Magnetic Resonance in Medicine 47:1083–1099 (2002)

Characterization of Anisotropy in High Angular Resolution Diffusion-Weighted MRI

Lawrence R. Frank*

The methods of group theory are applied to the problem of characterizing the diffusion measured in high angular resolution MR experiments. This leads to a natural representation of the local diffusion in terms of spherical harmonics. In this representation, it is shown that isotropic diffusion, anisotropic diffusion from a single fiber, and anisotropic diffusion from multiple fiber directions fall into distinct and separable channels. This decomposition can be determined for any voxel without any prior information by a spherical harmonic transform, and for special cases the magnitude and orientation of the local diffusion may be determined. Moreover, non-diffusion–related asymmetries produced by experimental artifacts fall into channels distinct from the fiber channels, thereby allowing their separation and a subsequent reduction in noise from the reconstructed fibers. In the case of a single fiber, the method reduces identically to the standard diffusion tensor method. The method is applied to normal volunteer brain data collected with a stimulated echo spiral high angular resolution diffusion-weighted (HARD) acquisition. Magn Reson Med 47:1083–1099, 2002. Published 2002 Wiley-Liss, Inc.†

Key words: high angular resolution diffusion; diffusion anisotropy; spherical harmonic decomposition; spherical harmonic transform; white matter diffusion

The sensitivity of the magnetic resonance (MR) signal to molecular diffusion provides the most sensitive noninvasive method for the measurement of local tissue diffusion characteristics (for a review, see Ref. 1). The basic effect from which diffusion information is derived is the signal diminution due to diffusive motions along the direction of an applied gradient field (2,3). The fact that diffusion can have a directional dependence was recognized early on (4), but there has been a great resurgence of interest in the application to diffusion-weighted imaging of human tissues, where inferences about tissue structure can be made from the directional dependence it imposes on the local diffusion. Anisotropic diffusion was first demonstrated in the brain by Moseley et al., (5,6) and has been used to study a variety of other tissues (7). The determination of anisotropy requires a reconstruction of the local apparent diffusion, D_{app}. If the diffusion process has spatially homogeneous Gaussian increments, the directional dependence can be completely characterized by the diffusion tensor D. This relates the signal loss along an applied gradient in an orthogonal Cartesian system defined by the imaging coordinate system to the diffusion along a direction in an arbitrarily rotated orthogonal system defined by the tissue (8). Reconstruction of local Gaussian diffusion can thus be posed as estimating the diffusion tensor (8), which, in principle, requires only six measurements plus an additional measurement for normalization (9). This technique is of particular interest in its application to the characterization of white matter tracts (10,11).

However, it has long been recognized that the Gaussian model for diffusion can be inappropriate within the complex structure of human tissues (8,12). One way in which this model can fail is the presence of multiple fiber directions within a single imaging voxel. Because the single fiber diffusion tensor model is no longer appropriate, the characterization of the diffusion in such voxels becomes problematic. A novel approach to this problem, proposed by Tuch et al.(13), is to map D_{app} at high angular resolution in order to more accurately detect variations in diffusion along different directions. This has been extended to a scheme by Wedeen et al. (14) in which measurements through a range of diffusion sensitivities are made. There remains, however, no method for characterizing the diffusion measured by these high angular resolution diffusion-weighted (HARD) methods.

In a recent paper (15) we proposed a simple method by which anisotropy can be detected from HARD measurements by using the variations of D_{app} from a sphere as a measure of deviation from isotropy. This method has the advantage that it does not depend upon any particular diffusion model, but simply indicates whether the diffusion is not isotropic. It has the disadvantage that it does not characterize the structure of the diffusion and is sensitive to artifacts. Motivated by the inherent spherical symmetry in the HARD technique, we present here a formalism to characterize diffusion based on the mathematical discipline of group theory, which is a powerful tool in the study of systems with intrinsic symmetry. We show that HARD measurements can be decomposed into isotropic, single-fiber, and multiple-fiber components. The theory reduces to the standard diffusion tensor model in the case of a single fiber. A new transform is introduced with which the composition of a voxel in terms of the three subspaces can be determined, as well as the magnitude and orientation of the local diffusion. Moreover, asymmetries produced by experimental artifacts fall into channels distinct from the fiber channels, thereby allowing their separation and a subsequent reduction in noise from the reconstructed fibers. The method is demonstrated by simulation and on normal human volunteers. A preliminary report of this method and its rationale were described in a recent abstract (16).

MR DIFFUSION MEASUREMENTS

Before presenting our strategy for characterizing diffusion anisotropy in multifiber systems, we summarize in this

Department of Radiology, University of California-San Diego, San Diego, California; VA San Diego Healthcare System, San Diego, California.

Grant sponsor: VA Merit Award.

*Correspondence to: Lawrence R. Frank, Ph.D., USCD Center for Functional MRI, 9500 Gilman Drive, Mail Code 0677, LaJolla, CA 92093-0677. E-mail: lfrank@ucsd.edu

Received 23 April 2001; revised 28 December 2001; accepted 4 January 2002.

DOI 10.1002/mrm.10156

Published online in Wiley InterScience (www.interscience.wiley.com).

section the basic mathematical underpinnings of MR diffusion measurements. In what follows, a "fiber" will be defined as a particular diffusion tensor $\boldsymbol{D}$. With this definition, a large bundle of parallel fibers would be synonymous with a single "fiber." While it is natural to confer cylindrical symmetry on a diffusion tensor as part of the definition of a fiber, we relax this restriction in order to more clearly establish where such symmetry actually affects the more general results.

Single Fiber

We first consider diffusion measurements made in a voxel containing a single fiber, following Hsu and Mori (17) throughout. The signal attenuation can be written

$$S(b, D) = S_0 e^{-bD_{app}} \quad [1]$$

where b is related to the k-space trajectory $k(t)$ by $b \equiv \int_0^{TE} k^2(t)\, dt$ and incorporates the gradient strengths. The gradient directions can be defined by the unit vectors $\boldsymbol{u}$. If the measurements are made along the principal axes, i.e. in the coordinate system in which the diffusion tensor is diagonal, the apparent diffusion coefficient D_{app} is (17)

$$D_{app} = \boldsymbol{u}^T \cdot \boldsymbol{D}_\Lambda \cdot \boldsymbol{u} \quad [2]$$

and the applied diffusion-encoding gradients $\boldsymbol{u}$ are in the direction of the eigenvectors of $\boldsymbol{D}$.

$$\boldsymbol{u} = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}, \quad \boldsymbol{D}_\Lambda = \begin{pmatrix} d_1 & 0 & 0 \\ 0 & d_2 & 0 \\ 0 & 0 & d_3 \end{pmatrix} \quad [3]$$

where the eigenvalues are the principal diffusivities $\{d_1, d_2, d_3\}$. Generally, however, this principal axis coordinate system is not known. The applied diffusion-encoding gradients v are therefore not coincident with the principal axis system, but are related to it by a rotation $\boldsymbol{R}$:

$$\boldsymbol{u} = \boldsymbol{R}\boldsymbol{v} \quad [4]$$

where $\boldsymbol{v}$ is a unit vector in the direction of diffusion encoding.

Thus, one usually wants to infer the principal diffusivities and the rotation R. From these can be determined the diffusion properties, such as the anisotropy, and the fiber directions. The rotation R is defined within the coordinate system shown in Fig. 1. The two angles that define the direction in this coordinate system are the polar angle $\theta \in [0, \pi]$, which is defined as the angle between the vector and the positive z-axis, and the azimuthal angle $\phi \in [0, 2\pi)$, which is defined as the angle in the $x-y$ plane relative to the positive x axis. It is also common to use the elevation angle $\delta = 90° - \theta$, which is the angle between the vector and the $x-y$ plane. These definitions are sometimes reversed in the mathematics literature. We retain the standard physics usage, depicted in Fig. 1, where (θ, ϕ) denote the polar and azimuthal angles, respectively. We will often employ the shorthand notation $\Omega \equiv (\theta, \phi)$.

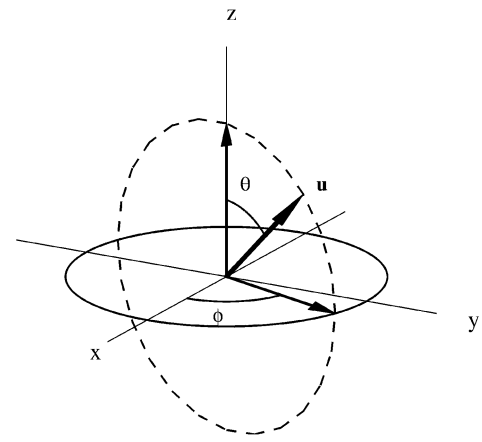

FIG. 1. Spherical coordinate system.

It is helpful to note that (θ, ϕ) are two of the Euler angles (18) used to described rotations in 3-D coordinates. These angles are typically denoted as (α, β, γ), where α is the azimuthal rotation angle, β is the polar rotation angle, and γ is a rotation about the new axis defined by the rotation through (α, β). For the description of a single point (i.e., a measurement) on a sphere, as is the case in this work, rotations about the final (radial) axis are unimportant, so the rotations can be described by the two angles (α, β). It is common in this case to denote these as (θ, ϕ).

The gradient direction vectors in the two coordinate systems are related by a rotation (17)

$$\boldsymbol{R} = \begin{pmatrix} \sin\phi & -\cos\phi & 0 \\ \cos\phi\cos\theta & \sin\phi\cos\theta & -\sin\theta \\ \cos\phi\sin\theta & \sin\phi\sin\theta & \cos\theta \end{pmatrix}. \quad [5]$$

The apparent diffusion coefficient for an arbitrary gradient direction v can thus be written (17)

$$D_{app} = \boldsymbol{v}^T \cdot \boldsymbol{D} \cdot \boldsymbol{v} \quad [6]$$

where

$$\boldsymbol{D} \equiv \boldsymbol{R}^T \boldsymbol{D}_\Lambda \boldsymbol{R}. \quad [7]$$

Equation [7] defines the diffusion tensor $\boldsymbol{D}$ in a rotated coordinate system. The signal from a single fiber is typically expressed in the form

$$\log(S/S_0) = -b\hat{D} \quad [8]$$

where we introduce the shorthand notation $\hat{D} \equiv D_{app} = \boldsymbol{v}^T\boldsymbol{D}\boldsymbol{v}$. For any symmetric matrix $\boldsymbol{D}$, such as the diffusion tensor, the product $\boldsymbol{x}^T\boldsymbol{D}\boldsymbol{x}$ is a pure quadratic form (19). From Eq. [8] one can estimate the diffusion tensor $\boldsymbol{D}$ by an eigenvalue decomposition in which the eigenvectors effectively determine the rotation of the fiber coordinate system relative to the laboratory system, and the eigenvalues determine the diffusivities (8). Since $\boldsymbol{D}$ is positive definite, it can be written in the form of Eq. [7], where $\boldsymbol{D}_\Lambda$ is diagonal

and the unit eigenvectors of $\boldsymbol{D}$ are the columns of $\boldsymbol{R}$. The rotation $\boldsymbol{y} = \boldsymbol{R}\boldsymbol{v}$ produces the sum of squares

$$\boldsymbol{v}^T\boldsymbol{D}\boldsymbol{v} = \boldsymbol{v}^T\,\boldsymbol{R}\boldsymbol{D}_\Lambda\boldsymbol{R}^T\boldsymbol{v} = \boldsymbol{y}^T\boldsymbol{D}_\Lambda\boldsymbol{y} = \sum_i^n \lambda_i y_i^2. \qquad [9]$$

The equation $\boldsymbol{v}^T\boldsymbol{D}\boldsymbol{v} = 1$ describes an ellipsoid whose axes end at the points where $\lambda_i y_i^2 = 1$, and the remaining y components are zero (19). Undoing the rotation, these points are in the directions of the eigenvectors and the axes have half length $1/\sqrt{\lambda_i}$. It is important to emphasize, however, that the ellipsoid that describes the eigenspace of the diffusion tensor is not a description of the shape of the measured local diffusion D_{app} (15).

Multiple Fibers

The above definition of a "fiber" can be extended to multiple fibers by defining the k'th fiber as synonymous with the k'th diffusion tensor $\boldsymbol{D}_k$, with no assumption of cylindrical symmetry until it is necessary. In addition, the assumption will be made throughout that there is no exchange between fibers so that the signals add independently.

In general, the signal from a voxel containing N fibers can be written

$$S = S_0 \sum_{k=1}^{n} f_k e^{-b\hat{D}_k} \qquad [10]$$

where f_k is the volume fraction of the k'th fiber ($\sum_{k=1}^{n} f_k = 1$). It will prove useful to express this in the following form (Taylor expanding about $b = 0$):

$$\log(S/S_0) \approx -b\sum_{k=1}^{n} f_k\hat{D}_k - \frac{b^2}{2}\left\{\sum_{k=2}^{n} f_k(1-f_k)\Delta\hat{D}_{k1}^2 - 2\sum_{\substack{i=2\\ i\neq j}}^{n}\sum_{\substack{j=2\\ j\neq i}}^{n} f_i f_j \Delta\hat{D}_{i1}\Delta\hat{D}_{j1}\right\} \qquad [11]$$

where we have defined the differential apparent diffusion $\Delta\hat{D}_{kj} \equiv \hat{D}_k - \hat{D}_j$. This form of writing the signal is useful in that it expresses the effect of additional fibers $k \geq 2$ relative to the first ($k = 1$). For a single fiber, the terms in the bracket drop out and the log signal assumes the pure quadratic form utilized in Eq. [8] in the standard diffusion tensor approach. Moreover, it makes explicit the higher order (in b) dependence of the terms that couple different diffusion tensors through the differential apparent diffusion, revealing why high b values bring out the complex diffusion structure, whereas for low b values, the weighted mean D (the term linear in b) is dominant. Each $\hat{D}$ is a quadratic, so each term in Eq. [11] is an even power polynomial in gradient direction v. The approximation to orders $O(b^2)$ yields even polynomials up to order 4.

Special Case: Two Fibers

For multiple fibers within a voxel, Eq. [11] can be complicated. However, the special case of two fibers is instructive, and is a useful approximation in practice. This is the simplest multifiber case, and the resulting equations have an intuitively clear interpretation and are numerically relatively simple to manage. From Eq. [11], the log signal be written

$$\log(S/S_0) \approx -b\,[f_1\,\hat{D}_1 + f_2\,\hat{D}_2] + f_1\,f_2 b^2\Delta\hat{D}_{21}^2 \qquad [12]$$

where f_1 and f_2 are the volume fractions in compartments 1 and 2, respectively, so that $f_1 + f_2 = 1$ and $\Delta\hat{D}_{21} = \hat{D}_2 - \hat{D}_1$, and we have kept terms up to the second order in b. The measurements $\hat{D}_i$ are composed of second-order polynomials, so the coupling term $\xi \equiv f_1\,f_2\,b^2\,\Delta\hat{D}_{21}^2$ is composed of even order polynomials up to the fourth order.

The two terms linear in b represent the "individual" fiber components and are pure quadratic forms. In addition, there is a "coupling" term, second order in b, with coefficient f_1f_2 that is of the fourth order. Note the interesting fact that the coefficient term f_1f_2 is a quadratic function, with a maximum at $f_1 = f_2$. The magnitude of the coupling thus depends on the volume fractions; however, the shape does not. This is clear from the fact that the volume fraction enters only as a multiplicative factor in ξ. Variations in the shape of the coupling term are more easily understood by rewriting ξ in a more illuminating form. First note, from Eqs. [6] and [7] that

$$\Delta\hat{D}_{21} = \boldsymbol{v}^T\boldsymbol{R}_1^T\,\hat{\boldsymbol{D}}_{12}\boldsymbol{R}_1\boldsymbol{v} \qquad [13]$$

where we have defined

$$\hat{\boldsymbol{D}}_{12} \equiv \boldsymbol{R}_{12}^T\boldsymbol{D}_{2,\Lambda}\boldsymbol{R}_{12} - \boldsymbol{D}_{1,\Lambda} \qquad [14]$$

$$\boldsymbol{R}_{12} \equiv \boldsymbol{R}_2\boldsymbol{R}_1^{-1}. \qquad [15]$$

The term $\boldsymbol{R}_{12}$ is the product of rotation matrices and is therefore also a rotation matrix (by virtue of the group properties of the 3D rotation group, denoted as $O(3)$ (20)). It first undoes the rotation $\boldsymbol{R}_1$, then applies the rotation $\boldsymbol{R}_2$. If the rotations $\boldsymbol{R}_1$ and $\boldsymbol{R}_2$ are the same, meaning that the fibers are pointing in the same direction, then $\boldsymbol{R}_{12} = 1$, and $\tilde{\boldsymbol{D}}_{12}$ is a diagonal matrix with eigenvalues equal to the difference in principal diffusivities between $\boldsymbol{D}_1$ and $\boldsymbol{D}_2$. If the fibers are identical in orientation and diffusivities (by our definition, $\boldsymbol{D}_1 = \boldsymbol{D}_2$), then $\tilde{\boldsymbol{D}}_{12} = 0$, so that $\Delta\tilde{D}_{12} = 0$ and the coupling term disappears, reducing the log signal to the standard single-fiber form (Eq. [8]). In general, though, the term $\Delta\hat{D}_{21}$ is of the same form as the rotated single-fiber diffusion matrices (e.g. Eq. [7]) but in terms of the new tensor $\tilde{\boldsymbol{D}}_{12}$, which we term the reduced diffusion tensor.

The accuracy of the approximation in Eq. [12] depends upon both the b-value and the relative orientation of the fibers. As shown in Fig. 2, this approximation is good up to quite high b-value. The manifestation of these errors is most clearly visualized by plotting D_{app} as a function of b-value for $\Delta\theta = 90°$ (the angle at which the influence of the coupling term is the greatest), an example of which is shown in Fig. 3.

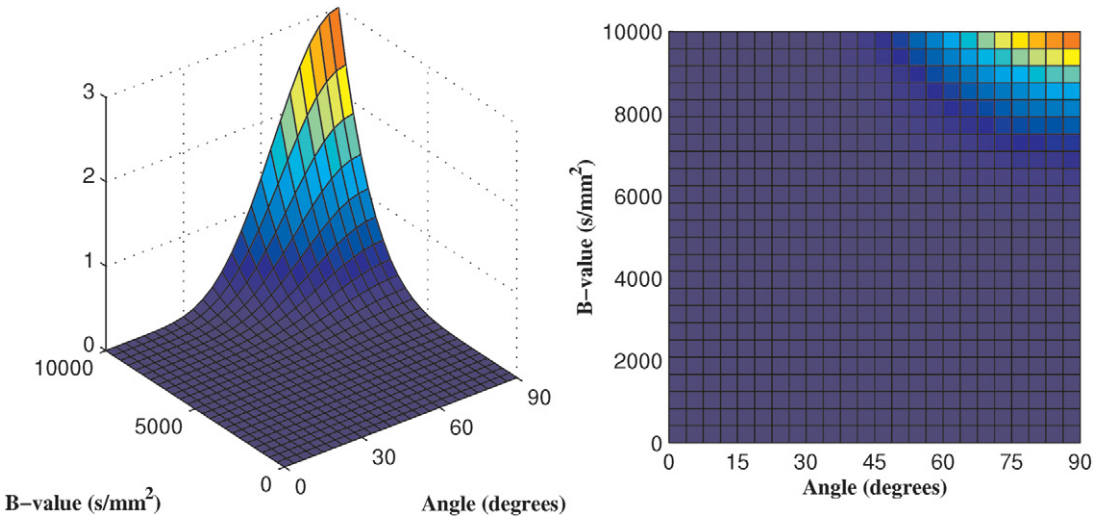

FIG. 2. Mean squared error between the true diffusion log (signal) (Eq. [11]) and the approximate form (Eq. [12]) as a function of *b*-values and azimuthal angle $\Delta\phi$ between two identical fibers oriented in the *x-y* plane (i.e., polar angle $\theta = 90°$). The approximation error gets worse for large *b* and large $\Delta\phi$, but is less than .05 for $b < 5000$ s/mm^2. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

The results for the HARD diffusion measurements can be summarized as follows:

1. The measured local diffusion is related to the diffusion in the fiber coordinate system by 3D rotations.
2. The measured log signal from a single fiber is a pure quadratic form.
3. The measured log signal from multiple fibers can be approximated by the sum of two pure quadratic forms and a coupling term of even polynomial up to order 4.

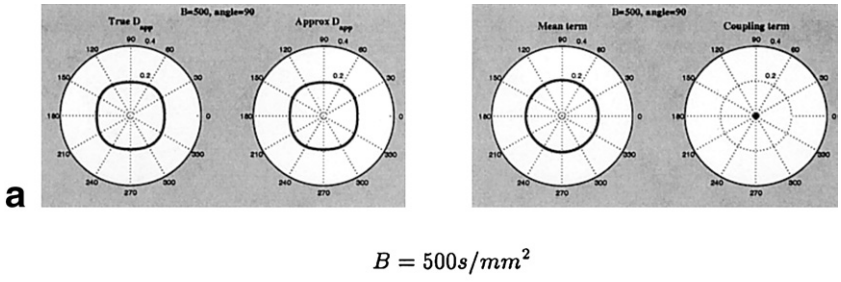

$B = 500 s/mm^2$

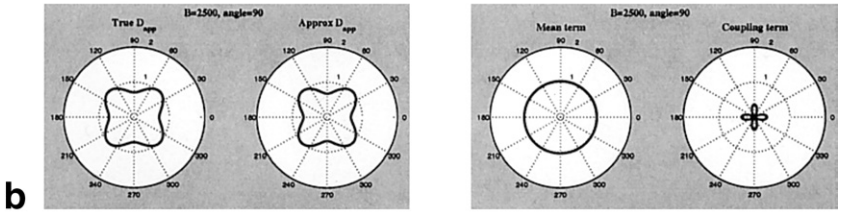

$B = 2500 s/mm^2$

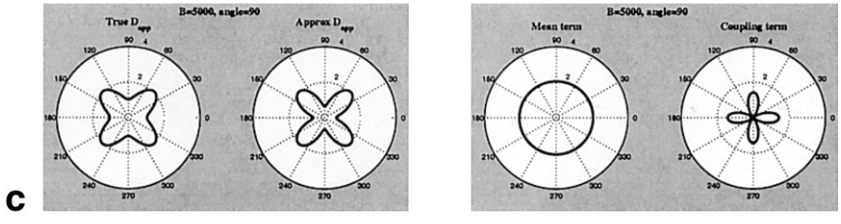

$B = 5000 s/mm^2$

FIG. 3. The effect of the coupling term in the approximation (Eq. [12]) in a voxel containing two fibers oriented in the *x-y* plane at an angle $\Delta\phi = 90°$ relative to one another for three *b*-values: $b = 500, 2500, 5000$ s/mm^2. (Left) The true D_{app} and the approximate D_{app} determined from Eq. [12]. (Right) The mean and coupling terms in Eq. [12]. The mean term is always circular. For increasing *b*-values the approximation is less accurate, as demonstrated by the increasing discrepancy between the true and the estimated shapes in Fig. 2.

The question then is this: If an eigenvector decomposition is sufficient to determine the diffusion tensor in the single-fiber case, is there a decomposition sufficient to determine the diffusion in the case of multiple fibers? We show in the next section that this is indeed possible.

SPHERICAL HARMONIC DECOMPOSITION

Spherical Tensor Representation of Rotations

As described in the previous section, the diffusion measurements along different encoding directions can be expressed as rotations in three dimensions relative to the (unknown) principal axis system of the diffusion. In the HARD technique, these measurements are along a set of directions covering the range of the spherical coordinates (θ, ϕ). The rotations $\boldsymbol{R}$ were expressed using a form, or representation, in terms of the Euler angles. However, the problem was formulated in the familiar Cartesian coordinate system, or basis: $\boldsymbol{e} = \{x, y, z\}$, because this is the natural coordinate system for imaging. Systems with spherical symmetry are often more conveniently handled in the spherical basis: $\boldsymbol{e} = \{r, \theta, \phi\}$, as defined in Fig. 1. As we show next, rotation matrices transformed into the spherical basis are the spherical harmonics (21). Tensors transformed into this representation are called spherical tensors (21).

The HARD diffusion measurements have an inherent spherical symmetry because they are made by a series of 3D rotations. The inherent symmetry in this problem can be elucidated through the theory of groups, which was originally developed for the purpose of characterizing symmetry (21). The concept and consequences of groups, although extremely powerful, are conceptually relatively simple. Their great power is that they facilitate the characterization and classification of mathematical structures into classes or groups with similar symmetry properties. All members within a particular group can then be treated as equivalent, even if their specific manifestations differ. For example, two points constrained to move on the surface of the same sphere can be seen as having identical symmetry properties even if their precise locations on that sphere differ.

The importance of the spherical tensor formulation is encapsulated in the following expression, which shows how a spherical tensor $\boldsymbol{T}$ is affected by a general rotation $\mathfrak{D}^{1}_{\mathrm{m'm}}(\Re)$ in some representation $\Re$ of the rotation group in the basis $\boldsymbol{e}$:

$$T_{lm} = \sum_{m'=-l}^{l} \mathfrak{D}^{l}_{m'm}(\Re)\, T_{lm'}. \qquad [16]$$

Notice that rotations of the individual m components within a particular component of rank l do not "mix" elements amongst components of different rank l. This quality of the $\boldsymbol{T}$ is called irreducibility. Equation [16] can be considered the defining quality of a spherical tensor: a tensor that transforms accordingly is by definition a spherical tensor. The spherical tensor representation is useful because rotations of the individual tensor components preserves the rank.

Both the spherical tensor and the rotations can be related to our specific problem as follows. A single fiber tensor is a second-rank tensor and so consists of nine components, represented by a 3×3 matrix. In its irreducible representation, the tensor is written as the sum of three terms:

$$\boldsymbol{T} = \boldsymbol{T}^{0} + \boldsymbol{T}^{a} + \boldsymbol{T}^{s} \qquad [17]$$

where $\boldsymbol{T}^{0} = T^{0}\boldsymbol{I}$ in which T^{0} is a rank 0 tensor (i.e., a scalar) and $\boldsymbol{I}$ is the 3×3 identity matrix; $\boldsymbol{T}^{a}$ is an antisymmetric rank 1 tensor (i.e., a vector); and $\boldsymbol{T}^{s}$ is a symmetric, rank 2 tensor. For a general spherical tensor, the rotations $\mathfrak{D}^{l}_{mm}$ (α, β, γ) in the spherical tensor basis expressed in terms of the Euler angles are called the *Wigner rotation matrices* (21). For a point in spherical coordinates, the Wigner rotation matrices are proportional to the spherical harmonics (22):

$$\mathfrak{D}^{l*}_{m0}(\alpha, \beta, \gamma) = \left(\frac{2l+1}{4\pi}\right)^{-1/2} Y_{lm}(\beta, \alpha). \qquad [18]$$

Therefore, the process of rotating a diffusion tensor can be reformulated by expressing the diffusion tensor in an irreducible form in which its individual components transform separately under rotations affected by spherical harmonic components. In this general formulation, the concept of the diffusion tensor can easily be extended to more complex structures by considering tensors of higher rank because now the transformation under rotations are of exactly the same form (Eq. [16]) and the functions that perform the rotations are exactly the same basis set (the spherical harmonics).

Application to HARD Measurements

Now, consider the general case of a HARD measurement of a voxel of unknown fiber composition. The measured apparent diffusion coefficient $D(\Omega)$ is then an arbitrary real function. The complex spherical harmonics form a complete orthonormal basis (18) so an arbitrary real function parameterized by the spherical coordinates (θ, ϕ) can be expanded in a Laplace series:

$$\boldsymbol{D}(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-1}^{l} a_{lm} Y_{l}^{m}(\theta, \phi). \qquad [19]$$

The coefficients a_{lm} are determined by (18) multiplying both sides of Eq. [19] by $Y_{l'}^{m'*}(\theta,\phi)$ and using the orthogonality condition

$$\int_{0}^{2\pi} \int_{0}^{\pi} Y_{l}^{m}(\theta,\phi)\, Y_{k}^{n}(\theta,\phi)\, \sin(\theta)\, d\theta d\phi = \delta_{lk}\, \delta_{mn}. \qquad [20]$$

The expansion coefficients are uniquely determined by multiplying each side of Eq. [19] by $Y_{k}^{n}(\theta, \phi)$ and integrating over the sphere. The result is that the coefficient can be determined by

$$a_{lm} = \int_0^{2\pi} \int_0^{\pi} \boldsymbol{D}_{app}(\theta, \phi) Y_l^m(\theta, \phi) \sin(\theta)\, d\theta d\phi. \quad [21]$$

This is precisely analogous to a Fourier decomposition of sinusoidal functions, but on the unit sphere. This will be called the spherical harmonic transform (SHT) of the measured apparent diffusion coefficient.

The utility of describing the measured HARD (log) signal in terms of group theoretical constructs can now be shown directly as follows. The measured HARD signal $\hat{D}(\Omega)$ is an arbitrary real function and so can be expanded in terms of a Laplace series (Eq. [19]) with the coefficients determined by the SHT (Eq. [21]). But the symmetry of HARD diffusion measurements imposes severe restrictions on the expansion coefficients that allow direct classification of the diffusion from the SHT. In particular, the SHT of $\hat{D}(\Omega)$ produces the coefficients of the irreducible representation of $\hat{D}(\Omega)$.

Specifically, the following is true: Isotropic diffusion is independent of direction, so the lowest order $Y_{00}(\Omega)$ is the basis for the representation of $\mathfrak{D}^{(0)}$. This is easily seen because $Y_{00}(\Omega)$ is just a sphere, so the calculated coefficient merely scales the radius of the sphere. This provides a $2L + 1 = 1$-dim representation of the $O(3)$. For a single fiber, the irreducible representation of $\hat{D}(\Omega)$ (Eq. [17]) provides a basis for a 6-dim representation of $O(3)$ (with T^o providing $2 * 0 + 1 = 1$, and T^s providing $2 * 2 + 1 = 5$). The basis functions are then the spherical harmonics $Y_{lm}(\Omega)$ of order $L = 0, 2$. Because this is an irreducible representation, Eq. [16] expresses the fact that fiber rotations do not alter the basis functions. That is, they only produce a redistribution of energy amongst the M components or variations in the phase of the M components, but not the order L. For multiple fibers, the log signal from multiple fibers can be expressed in terms of an expansion in even-power polynomials (Eq. [11]) and approximated up to relatively high b-values by keeping terms up to the second order in b, which corresponds to terms of 4'th-order polynomials. A multifiber voxel, therefore, provides a basis for a $2 * (0 + 2 + 4) + 3 = 15$-dim representation of $O(3)$, with basis functions being the spherical harmonics of even orders up to order $L = 4$. The dimensions of the representation are the number of measurements that are required to characterize the $\hat{D}(\Omega)$. It is for this reason that $2 * (0 + 2) + 2 = 6$ measurements are required to characterize the standard single-fiber diffusion tensor.

The results can be summarized by the following rather remarkable conclusions:

1. Isotropic diffusion is described by $a_{00} Y_0^0(\Omega)$.
2. Single-fiber diffusion is described by $\sum_{l=0,2} a_{lm} Y_l^m(\Omega)$.
3. Multiple-fiber diffusion is approximately described by $\sum_{l=0,2,4} \sum_m a_{lm} Y_l^m(\Omega)$.
4. In general, the local diffusion, including the magnitude and orientation, can be described by sum of spherical harmonics of even order, i.e., $\sum_l \sum_m a_{lm} Y_l^m(\Omega)$, $l = 0,2,4, \ldots$.
5. Odd orders of the spherical harmonics describe asymmetric components and therefore imaging artifacts.
6. The coefficients a_{lm} are determined by the spherical harmonic transform of $\hat{D}(\Omega)$.
7. The dimension of the representation is the number of measurements required to characterize the apparent diffusion coefficient $\hat{D}(\Omega)$.

It is important to point out that the order L required to characterize the diffusion in a multifiber voxel depends upon the orientation of the two fibers. For fibers more closely aligned, higher orders of L will be required. This can be seen by considering the simplest case of two identical fibers lying in the $\theta = 90°$ plane, oriented nearly parallel to one another. Distinguishing between the two fibers requires high resolution in the azimuthal angle ϕ. Since the azimuthal dependence of the spherical harmonics is proportional to $\exp(im\phi)$, higher frequency variations in ϕ require larger values of m and thus higher orders of L in the basis functions. By 7 above, more dense sampling is then required. We have focused on the simplest case in the present work where $L = 4$ is considered sufficient, but this is not required.

The group theory arguments provide a simple and concise description of the categorization of voxel diffusion characteristics, since the above results are expressible as direct sum subspaces. Let us call the "state" of a voxel with k fibers ψ_k, where $k = 0$ means isotropic diffusion. Then we can write

1. Isotropic diffusion: $\psi_0 = \mathfrak{D}^{(0)}$.
2. Single-fiber diffusion: $\psi_1 = \mathfrak{D}^{(0)} \oplus \mathfrak{D}^{(2)}$.
3. Multiple-fiber diffusion; $\psi_2 \approx \mathfrak{D}^{(0)} \oplus \mathfrak{D}^{(2)} \oplus \mathfrak{D}^{(4)}$

where the $\mathfrak{D}^{(j)}$ are the irreducible representations of the rotation group, and $\oplus$ denotes the direct sum subspace. In principle, the composition of a voxel in terms of these can be determined, as well as the magnitude and orientation of the local diffusion.

The symmetry inherent in this problem precludes the existence of power in the odd L channels. Energy in these channels in the SHD of actual experimental data is therefore produced by non diffusion effects, such as subject motion or eddy currents, which are not constrained to the same symmetry properties. This can be used to advantage as a means of identifying nondiffusion-related variations. Of course, such variations may also have power in the diffusion-related channels, so their reconstruction may not be trivial. The information in the odd channel might also be useful in incorporating eddy current correction into the image reconstruction.

The ability to characterize the diffusion does not imply that extraction of fiber information is easy, however. This becomes apparent in examining the coupling term ξ, about which can be said:

1. The shape of ξ depends on the eigenvalues of the reduced diffusion tensor, and thus on the relative anisotropies of two fibers.
2. The orientation (and hence phase) of ξ depends upon the mean orientation of fibers.
3. The magnitude of ξ depends on the volume fractions and the relative orientations of the fibers.

The last item underscores a basic ambiguity in the diffusion measurements: the volume fractions and relative orientations can confound information about one another in the measured signal. However, the orientations affect the phase, whereas volume fraction changes do not.

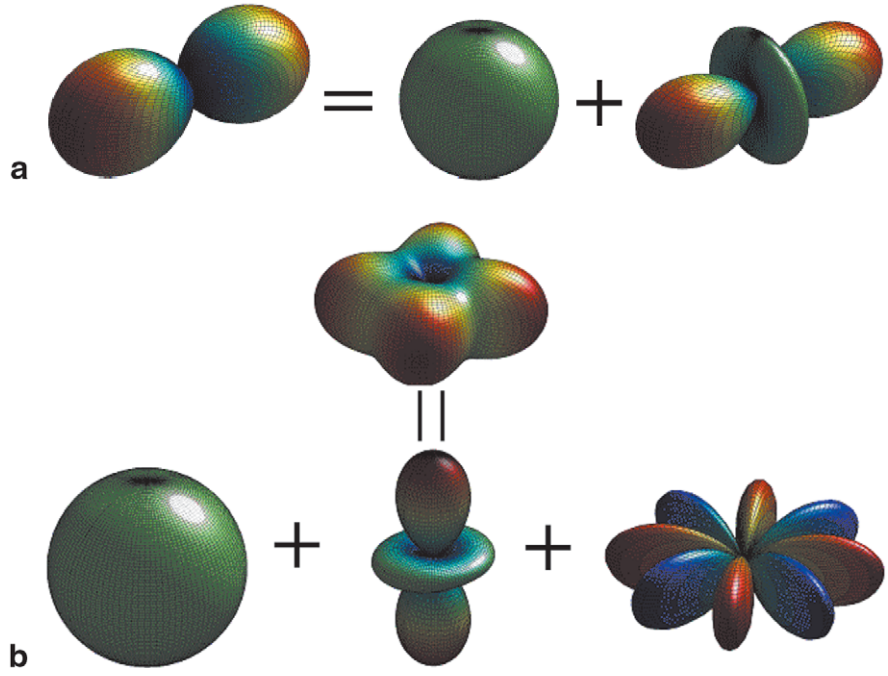

FIG. 4. Graphical representation of the SHD. The contributions to D_{app} from the different L orders of the SHD are shown for (**a**) a single-fiber and (**b**) a double-fiber voxel. **a**: The D_{app} contains contributions only from the $L = 0$ term (the sphere) and a more complex shape produced by the $L = 2$ spherical harmonic components. **b**: The D_{app} for two identical fibers oriented $(\Delta\phi, \Delta\theta) = (90°, 0°)$ relative to one another contains contributions only from the $L = 0$ and $L = 2$ spherical harmonic components, as in **a**, as well as a contribution from the $L = 4$ component (bottom right). Note that the four-lobe structure of D_{app} in the equatorial plane is generated by the addition of the four positive red lobes and the four negative blue lobes of the $L = 4$ component to the $L = 0$ component sphere. The color scheme has red at maximum positive, and the blue at maximum negative. Each shape is scaled independently for display.

Simulation and Computation of the SHD

Display of the SHD

Before proceeding, it is important to outline our basic method of displaying the SHD data. Because there is a great deal of information produced by this method, we have found it important to have a concise method for displaying the results. The SHD calculates components up to a user-specified value of L, for all the $2L + 1$ values of M associated with each L (because $M = -L, \ldots 0, \ldots + L$). A useful way to display the coefficients is thus on an array of coordinates (L, M), as shown in Fig. 5. Since $M \leq L$, the diagram has a characteristic triangular shape (white boxes have no coefficients).

This theory is easily confirmed by simulation, as shown in Fig. 4 where the D_{app} for a single fiber and for two identical perpendicular fibers are broken down in terms of the separate contributions from the different spherical harmonic components. In this simulation we use a fiber model that possesses cylindrical symmetry, since this is a reasonable physical model for white matter geometry (e.g., Ref. 17), though the above results do not require this. The corresponding SHD is shown in Fig. 6, which shows on the right the SHT of the D_{app} from a single fiber (top) rotated through the entire range of both θ and ϕ, and two fibers (bottom) with one fixed along the x-axis and the relative angle between the fibers rotated through the entire range of both θ and ϕ. The D_{app} for one particular orientation is shown on the left. No energy is produced in channels other than those predicted by the group arguments above.

Generally, the coefficients are determined for a range of parameters, in which case the boxes become arrays wherein the parameters are varied. The first is where the parameters are the spatial coordinates (x, y). Then each "box" on the (L, M) plot is just an image of the spatial distribution of the amplitude of that specific (L, M) component of the SHD. The other case of importance is when the parameters are the azimuthal and polar fiber angles (ϕ, θ) as shown in the example in Fig. 6. This is useful, for example, in showing that as a single fiber is rotated arbitrarily, there is a redistribution of amplitudes within the ψ_1 subspace, but the energy remains completely contained with this channel, as predicted.

Numerical Methods

The spherical harmonic decomposition is achieved by computing the spherical harmonic transform of the measured (i.e., apparent) diffusion coefficient (Eq. [21]). Unfortunately, unlike the discrete Fourier transform (DFT), for which there exists a matrix decomposition that allows the fast Fourier transform (FFT), no such algorithm exists for the SHT. Although a variety of algorithms have been proposed for the computation of spherical harmonic coefficients (23,24), no clear algorithm has emerged as clearly superior, and the subject remains an area of active research. Therefore, for the present work the coefficients were determined by direct computation of Eq. [21].

The direct computation of Eq. [21] to order $L_{\max}$ on a grid of N values of ϕ and M values of θ requires the computation of $\sum_{l=0}^{L_{max}} (2l + 1)$ spherical harmonics evaluated at $N\,M$ points. The spherical harmonics involve multiple expensive trigonometric evaluations. Many of the trigonometric evaluations are redundant, however, so precomputation of these values can be used to speed up the computation. Direct computation of the measure $\sin(\theta) d\theta d\phi$ requires M trigonometric evaluations and $N\,M^2$ multiplies. However, an efficient algorithm for computing the measures was developed by noting that these weights in the summation that approximates the integral are equal to the Voronoi areas for the sampling points on the unit sphere.[1] Precomputation of the trigonometric functions and

[1]The code for computing spherical voronoi areas was written by R. W. Cox.

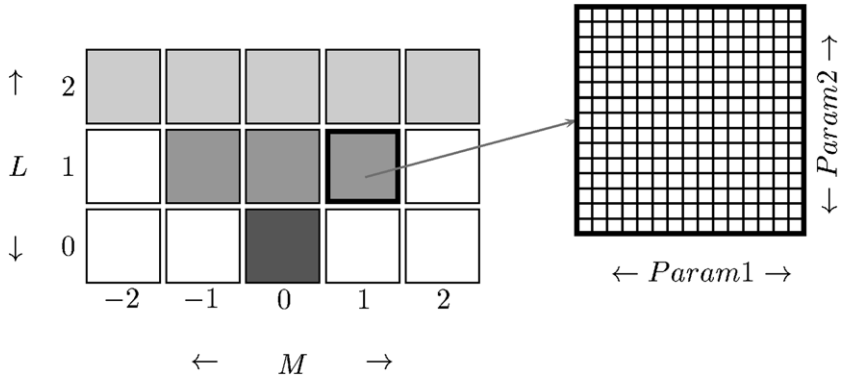

FIG. 5. Graphical representation of the spherical harmonic transform. For each value of L there corresponds $2L + 1$ values of M : $-L, \ldots, 0, \ldots, L$. The box at each (L, M) coordinate contains an array (upper right diagram) whose coordinates are varied parameters at the fixed values (L, M). The transform is identically zero for $|M| > L$ (blank boxes).

the weights therefore allowed an efficient SHT. This algorithm has been incorporated into the author's diffusion plug-in module in AFNI (25) and is sufficiently fast to rapidly process reasonably large HARD data sets. For example, the data shown in the present work are comprised of 10 slices and 43 directions. The SHT on the entire data set took only 1.25 min on an SGI Octane2 with dual R12000 processors.

STRUCTURE OF THE SHD OF DIFFUSION MEASUREMENTS

Determining Significance of the Multiple-Fiber Channel

The categorization of voxel fiber composition outlined above suggests a strategy for the analysis of HARD data. The SHT is taken for each voxel in the image and sorted into even- and odd-order L. The odd orders represent artifacts and therefore can be eliminated from the analysis. The remaining even orders up to order $L = 4$ are then sorted in the following manner. Voxels with significant power in $L = 4$ are classified as "multiple fibers," voxels with significant power in $L = 2$ but not $L = 4$ are classified as "single fibers," and voxels with significant power in $L = 0$ but not $L = 2$ or $L = 4$ are classified as "isotropic." The proper method for determining significance is to determine which fiber model is most probable in any particular voxel. This approach will be pursued in a future work. Having determined into which classification each voxel falls, the voxel local diffusion is determined from the sum of the appropriate spherical harmonics components (isotropic: $L = 0$; single fiber: $L = 0,2$; multiple fiber: $L = 0, 2, 4$), with the coefficients determined from the spherical harmonic transform.

In lieu of a more complete analysis, a simple method for the determination of significance is suggested by the results above that show that the magnitude of the $L = 4$ term increases with increasing relative orientation of the fibers. Therefore, a comparison of energy in the $L = 4$ to the $L = 2$ channel can be used to gauge of whether or not a voxel is of state ψ_0 or ψ_1. One measure of the significance of a multiple-fiber channel is the fractional even order greater than 0 in that channel. We can define the fractional multifiber index (FMI) as

$$FMI \equiv \frac{\sum_{L\geq 4} \sum_M |A_{L,M}|^2}{\sum_M |A_{L=2,M}|^2}, \quad \text{L even.} \qquad [22]$$

An example is shown in Fig. 10 (bottom) for the simplest example of two fibers at a range of relative orientations ranging from parallel to perpendicular. This is a reasonable measure of comparison and means of separating single- and multifiber channels, raising the question of which threshold to choose. We will use this here in lieu of a more complete probabilistic model, which will be pursued in the future.

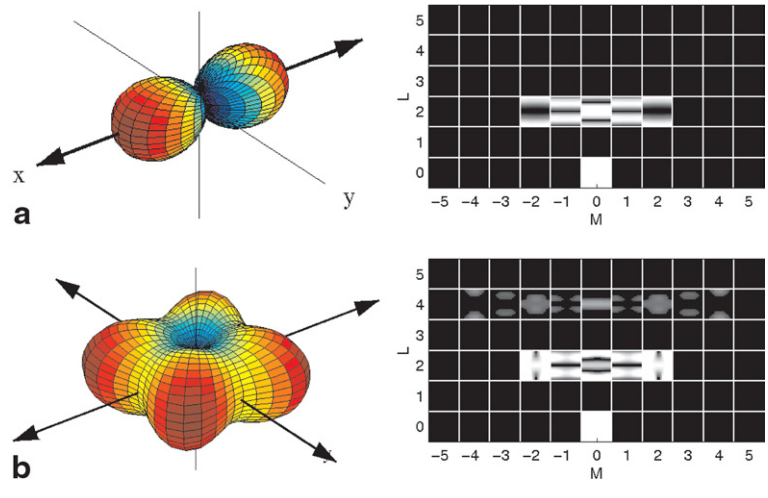

FIG. 6. Rotational variations of the SHD. Reconstructed D_{app}'s are on the left. The SHT coefficients are on the right, with the varied parameters azimuthal ($\Delta\phi$) angle on the horizontal axis and the polar ($\Delta\theta$) angle on the vertical axis, and displayed as described in Fig. 5. **a**: A single fiber rotated through the full range of (ϕ, θ). **b**: Two identical fibers oriented ($\Delta\phi$, $\Delta\theta$) relative to one another: (left) D_{app} for ($\Delta\phi$, $\Delta\theta$) = (90°, 0°), (right) spherical harmonic transform of two fibers with one fiber rotated through the full range of ($\Delta\phi$, $\Delta\theta$). The angular variations are with respect to the first fiber, which is fixed along $x = 0$. The SHT in **a** and **b** shows that only $L = 0, 2$ components arise from the single fiber and only $L = 0, 2, 4$ arise from the multiple fiber. This is a consequence of the fact that rotations only mix M components.

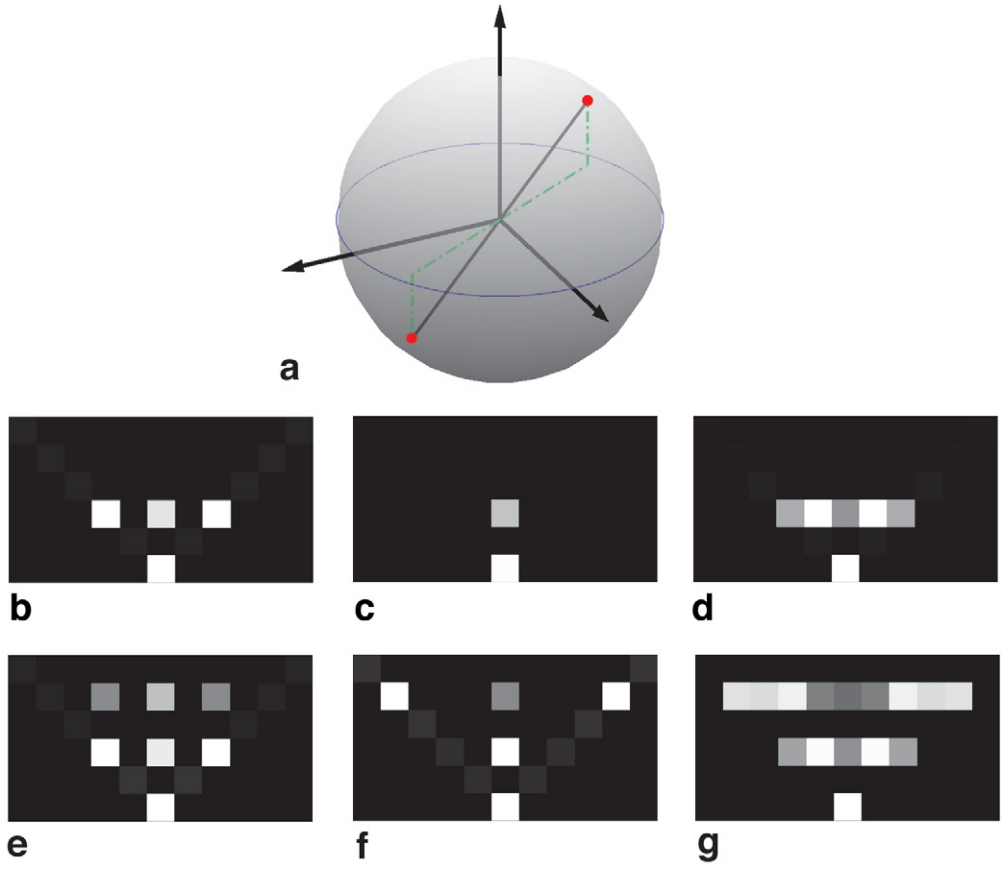

FIG. 7. MR diffusion measurements are sensitive only to the absolute value of the direction: motion in $+x$ has the same effect as motion in $-x$. This imposes a projective symmetry on the measurements so that antipodes on the measurement sphere (**a**) are identical. **b–d:** Projective subspace of the single fiber with $\phi = 0$, $\theta =$ **b**) 90°, **c**) 0°, **d**) 45°. **e–g:** Projective subspace behavior of two fibers with $(\phi_1, \theta_1) = (90°, 0°)$ and $\phi_2 = 90°$ with $\theta_2 =$ **e**) 0°, **f**) 90°, **g**) 45°. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Symmetry of Diffusion Weighting: Projective Subspaces

Decomposition of the diffusion into separate isotropic, single, and multifiber channels is a consequence of the group algebra, which generates additive subspaces that depend upon the degree of polynomial necessary to describe the measurements. However, this is far from the complete story, for it is just the L story. There is structure within the SHD contained in the way in which the energy is distributed among the M components.

There is a fundamental symmetry in the D_{app} imposed by the imaging process because the signal loss due to diffusion along the direction of a gradient is insensitive to the sign of the motion. That is, equal diffusive motion in both the $+x$ and $-x$ directions produce the same diffusion-related signal loss. This results in a projective symmetry, which can be visualized by the diagram in Fig. 7a. In the mathematics literature, *projective* is synonymous with *antipodal*, and is unrelated with projection in the sense commonly used in the physics literature (i.e., the component along a chosen axis). Diffusion measurements are therefore represented by the *projective* subgroup of $O(3)$ (denoted $PSO(3) = O(3)/Z(2)$ where "/" can be thought of as "mod" and $Z(2)$ represents the group of two integers). This means that two antipodal points (the two integers) on a sphere are indistinguishable. The effect of projective symmetry is that it restricts the energy to even values of M because odd values are not symmetric in ϕ. However, this symmetry imposed on the $PSO(2)$ subgroup of $O(3)$ is only evident for a single fiber in the equatorial plane ($\theta = 90°$) because arbitrary rotations can be represented by mixtures of (ϕ, θ) components. These effects are demonstrated by simulation in Fig. 7b–d.

The SHD for two fibers exhibits a similar symmetry. For two fibers in the equatorial plane oriented at $\Delta\phi = 90°$ to one another, the mean component is cylindrically symmetric and therefore does not possess $(L, M) = (2, \pm 2)$ components. For two fibers in the meridian plane oriented at $\Delta\theta = 90°$ to one another, the coupling term generates all even M terms for $L = 0, 2, 4$ except for $(L, M) = (L, \pm 4)$. Two fibers in the equatorial plane generate only one set of non-zero M components: $(L, M) = (4, \pm 4)$ components, similar to the single fiber, and by symmetry do not possess $(L, M) = (2, \pm 2)$ components. Again, arbitrary rotations produce mixing into all available M components. The projective subspace behavior of two identical crossed fibers is shown in Fig. 7e–g.

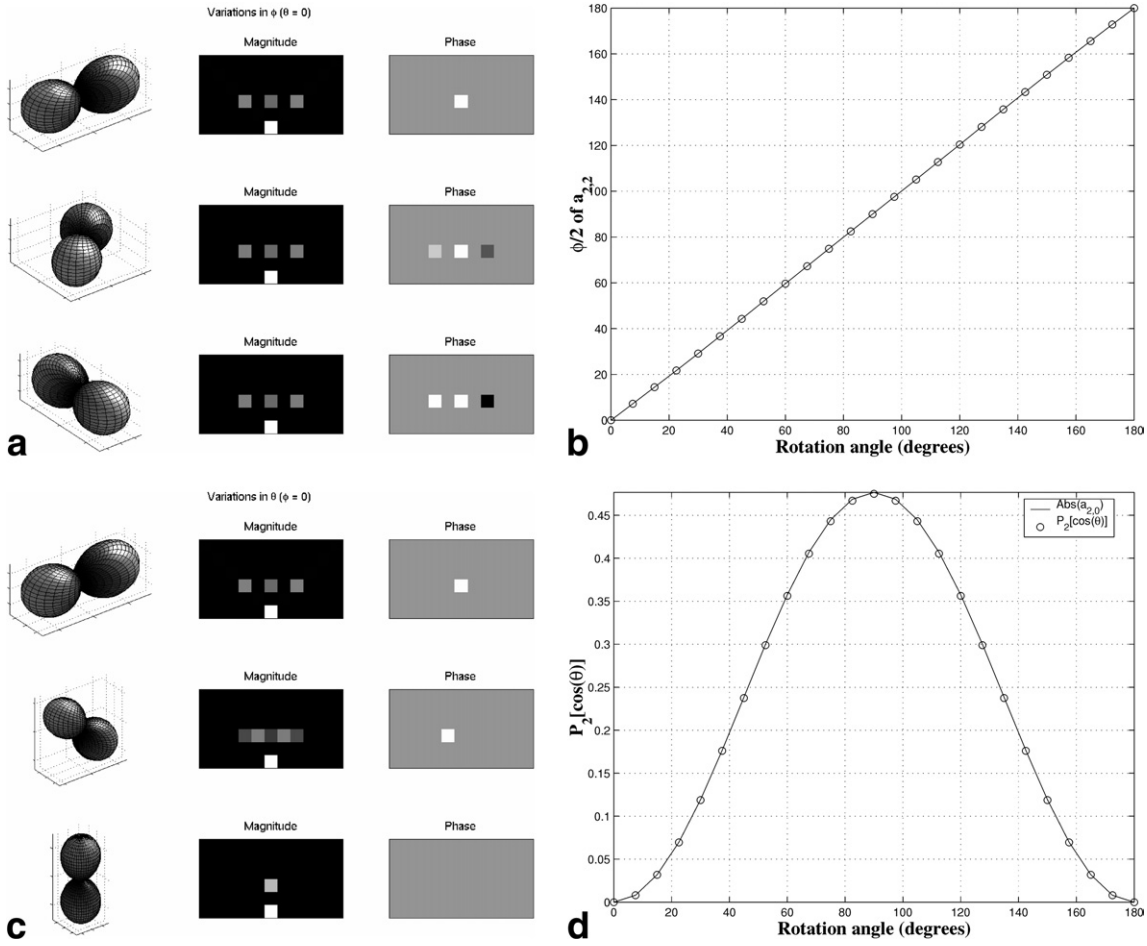

FIG. 8. Single-fiber rotations. **a** and **b:** Single-fiber azimuthal rotation. Rotations in the equatorial plane produce only phase changes. **c** and **d:** Single-fiber polar rotation. Rotations in the meridian plane produce only magnitude changes.

Fiber Orientation

Characterizing the local diffusion amounts to determining diffusion tensor(s) and their orientation(s) relative to the laboratory system. In the special case of cylindrically symmetric diffusion, a natural definition of the fiber orientation is the direction of the principal (i.e., largest) eigenvector e_1 of the diffusion tensor (8). For the problem of multiple-fiber voxels, the issue of orientation becomes more complicated, for one can examine the orientation of the individual fibers as well as their orientation relative to one another. In this section we reconsider the case of the single fiber, but now in the context of the SHD. We then consider the simplest multifiber case of two cylindrically symmetric fibers at arbitrary orientations.

Single-Fiber Orientation From the SHD

Because the SHD results for a single fiber identically reproduce those of the diffusion tensor, a single-fiber orientation can always be determined from the SHD by transforming from the spherical basis back to the Cartesian basis, thereby reconstructing the diffusion tensor, and then determining the fiber orientation from the principal eigenvector.

The key to determining the fiber orientation from the SHD is to recognize that a rotation R_ϕ of the fiber by an azimuthal angle ϕ modulates that phase of the SHT components because the ϕ dependence of the coefficients is of the form $\exp(i\phi)$. On the other hand, fiber rotations R_θ by a polar angle produce amplitude variations because the θ dependence is proportional to $P_l(\cos\theta)$, the Legendre polynomial, which is a polynomial in $\cos\theta$ of order L. As a consequence of Eq. [16], the θ variations mix energy among the available M components for a particular L component, but does not exchange energy among the L components.

These variations are illustrated in Fig. 8(a and b), where a single cylindrically symmetric anisotropic fiber is rotated through azimuthal angles $\phi = \{0°, 45°, 90°\}$ while fixed at the equatorial plane (i.e., $\theta = 90°$) and through polar angles $\theta = \{0°, 45°, 90°\}$ while fixed at the prime meridian (i.e. $\phi = 0°$). The phase and magnitude for the single-fiber case shown in Fig. 8a–d demonstrate that the ϕ orientation can be determined from the phase, while the θ orientation can be determined from the amplitude.

In this example, the phase (Fig. 8a and b) is determined from the coefficient $a_{2,2}$ that corresponds to the spherical harmonic $Y_2^2 = \sqrt{15/32\pi}\sin^2(\theta)e^{i2\phi}$ so that the estimate of the azimuthal orientation is $\hat{\phi} = \arg[a_{2,2}]/2$, where arg

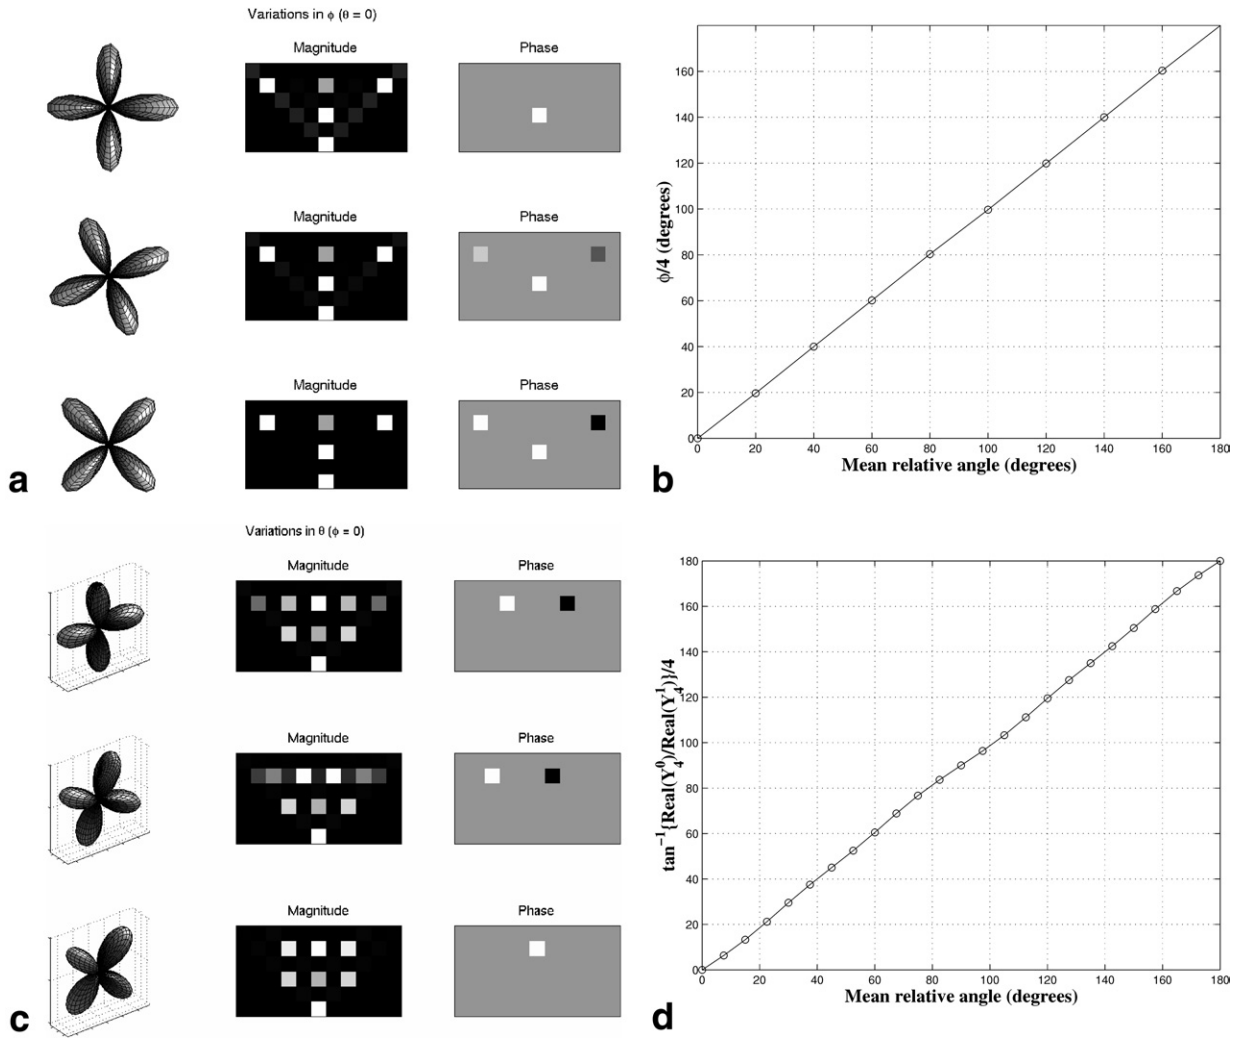

FIG. 9. Double-fiber rotations. **a** and **b:** Double-fiber azimuthal rotation. Rotations in the equatorial plane produce only phase changes in the coupling term. **c** and **d:** Double-fiber polar rotation. Rotations in the meridian plane produce only magnitude changes in the coupling term and a mixing of energy amongst the *M* components.

denotes the phase angle. The magnitude variation of $a_{2,0}$ as a function of $0 \leq \theta \leq \pi$ is seen (Fig. 8c and d) to follow $P_l(\cos\theta)$ so that the estimate of the polar orientation $\hat{\theta}$ can be made by noting that the $a_{2,1}$ component has no energy for $\theta = 90°$, so that the ratio of $a_{2,1}$ to $a_{2,0}$ can be related to the polar angle by: $\hat{\theta} = \tan^{-1}(a_{2,0}/a_{2,1})$.

Two-Fiber Orientation From the SHD

Determination of the fiber orientations in a multifiber voxel is complicated by the fact that the measured signal is not just a function of the individual fiber orientations, but of their relative orientations and their volume fractions. However, the coupling term ξ that generates the $L = 4$ components have some remarkable properties that make this problem tractable, at least for the two fiber case. Three limiting cases illustrate these properties.

The first is two fibers coplanar in the equatorial plane (i.e., $\theta = 90°$). If the angle between the two fibers is kept constant but the two fibers are rotated relative to the laboratory coordinate system, then the coupling term ξ rotates, causing a phase change in the $L = 4$ components, but no magnitude change. This is shown in Fig. 9a and b.

The second is two fibers coplanar in the meridian plane (i.e., $\theta = 0°$). If the angle between the two fibers is kept constant but the two fibers are rotated relative to the laboratory coordinate system, then the coupling term ξ rotates, causing a phase change in the $L = 4$ components, but no magnitude change. This is shown in Fig. 9c and d.

The third is two fibers coplanar in the meridian plane (i.e., $\phi = 0°$). If the angle between the two fibers is varied but the mean angle between the fibers is kept fixed, the size of the coupling term ξ changes, disappearing when the fibers are aligned (since this is identical to a single fiber). This is shown in Fig. 10a and b (top).

The orientation results for the special case of two fibers can be summarized as follows:

1. Azimuthal rotations R_ϕ of two fibers together (i.e., fixed relative orientation but variable mean ϕ orien-

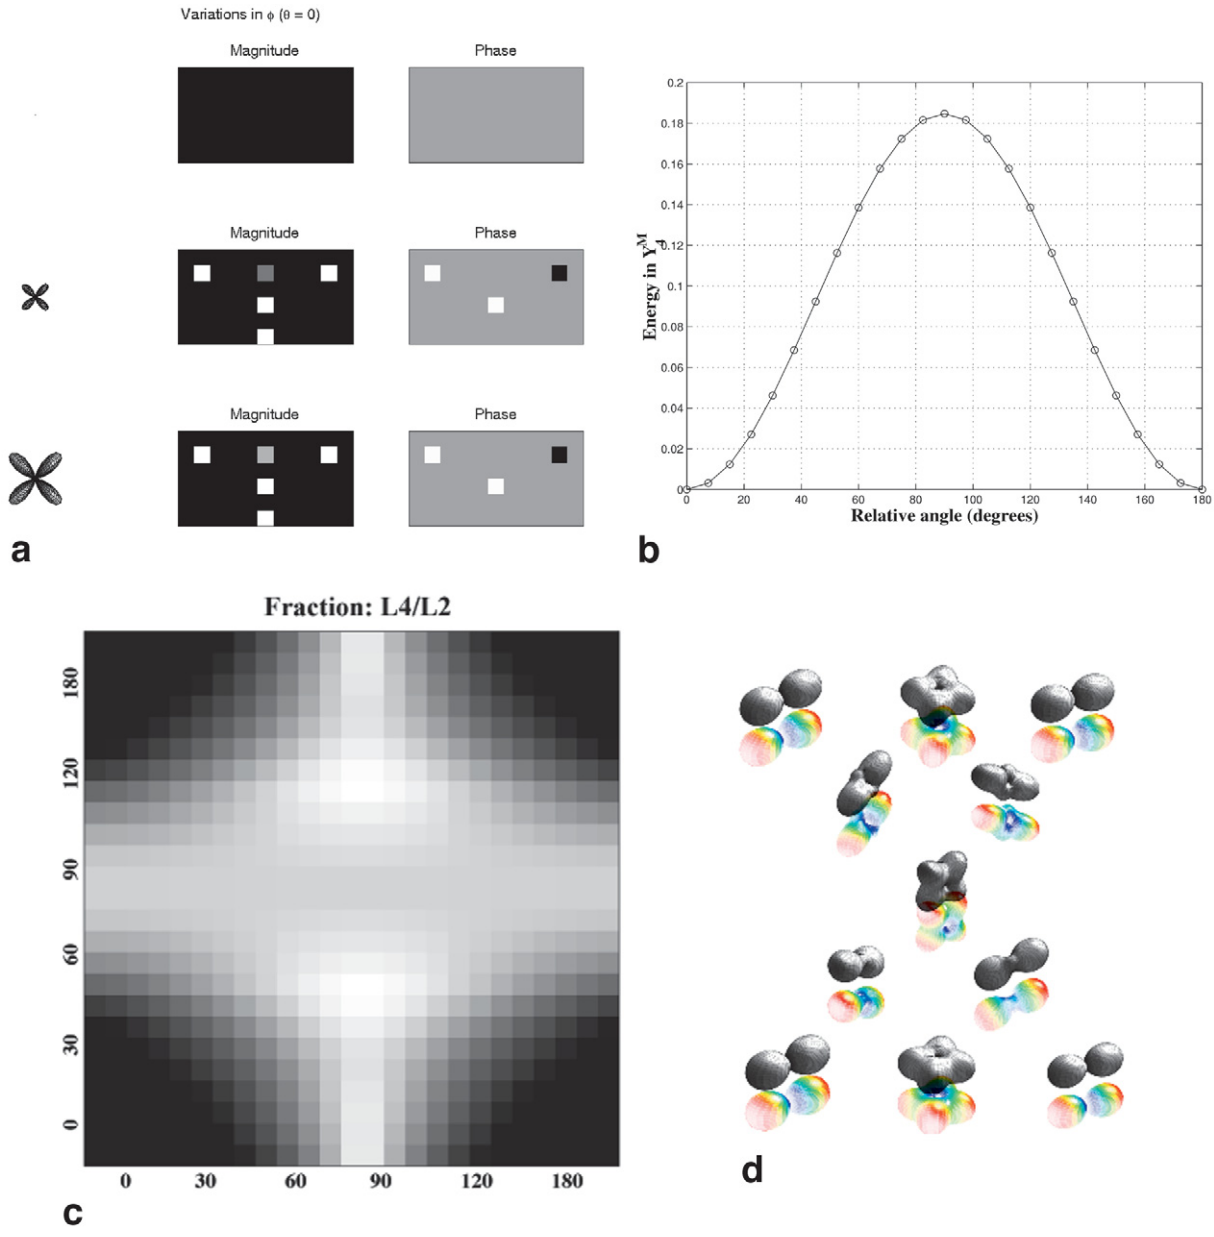

FIG. 10. Double-fiber relative rotations and the FMI. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

tation relative to laboratory frame) produces only phase change in the components.

2. Polar rotations R_θ of two fibers together (i.e., fixed relative orientation but variable mean θ orientation relative to laboratory frame) produces a redistribution of magnitude change in components, but no phase changes.
3. Variable relative orientation but fixed mean orientation relative to the laboratory frame produces only magnitude changes in the coupling component, but the relative amplitudes of the components (i.e., the pattern) remains unchanged.

Utilizing these results, however, necessitates the assumption of a two-fiber model.

METHODS

To test the application of this method on actual data, images were acquired on a GE SIGNA 1.5T Clinical Imager with high-speed gradient hardware using a previously described stimulated-echo spiral acquisition (15). Diffusion-sensitive images were acquired on five normal human subjects, with approval from the Human Subjects Committee of UC–San Diego and the VA San Diego Healthcare System.

HARD encoding was achieved by generating gradient directions equally spaced on a sphere by tessellations of an icosahedron (13,15), as shown in Fig. 11. This procedure produces directions that are equally separated in angle on the surface of a sphere. Single-shot images were acquired at nine

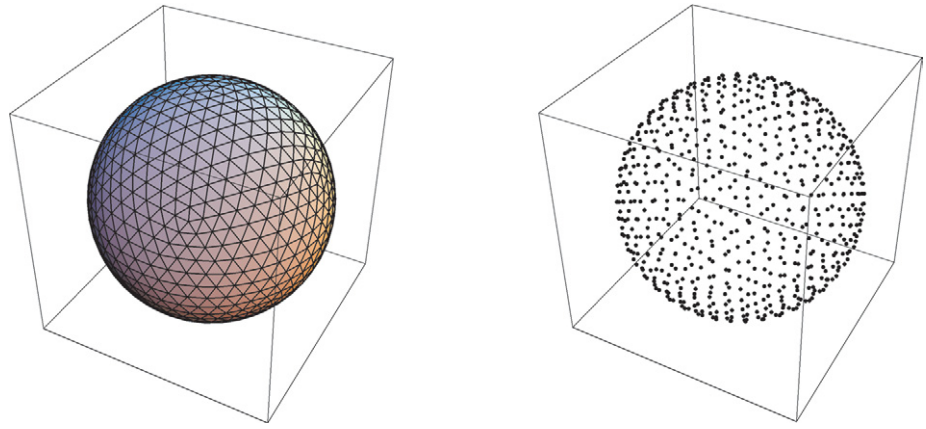

FIG. 11. Diffusion-encoding directions generated by the pulse sequence are spherical tessellations of an icosahedron (13) of user-specified order (shown here for clarity is the fifth-order tessellation). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley-.com.]

slices with the following parameters: FOV = 24 cm, slice thickness = 3.8 mm, and matrix size 64 × 64 for approximately (3.75 mm × 3.75 mm × 3.8 mm) isotropic resolution, TR = 2700 ms, TE = 52 ms. The diffusion parameters were: diffusion gradient duration, δ = 20 ms, stimulated-echo mixing time TM = 200 ms, and $b \approx 3000$ s/mm^2, and 43 diffusion directions determined by the icosahedral tessellations of a sphere. Twenty averages at each diffusion direction were collected to ensure high signal-to-noise ratios (SNRs), and resulted in a total scan time of ≈ 34 min.

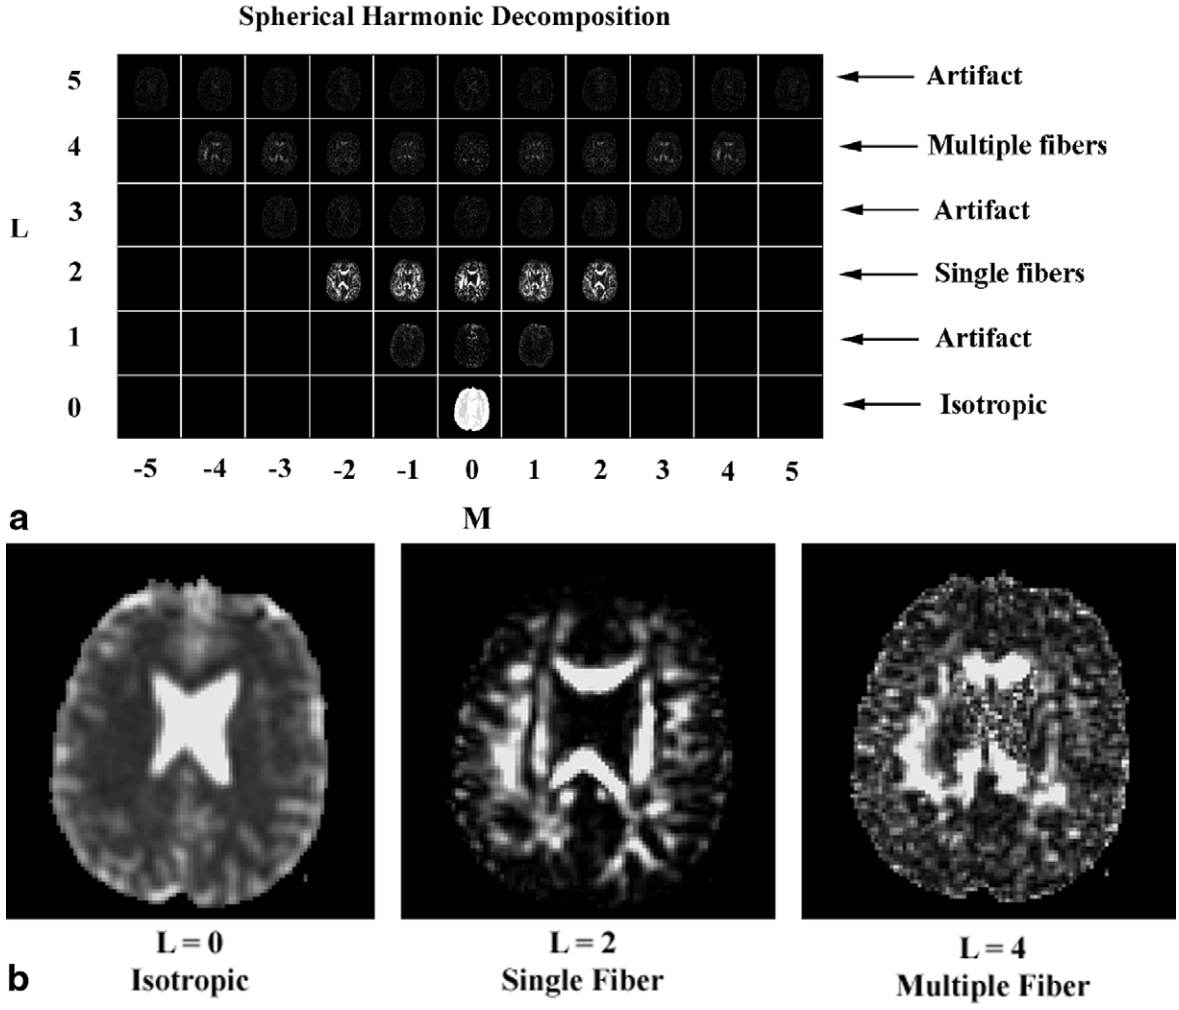

FIG. 12. SHD and the separate fiber channels it produces.

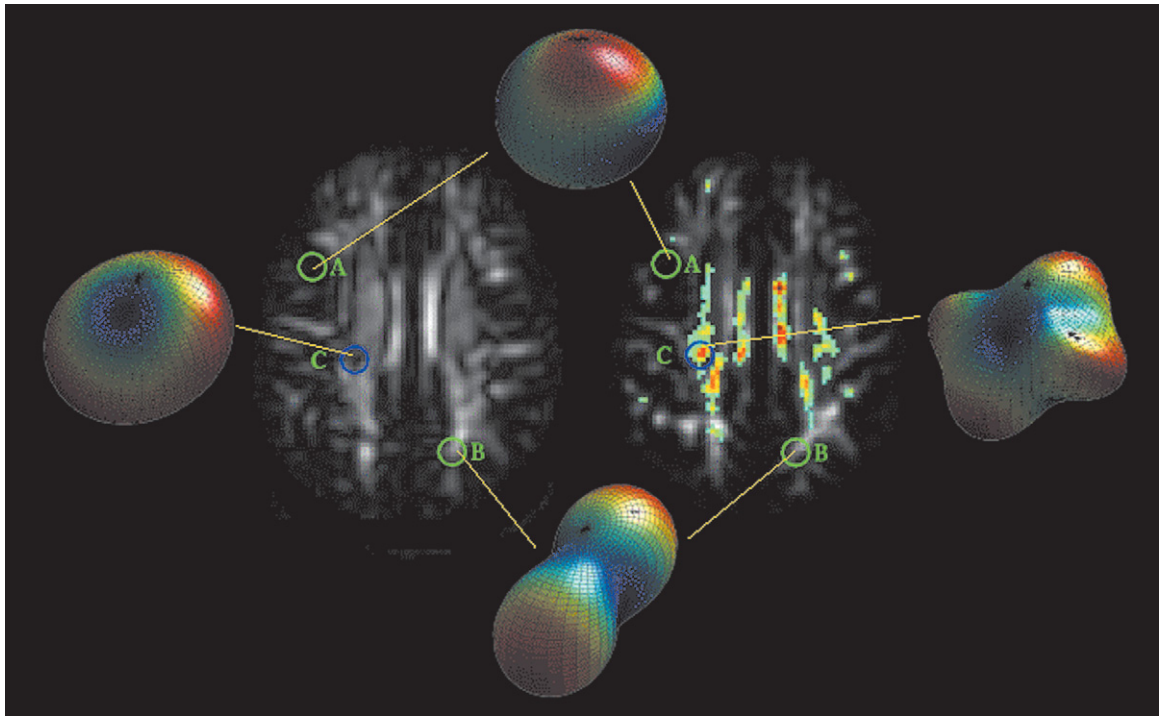

FIG. 13. D_{app} estimated from the L = 0, 2, 4 fiber channels. (Top left) Single-fiber channel (grayscale) overlaied in color with multifiber channel (FMI > .4) (color scale). (Top right) Grayscale image is relative anisotropy index determined from diffusion tensor calculation from the same data. In circles are three points from the (**A**) isotropic, (**B**) single fiber, and (**C**) multifiber channels for which D_{app} is (colored shapes) reconstructed from all three SHT components $L = 0, 2, 4$. The gray matter voxel (**A**) is essentially isotropic, so that D_{app} is a sphere and is unaffected by the inclusion of $L = 2, 4$. The single-fiber voxel (**B**) has the characteristic peanut shape, which is unaffected by the inclusion of $L = 4$. The voxel in **C** requires $L = 0, 2, 4$ to represent D_{app}, which in this case is consistent with two identical but perpendicular fibers. **c**: The DTI reconstruction, on the other hand, can only accurately reconstruct the isotropic and single-fiber channels. Both (**A**) isotropic and (**B**) single-fiber voxels are correct, but (**C**) the multifiber voxel produces a nearly spherical D_{app}, with very low anisotropy and a subsequent "black hole" in the anisotropy image.

RESULTS

The spherical harmonic decomposition (up to order $L = 5$) of a single slice from a HARD data set collected in a normal human volunteer is shown in Fig. 12. Note that the energy is indeed confined to the $L = 0,2,4$. The isotropic, single-fiber, and multiple-fiber channels reconstructed from this transform are shown at the bottom of the figure. As predicted, the isotropic channel appears to consist of gray

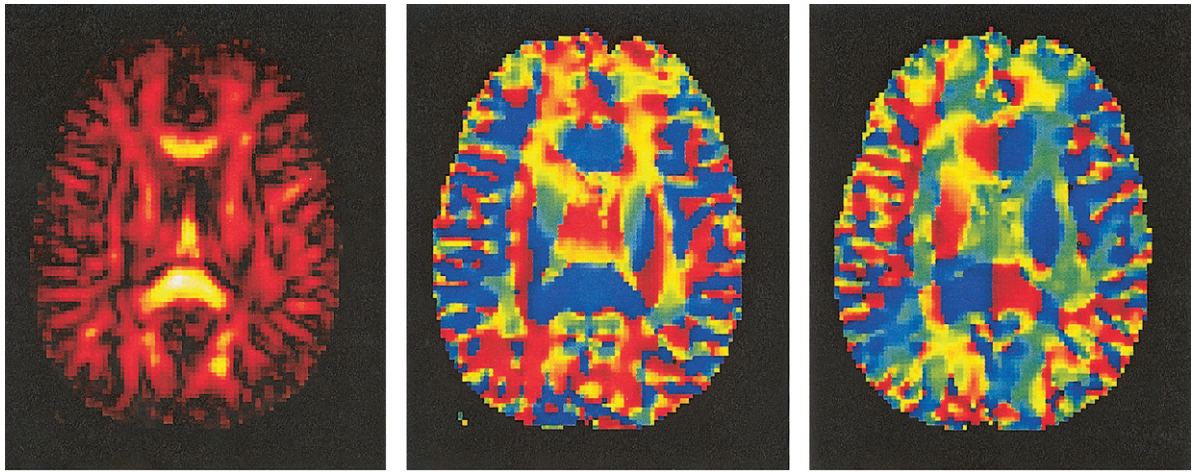

FIG. 14. Fiber magnitude and orientation. **a:** Magnitude of a single-fiber channel, representing the presence of single fibers, **b:** Map of estimated θ. **c:** Map of estimated ϕ.

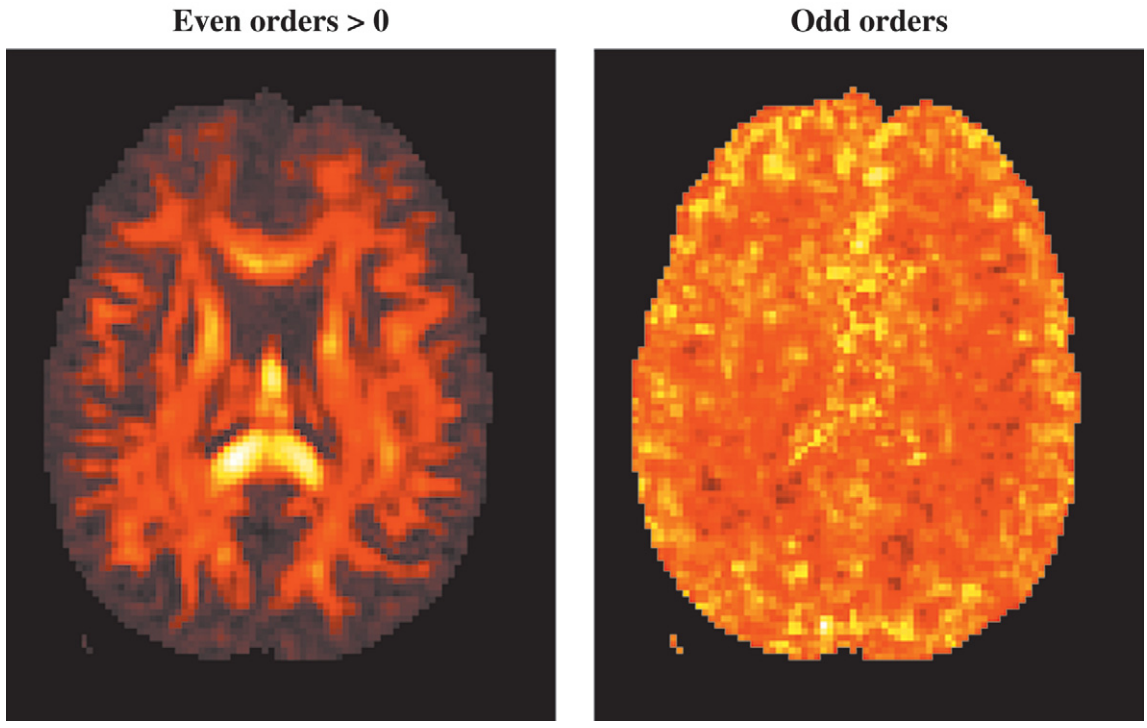

FIG. 15. Left: A white-matter map reconstructed by adding the energy in the $L = 2, 4$ channels. Right: A map made from the energy in the odd-order components ($L = 1, 3, 5$) shows little structure. These images are not on the same scale, but are scaled independently for display. [Color figure can be viewed in the online issue; which is available at www.interscience.wiley.com.]

matter, the single-fiber channel looks like a white-matter map, and the multifiber channel corresponds anatomically to regions of complex fiber geometry. The measured $D_{app}(\theta, \phi)$ from the multifiber channel of Fig. 12 are compared with simulations in Fig. 16, assuming two fibers oriented $\Delta\phi$ relative to one another. From the coefficients of the SHT the local D_{app} can be calculated. In Fig. 13 are shown representative shapes from the three different fiber channels in another slice from the same data set.

For the representative voxel from the multiple-fiber voxel the extreme case of nearly identical fibers crossed at 90° was chosen, because the failure of the DTI method is most apparent in this circumstance. However, it is important to note that regions of significant energy in the multifiber channel do not necessarily correspond to "black holes" in the standard DTI maps (such as those shown in Ref. 15) since this effect is produced only in the specific case of identical fibers oriented at 90° to one another. However, these regions do correspond to regions in which characterization by a single diffusion tensor is incorrect, and will produce spurious results in the estimation of diffusivities and orientations.

The magnitude of the single-fiber channel, and the estimated maps of θ and ϕ, are shown in Fig. 14 for a different slice from the same data set. A white-matter map is reconstructed from the energy in the $L = 2, 4$ channels, regardless of any categorization as single or multiple fiber, and shown in Fig. 15, along with a map of the energy in the odd channels which shows little structure.

DISCUSSION AND CONCLUSIONS

MR's sensitivity to diffusion anisotropy (4) was recognized shortly after the initial studies of isotropic diffusion (2,3), but it really blossomed with the recognition of its utility for the study of fibrous biological systems (10,11). The natural first step in quantifying anisotropic diffusion is to assume a Gaussian model for a single fiber, from which a diffusion tensor model ensues (8). From this model can be estimated the fiber diffusivities and orientation, and the conditions sufficient to capture this information (9). While it was recognized that this method is restricted to single fibers (8,12), development of a technique to measure more complex systems was only recently proposed (13). With the general goal being to investigate complex anatomical structures with no a priori assumptions about the diffusion characteristics, the sampling criterion is no longer based upon a presumed model, and this is phrased in terms of being "unbiased" in the sense that no direction is assumed to be preferable. The result is sampling on a sphere with a radius defined by the *b*-value with sufficient density to detect diffusion changes in different directions (13). The question then arises: How does one characterize the diffusion from such measurements? One approach, suggested by Wedeen et al. (14), is to extend the measurements to several radii by collecting "shells" of high angular spherical resolution data and Fourier transforming the data to produce a "*q*-space" image (26). This method is sensitive to restricted diffusion because it samples a range of *b*-values, but is unnecessarily complicated for the determination of non-Gaussian diffusion arising, for instance, from multiple-fiber directions. As pointed out in Ref. 15, the magnitude of *b*, in conjunction with high angular resolution sampling, is sufficient for this purpose and allows data acquisition in a time that is clincally feasible. The more pertinent question is how to characterize HARD mea-

surements made at high *b*-values, which is the subject of the present work.

Our approach to characterizing the diffusion measured with the HARD technique is based on characterizing the shape of the measured apparent diffusion coefficient. Using the methods of group theory, this characterization is that HARD measurements can be decomposed into irreducible representations of the rotation group in which isotropic, single-fiber, and multiple-fiber components are three, separable, direct-sum subspaces. In our initial investigation of this problem, deviations from a spherical surface in the form of the variance of the measurements were used as a measure of anisotropy, the spherical diffusion variance (15). While this has the advantage over the diffusion tensor method in identifying regions of anisotropy not well characterized by a single diffusion tensor, it is unable to quantify the anisotropy in any meaningful way. In particular, it does not allow the quantitation of either the magnitude or the direction of diffusion. In addition, all non-spherical noise or artifacts contribute to this variance. The present work can be seen as a formalization of this approach, since the spherical harmonics higher than Y_0^0 that contribute to the spherical diffusion variance characterize how the various anisotropic components contribute to the variance.

A strength of the approach is that it does not require any a priori information about the diffusion. The utility of the decomposition results from the group algebra imposed by the symmetries of both the measurement scheme and the diffusion. The decomposition allows distinction of diffusive and nondiffusive signal variations, as well as distinction among diffusion variations. In particular, the diffusion channels can be broken down into direct-sum subspaces representing isotropic, single-fiber, and multiple-fiber components. Asymmetries produced by experimental artifacts fall into channels that are impossible to reach by diffusion, thereby providing a direct means of noise reduction within the diffusion channels as well as a means of identifying artifactual effects. The technique lends itself to prior information and model selection, which will be required to properly estimate voxel composition.

The numerical computation of the SHT was implemented, by direct computation of the coefficients by a discretized version of the integral Eq. [21], in a fashion analogous to a DFT on a sphere. Unfortunately, there is no matrix decomposition of the SHT analogous to that used to implement the FFT from the DFT that would facilitate a fast SHT. Several methods for fast SHT computation have been proposed (e.g., Refs. 23,24, and 27–30). Our implementation based on direct integration using spherical Voronoi weight is sufficiently fast and accurate for most practical sampling schemes.

A very general problem that arises is the determination of the voxel composition from the SHD. The solution to this problem hinges on the ability to estimate the parameters of a voxel, in particular, the number of fibers, their volume fractions, anisotropies, and orientations. Even in the simple example discussed above of two fibers, it was shown that it is not possible to determine all of these uniquely, as the volume fraction and relative fiber orientation both affect the higher-order SHD components in a similar fashion. In practice, fiber configurations within a voxel may be much more complicated than the simple two-fiber model, making the problem of parameter determination exceedingly complicated.

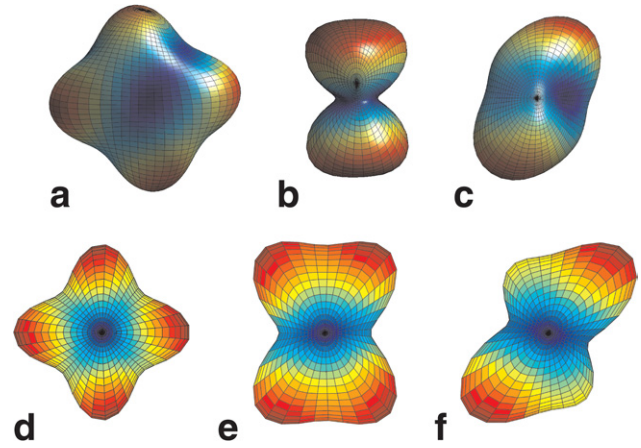

FIG. 16. Evidence of consistency with the two-fiber model. D_{app} measured from multifiber channel in Fig. 12 (top) and simulated (bottom) for (**d**) $f_1 = f_2$, $\delta\phi = 90°$, (**e**) $f_1 = 2f_2$, $\delta\phi = 90°$, (**f**) $f_1 = 2f_2$, and $\delta\phi = 75°$. The shapes have been interpolated from the diffusion-directional sampling for clarity. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

The power of the SHD, however, is the fact that the identification of the existence of multiple fibers is not dependent upon making this distinction: multiple fibers of any sort show energy in the higher channels. Moreover, even in lieu of a particular fiber model, the SHD allows the shape of the diffusion measurements to be quantified by the coefficients of the SHT. These may be then used to reconstruct the diffusion structure in each voxel depending upon the model used. With constraints, such as on the number of fibers, the computational complexity can be reduced. An example is shown in Fig. 16, in which two fibers were assumed. One solution for the composition in both angles and volume fractions was estimated simply by trial and error in the simulation. However, one can imagine formalizing this process by choosing a limited maximum number of fibers and searching for the relative angular displacements between the fibers and the fiber volume fractions in order to estimate these quantities. Incorporation of other imaging information may augment this estimation by determining which fiber model is most probable in any particular voxel.

To assess the significance of the power in the multiple-fiber channels in order to determine whether a voxel actually contains multiple-fiber directions, some comparison with the single-fiber channel is required. For this purpose, a simple statistic for determining the significance of the multifiber channel, the FMI, was introduced. While this is a natural measure, there is no indication of its optimality. A more formal probabilistic analysis needs to be undertaken to determine a method for determining the significance of energy among the channels.

Remember that what we have loosely referred to as "fibers" are really fiber bundles, within which the water movement that produces the diffusion signal is most likely

complex. We have implicitly assumed that all the fibers that make up a bundle are essentially identical. Moreover, we have also made the assumption throughout this work is that there is no exchange between fibers, and thus the signals from the individual fibers add independently. This will not be true in general. Both of these assumptions are made here, as they are in most of the DTI literature, not for lack of recognition, but because the true nature of the diffusion signal can be exceedingly complicated and is beyond the scope of the current study.

The characterization of fibers in terms of even orders of L is a consequence of the symmetry. Experimental artifacts, such as eddy current effects, that do not possess such symmetry are not confined to the even channels, and therefore appear in the odd L channels. Energy in these channels therefore indicates the presence of nondiffusion effects. The SHD, by automatically separating out some fraction of the non-diffusion energy, effectively reduces the noise in the diffusion channels. This has great potential for use on systems for which artifacts are present.

In practice, the issue of multiple fiber directions within a voxel is intimately tied to the image resolution: for higher resolution there will be fewer voxels with multiple directions. Nevertheless, there will always remain heterodirectional voxels at any resolution. Moreover, the penalty in SNR per unit time can make high-resolution diffusion imaging over a large region prohibitive. A postprocessing scheme capable of accurately identifying and quantifying multiple-fiber voxels may lead to more efficient acquisition protocols.

One important application of our method is its incorporation of multifiber voxels into fiber-tract mapping schemes. This will require utilizing estimates of the individual fiber orientations and volume fractions determined from the SHD of individual voxels. Additional machinery to keep track of multiple possible pathways of fibers passing through such voxels will then be necessary.

It is worth reiterating, in summary, that the proposed method reduces to the standard diffusion tensor method in the presence of single-fiber voxels, so no penalty of information is imposed by its usage. Rather, deviations from the DTI model due to artifacts or multiple fiber directions are readily extracted and quantified, allowing a more complete description of complex diffusion processes in tissues using a clinically efficacious diffusion encoding scheme.

ACKNOWLEDGMENTS

I thank Dr. Gary Glover, Stanford University, for graciously providing the basic spiral pulse sequence and reconstruction software upon which the sequence was developed, and Drs. Eric Wong, and Tom Liu for helpful discussions on diffusion. I am grateful to Drs. Richard Buxton and Robert Cox for reading versions of the manuscript and offering several useful suggestions. I would also like to thank Dr. Pete Basser, NIH, for inspiring me to clarify certain key points, and Dr. David Meyer, Dept. of Mathematics, UCSD, for lively discussions on group theory.

REFERENCES

1. LeBihan D, Basser PJ. Molecular diffusion and nuclear magnetic resonance. In: LeBihan D, editor. Diffusion and perfusion magnetic resonance imaging, New York: Raven Press; 1995.
2. Carr HY, Purcell EM. Effects of diffusion on free precession in nuclear magnetic resonance experiments. Phys Rev 1954;94:630–638.
3. Hahn EL. Detection of sea-water motion by nuclear precession. J Geophys Res 1960;65:776.
4. Stejskal EO. Use of spin echoes in a pulsed magnetic field gradient to study anisotropic, restricted diffusion and flow. J Chem Phys 1965;43: 3597–3603.
5. Moseley ME, Cohen Y, Mintorovitch J, Chileuitt L, Tsuruda J, Norman D, Weinstein P. Evidence of anisotropic self-diffusion in cat brain. In: Proceedings of the 8th Annual Meeting of SMRM, Amsterdam, 1989. p 136.
6. Moseley ME, Kucharczyk J, Asgari HS, Norman D. Anisotropy in diffusion weighted MRI. Magn Reson Med 1991;19:321–326.
7. Basser PJ, Pierpaoli C. Microstructual and physiological features of tissues elucidated by quantitative diffusion tensor MRI. J Magn Reson B 1996;111:209–219.
8. Basser PJ, Mattiello J, LeBihan D. Estimation of the effective self-diffusion tensor from the NMR spin echo. J Magn Reson 1994;103:247–254.
9. Basser PJ, Pierpaoli C. A simplified method to measure the diffusion tensor from seven MR images. Magn Reson Med 1998;39:928–934.
10. Chenevert TL, Brunberg JA, Pipe JG. Anisotropic diffusion in human white matter: demonstration with MR techniques in vivo. Radiology 1990;177:401–405.
11. Doran M, Hajnal JV, Van Bruggen N, King MD, Young IR, Bydder G. Normal and abnormal human white matter tracts shown by MR imaging using directional diffusion weighted sequences. J Comput Assist Tomogr 1990;14:865–873.
12. Pierpaoli C, Jezzard P, Basser PJ, Barnett A, DiChiro G. Diffusion tensor MR imaging of the human brain. Radiology 1996;201:637–648.
13. Tuch DS, Weisskoff RM, Belliveau JW, Wedeen VJ. High angular resolution diffusion imaging of the human brain. In: Proceedings of the 7th Annual Meeting of ISMRM, Philadelphia, 1999. p 321.
14. Wedeen VJ, Reese TG, Tuch DS, Weigel MR, Dou J-G, Weiskoff RM, Chesler D. Mapping fiber orientation spectra in cerebral white matter with Fourier transform diffusion MRI. In: Proceedings of the 8th Annual Meeting of ISMRM, Denver, 2000. p 82.
15. Frank LR. Anisotropy in high angular resolution diffusion weighted MRI. Magn Reson Med 2001;45:935–939.
16. Frank LR. Characterization of anisotropy in high angular resolution diffusion weighted MRI. In: Proceedings of the 9th Annual Meeting of ISMRM, Glasgow, Scotland, 2001. p 1531.
17. Hsu EW, Mori S. Analytical expressions for the NMR apparent diffusion coefficients in an anisotropic system and a simplified method for determining fiber orientation. Magn Reson Med 1995;34:194–200.
18. Arfken G. Mathematical methods for physicists. 2nd ed. New York: Academic Press; 1970.
19. Strang G. Linear algebra and its applications. 1st ed. New York: Academic Press, New York; 1976.
20. Elliott JP, Dawber PG. Symmetry in physics. Vol. I. Principles and simple applications. New York: MacMillan; 1979.
21. Wigner EP. Group theory and its application to the quantum mechanics of atomic spectra. New York: Academic Press; 1959. Translation by J.J. Griffin of the 1931 German edition.
22. Biedenharn LC, Louck JD. Angular momentum in quantum physics: theory and applications. Vol. VIII. Encyclopedia of mathematics and its applications. Reading, MA: Addison-Wesley, 1981.
23. Ritchie DW, Kemp GJL. Fast computation, rotation, and comparison of low resolution spherical harmonic molecular surfaces. J Math Chem 1999;20:383–395.
24. Mohlenkamp MJ. A fast transform for spherical harmonics. J Fourier Anal Appl 1999;5:159–184.
25. Cox RW. AFNI: software for the analysis and visualization of functional magnetic resonance neuroimages. Comput Biomed Res 1996;29:162–173.
26. Callaghan PT. Principles of nuclear magnetic resonance microscopy. Oxford: Oxford University Press; 1991.
27. Dilts GA. Computation of spherical harmonic expansion coefficients via FFT's. J Comput Phys 1985;57:439–453.
28. Swarztrauber PN. The vector harmonic transform method for solving partial differential equations in spherical geometry. Mon Wea Rev 1993;121:3415–3437.
29. Driscoll JR, Healy DM. Computing Fourier transforms and convolutions on the 2-sphere. Adv Appl Math 1994;15:202–250.
30. Adams JC, Swarztrauber P.N. SPHEREPACK 2.0: a model development facility. Technical Report TN-436-STR, NCAR, 1997.