figure 29.11 Logarithmic sampling of the exponential decay for fixed $SNR = 10$ decay rate $d = .050$ as a function of the number os samples. The amplitude is $A = 100$. Reasonable estimates are possible even with very few samples.