Magnetic Resonance in Medicine 44:41–50 (2000)

Statistical Artifacts in Diffusion Tensor MRI (DT-MRI) Caused by Background Noise

Peter J. Basser[1]* and Sinisa Pajevic[2]

This work helps elucidate how background noise introduces statistical artifacts in the distribution of the sorted eigenvalues and eigenvectors in diffusion tensor MRI (DT-MRI) data. Although it was known that sorting eigenvalues (principal diffusivities) by magnitude introduces a bias in their sample mean within a homogeneous region of interest (ROI), here it is shown that magnitude sorting also introduces a significant bias in the variance of the sample mean eigenvalues. New methods are presented to calculate the mean and variance of the eigenvectors of the diffusion tensor, based on a dyadic tensor representation of eigenvalue–eigenvector pairs. Based on their use it is shown that sorting eigenvalues by magnitude also introduces a bias in the mean and the variance of the sample eigenvectors (principal directions). This required the development of new methods to calculate the mean and variance of the eigenvectors of the diffusion tensor, based on a dyadic tensor representation of eigenvalue–eigenvector pairs. Moreover, a new approach is proposed to order these pairs within an ROI. To do this, a correspondence between each principal axis of the diffusion ellipsoid, an eigenvalue–eigenvector pair, and a dyadic tensor constructed from it is exploited. A measure of overlap between principal axes of diffusion ellipsoids in different voxels is defined that employs projections between these dyadic tensors. The optimal eigenvalue assignment within an ROI maximizes this overlap. Bias in the estimate of the mean and of the variance of the eigenvalues and of their corresponding eigenvectors is reduced in DT-MRI experiments and in Monte Carlo simulations of such experiments. Improvement is most significant in isotropic regions, but some is also observed in anisotropic regions. This statistical framework should enhance our ability to characterize microstructure and architecture of healthy tissue, and help to assess its changes in development, disease, and degeneration. Mitigating these artifacts should also improve the characterization of diffusion anisotropy and the elucidation of fiber-tract trajectories in the brain and in other fibrous tissues. Magn Reson Med 44:41–50, 2000. Published 2000 Wiley-Liss, Inc.†

Key words: MRI; diffusion; DTI; DT-MRI; tensor; dyadic; bias; noise; eigenvalue; eigenvector; artifact

Diffusion tensor MRI (DT-MRI) (1) combines a measurement of the effective diffusion tensor (2) and conventional MRI. From this tensor, three eigenvalues (principal diffusivities) and three eigenvectors (principal directions) are calculated in each voxel (1). These eigenvalues represent the effective (scalar) diffusivities along the three corresponding principal directions that define the local "fiber" coordinate system (1). The eigenvalues of the diffusion tensor, along with other scalar quantities calculated from them (such as the orientationally averaged diffusivity, $\frac{1}{3}$ Trace(**D**); as well as measures of diffusion anisotropy, similarity, and organization (1,3,4) characterize distinct microstructural and architectural features of tissue. In general, these MRI "stains" provide useful physiological information noninvasively, not only about the state of normal tissue, but about its changes in development, aging, disease, and degeneration.

When there is background noise present in diffusion-weighted images (DWI), it is not clear how to sort the eigenvalues consistently within a homogeneous region of interest (ROI). Ordering eigenvalues by increasing magnitude in each voxel results in a "sorting bias," whose severity increases as the signal-to-noise ratio (SNR) decreases (5). Using Monte Carlo simulations of DT-MRI experiments in ROIs containing isotropic or anisotropic media, Pierpaoli and Basser (6) showed that the sample mean of the largest sorted eigenvalue in an ROI, λ_1, is always larger than its true mean, whereas the sample mean of the smallest sorted eigenvalue within an ROI, λ_3, is always smaller than its true mean. This artifact results in an overestimate of the degree of diffusion anisotropy within each voxel at all SNRs (e.g., as measured by λ_1/λ_3) (6), and may lead one to the erroneous conclusion that differences among sorted eigenvalues within an ROI are statistically significant, even in isotropic media whose "true" eigenvalues are all equal. This artifact also has made it difficult to interpret whether differences among sorted eigenvalues measured in skeletal (7) and cardiac muscle (8) were biologically meaningful or were the result of noise.

In considering the problems of sorting bias, it was clear that there were no methods available to quantify, represent, and display bias in the distributions of both eigenvalues and eigenvectors of the diffusion tensor in a self-consistent manner. Here we propose new analytical and graphical methods for doing so. Correlations between sorted eigenvalues and eigenvectors (9) suggest that, if the diffusion tensor field within an ROI is homogeneous, then sorting eigenvalues and their corresponding eigenvectors in tandem could reduce these statistical artifacts. However, implementing this schema entails solving several challenging problems: to represent geometric and algebraic features of eigenvalue–eigenvector pairs (the former being scalars and the latter, vectors), and to develop a new measure of overlap between such eigenvalue–eigenvector pairs in different voxels. Using this framework, we can then evaluate the improvement in this proposed method over magnitude sorting. Some material here was previously presented in abstract form (10).

[1]Section on Tissue Biophysics and Biomimetics, Laboratory of Integrative and Medical Biophysics, National Institute of Child Health and Human Development (STBB/LIMB/NICHD), National Institutes of Health, Bethesda, Maryland.

[2]Mathematical and Statistical Computing Laboratory, Center for Information Technology (MSCL/CIT), National Institutes of Health, Bethesda, Maryland.

*Correspondence to: Peter J. Basser, STBB/LIMB/NICHD, National Institutes of Health, Bldg. 13, Rm. 3W16, 13 South Drive, Bethesda, MD 20892-5772. E-mail: pjbasser@helix.nih.gov

Received 21 May 1999; revised 17 February 2000; accepted 17 February 2000.

THEORY

Estimates of the Mean and Variance of the Distribution of Eigenvalues Within an ROI

To assess the severity of the sorting artifact, it is helpful to define various measures to characterize the bias in the mean and variance of the eigenvalues within an ROI. In Monte Carlo simulations, bias of the mean can be assessed easily, since the "true" eigenvalues are always prescribed or known in advance. For the three eigenvalues, the fractional bias in the mean is given by:

$$\frac{\lambda_1^{\mathrm{ROI}} - \lambda_1^{\mathrm{True}}}{\lambda_1^{\mathrm{True}}}; \quad \frac{\lambda_2^{\mathrm{ROI}} - \lambda_2^{\mathrm{True}}}{\lambda_2^{\mathrm{True}}}; \quad \frac{\lambda_3^{\mathrm{ROI}} - \lambda_3^{\mathrm{True}}}{\lambda_3^{\mathrm{True}}} \qquad [1]$$

where λ_i^{ROI} represents the sample ROI-averaged eigenvalue, the sample mean eigenvalue is

$$\lambda_i^{\mathrm{ROI}} = \frac{1}{N}\sum_{j=1}^{N} \lambda_i^j \qquad [2]$$

and $\lambda_i^{\mathrm{True}}$ represents the true eigenvalue. The unbiased sample variance of each eigenvalue is

$$\langle \sigma_i^2 \rangle = \left(\frac{1}{N-1}\right)\sum_{j=1}^{N} (\lambda_i^j - \lambda_i^{\mathrm{ROI}})^2 \qquad [3]$$

where it is assumed that the ROI contains N voxels.

A measure of the dispersion of the distribution of eigenvalues, which is sensitive to their order or assignment, is the range. Dividing this quantity by the mean of this distribution yields

$$\mathrm{Range/Mean} = \frac{\lambda_1 - \lambda_3}{\langle \lambda \rangle} \qquad [4]$$

in which it is assumed that $\lambda_1 > \lambda_3$. The mean diffusivity $\langle \lambda \rangle$, which is not sensitive to sorting order (1), provides a reliable reference standard for comparing the range of eigenvalues in different tissue compartments, since $\langle \lambda \rangle$ is relatively unbiased for SNR above about 5 (11), and has been shown to be virtually indistinguishable in normal gray and white matter in cat (12) and in human (13) brain. Moreover, in cerebrospinal fluid (CSF) the larger expected difference between λ_1 and λ_3 is normalized by a proportionately larger denominator. Ideally, in isotropic regions, such as gray matter and CSF, Range/Mean (R/M) should vanish. Finally, R/M grows linearly with the sorting bias in λ_1 and λ_3. Measures of diffusion anisotropy that we proposed previously, such as the relative and fractional anisotropy (3,4), by design are insensitive to the order in which eigenvalues are sorted, and unlike R/M cannot be used to assess the efficacy of an eigenvalue–eigenvector sorting algorithm.

Estimates of the Mean and Variance of the Distribution of Eigenvectors Within an ROI

To date, it has not been possible to reliably report statistics about the eigenvectors within an ROI because of the inherent sign ambiguity of each eigenvector. For a given principal diffusivity (or eigenvalue λ_i), the sign of its corresponding eigenvector ε_i is indeterminate (i.e., one is free to choose either $+\varepsilon_i$ or $-\varepsilon_i$). Data with such antipodal symmetry are called *axial* data, the statistics of which are described elsewhere (14). This sign ambiguity complicates the calculation of ROI-averaged eigenvectors, because one can use either the positive or the negative eigenvector when calculating the sample mean (ROI-averaged eigenvector). So, arithmetically averaging eigenvectors within an ROI produces a poor estimate of their mean.

However, representing each eigenvector as a second-order dyadic tensor[3] allows us to calculate a sample mean eigenvector unambiguously, as well as to quantify the dispersion about the mean within an ROI. Rather than averaging the eigenvectors themselves, we first average their dyadics within an ROI:

$$\langle \varepsilon_i\, \varepsilon_i^T \rangle = \left\langle \begin{pmatrix} \varepsilon_{i_x}^2 & \varepsilon_{i_x}\varepsilon_{i_y} & \varepsilon_{i_x}\varepsilon_{i_z} \\ \varepsilon_{i_x}\varepsilon_{i_y} & \varepsilon_{i_y}^2 & \varepsilon_{i_y}\varepsilon_{i_z} \\ \varepsilon_{i_x}\varepsilon_{i_z} & \varepsilon_{i_y}\varepsilon_{i_z} & \varepsilon_{i_z}^2 \end{pmatrix} \right\rangle = \frac{1}{N}\sum_{j=1}^{N} \varepsilon_i^j\, \varepsilon_i^{jT}. \qquad [5]$$

Here we have represented the positive, semidefinite second-order dyadic tensor as a 3×3 symmetric matrix whose elements contain products of the x, y, and z components of the eigenvector ε_i. To find the ROI-averaged eigenvector $\langle \varepsilon_i \rangle$, we calculate the eigenvector of $\langle \varepsilon_i \varepsilon_i^T \rangle$ above that corresponds to its largest eigenvalue β_1^i (14). We quantify the bias in $\langle \varepsilon_i \rangle$ by using the angle between the true eigenvector and the sample mean eigenvector, $\Delta\theta_i$, (obtained from $\langle \varepsilon_i \varepsilon_i^T \rangle$):

$$\Delta\theta_i = \arccos(|\langle \varepsilon_i \rangle \cdot \varepsilon_i^{\mathrm{True}}|). \qquad [6]$$

Moreover, the two remaining eigenvalues of $\langle \varepsilon_i \varepsilon_i^T \rangle$, β_2^i and β_3^i (14), can be used to characterize the dispersion about the sample mean eigenvector $\langle \varepsilon_i \rangle$. A dispersion measure we propose is the square root of the ratio of the average of the two smaller eigenvalues and the largest one:

$$\sqrt{\frac{\beta_2^i + \beta_3^i}{2\beta_1^i}}. \qquad [7]$$

This quantity attains a minimum value of 0 when there is no scatter about the mean eigenvector (i.e., $\beta_2^i = \beta_3^i = 0$) and attains a maximum of 1 when the eigenvectors are uniformly distributed about the sphere (i.e., $\beta_1^i = \beta_2^i = \beta_3^i$). Geometrically, Eq. [7] is akin to the radius of the "cone of uncertainty" around an eigenvector having a unit length (16). This dispersion measure also grows approximately linearly with the standard deviation of normally distributed angular data ($\alpha \ll \pi/2$).

[3] The first use of the dyadic tensor in DT-MRI was in Ref. 15 to classify different types of diffusion anisotropy. Here the dyadic tensor is used in an altogether different context.

Measuring Overlap Between Eigenvalue–Eigenvector Pairs

Our proposed method to reduce artifacts in magnitude sorting of eigenvalues is to sort them and their corresponding eigenvectors in pairs, because in a homogeneous ROI, eigenvalues and eigenvectors contain complementary information that together specify the size, shape, and orientation of the diffusion ellipsoid that characterizes diffusion in each voxel. If we consistently assign or order the principal axes of the ellipsoids within an ROI, so as to maximize the overlap among these ellipsoids, we could ameliorate the sorting bias that results from considering only the lengths of the principal axes of these diffusion ellipsoids.

This method entails (1) establishing a correspondence between a principal axis of a diffusion ellipsoid, an eigenvalue–eigenvector pair, and a dyadic tensor; (2) defining a new quantity that measures the degree of overlap between dyadic tensors; and (3) consistently sorting dyadics of the diffusion tensor within an ROI by maximizing this overlap measure.

The dyadic tensor framework described previously also assists us in developing a means to measure the degree of overlap of the principal axes of the diffusion tensor in different voxels. A natural way to represent eigenvalue–eigenvector pairs is by using a second-order dyadic tensor space. We can construct a dyadic tensor from each eigenvalue–eigenvector pair λ_i and ε_i, by taking the outer product of the weighted eigenvector, as follows:

$$(\sqrt{\lambda_i}\,\varepsilon_i)(\sqrt{\lambda_i}\,\varepsilon_i)^T = \lambda_i\,\varepsilon_i\,\varepsilon_i^T = \lambda_i \begin{pmatrix} \varepsilon_{i_x}^2 & \varepsilon_{i_x}\varepsilon_{i_y} & \varepsilon_{i_x}\varepsilon_{i_z} \\ \varepsilon_{i_x}\varepsilon_{i_y} & \varepsilon_{i_y}^2 & \varepsilon_{i_y}\varepsilon_{i_z} \\ \varepsilon_{i_x}\varepsilon_{i_z} & \varepsilon_{i_y}\varepsilon_{i_z} & \varepsilon_{i_z}^2 \end{pmatrix} \quad [8]$$

The vectors $\sqrt{\lambda_i}\varepsilon_i$ coincide with the principal axes of a diffusion ellipsoid (1,17).

Just as one uses the vector dot product to determine the overlap between two vectors $\sqrt{\lambda_i}\varepsilon_i$ and $\sqrt{\lambda'_i}\varepsilon'_i$, one uses the tensor dot product ":" to determine the overlap between two dyadic tensors (18) $\lambda_i\varepsilon_i\varepsilon_i^T$ and $\lambda'_i\varepsilon'_i\varepsilon'^T_i$,

$$\lambda_i\,\varepsilon_i\,\varepsilon_i^T : \lambda'_i\,\varepsilon'_i\,\varepsilon'^T_i = \mathrm{Trace}(\lambda_i\,\varepsilon_i\,\varepsilon_i^T\,\lambda'_i\,\varepsilon'_i\,\varepsilon'^T_i) = \lambda_i\,\lambda'_i\,(\varepsilon_i\cdot\varepsilon'_i)^2 \quad [9]$$

which represent eigenvalue–eigenvector pairs in different voxels. Geometrically, we can view "overlap" as the degree to which two diffusion ellipsoids (whose three major axes have lengths of $\sqrt{\lambda_1}$, $\sqrt{\lambda_2}$, $\sqrt{\lambda_3}$ and $\sqrt{\lambda'_1}$, $\sqrt{\lambda'_2}$, $\sqrt{\lambda'_3}$, respectively) resemble each other, i.e., have similar size, shape, and orientation (3,19). Other properties of these dyadic tensors that allow us to establish the correspondence between them and the principal axes of the diffusion ellipsoid are described in the Appendix.

To account for the three-dimensional character of anisotropic diffusion, we sum the "overlaps" between corresponding eigenvalue–eigenvector pairs or dyadics in two different voxels:

$$\sum_{i=1}^{3} \lambda_i\,\lambda'_i\,(\varepsilon'_i\cdot\varepsilon_i)^2. \quad [10]$$

This quantity contains all eigenvalue–eigenvector pairs, weighting the dot product between eigenvectors by the product of the square root of their corresponding eigenvalues. Normalizing this quantity using its global maximal value $\sum_{i=1}^{3}\lambda_i\lambda'_i$, we obtain a new nondimensional scalar measure of intervoxel overlap between two diffusion tensors, C_t:

$$C_t = \frac{\sum_{i=1}^{3} \lambda_i\,\lambda'_i\,(\varepsilon_i\cdot\varepsilon'_i)^2}{\sum_{i=1}^{3} \lambda_i\,\lambda'_i}. \quad [11]$$

By construction, $0 < C_t \leq 1$, where 0 indicates no overlap and 1 indicates complete overlap.

Note that the dot product of corresponding eigenvectors is squared. This removes the sign ambiguity discussed earlier, since whatever the signs of ε_i and ε'_i, the square of their dot product is always the same value. The tensor space distance measure has the additional desirable feature of always yielding the smallest distance between any three pairs of dyadics.

It is also possible to consider an overlap function that uses only the eigenvectors to sort the eigenvalues. This quantity measures the degree of overlap between unit dyadics representing the principal axes of diffusion tensors in different voxels:

$$C_t^* = \frac{1}{3}\sum_{i=1}^{3} (\varepsilon_i\cdot\varepsilon'_i)^2. \quad [12]$$

The scheme in Eq. [11] appears reasonable in some ordered anisotropic fibrous tissues, where it is usually more important to assign the (largest) eigenvalue (i.e., the one associated with the fiber-tract direction) correctly. (A notable exception is anisotropic diffusion, in which the diffusion ellipsoid assumes a pancake geometry, where the eigenvector associated with the smallest eigenvalue defines the axis of symmetry of diffusion.) But, in isotropic media, where all eigenvalues are equal, these methods properly assign each eigenvalue–eigenvector pair an equal weight. Second, particularly when SNR is low, some eigenvalues estimated using multivariate linear regression (2) could be negative (6). Although this is not plausible physically, it is nonetheless possible mathematically, since each diffusion tensor is not explicitly constrained to be positive definite (i.e., having all positive eigenvalues). If an eigenvalue is negative, then negative cross-terms of the form $\lambda_i\lambda'_j$ could appear in Eq. [11], contrary to our earlier assumption; so when calculating C_t, we must first check whether any eigenvalue is negative. If so, we set it to a negligibly small positive number (e.g., 10^{-6} μm^2/sec). This problem does not arise when using Eq. [12].

As an aside, since we plan to use this overlap measure to sort eigenvalues and eigenvectors of the diffusion tensor, it must be sensitive to their order. Therefore, we could not

use previously proposed measures of diffusion similarity (4), such as

$$S = \frac{\sum_{j=1}^{3}\sum_{i=1}^{3} \lambda_i \lambda'_j (\varepsilon_i \cdot \varepsilon'_j)^2}{\sum_{j=1}^{3} \lambda'_j \sum_{i=1}^{3} \lambda_i} \qquad [13]$$

which are insensitive to the order or assignment of the eigenvalue–eigenvector pairs (1,3,4).

Overlap Measurements Within an ROI

So far, we have described a method to compare the degree of overlap between pairs of dyadic tensors in different voxels. To apply this to sorting eigenvalue–eigenvector pairs within an ROI, we first take the arithmetic average of all the diffusion tensors within an ROI to determine the ROI-averaged diffusion tensor. We then find and sort its eigenvalues by magnitude. This average tensor becomes the reference tensor we use subsequently to sort each eigenvalue–eigenvector pair within each voxel of the ROI. We choose the order of the eigenvalue–eigenvector pairs within each voxel that maximizes the intervoxel overlap function with respect to the reference tensor. Using these new assignments, we now obtain ROI-sorted values of λ_1, λ_2, and λ_3, as well as ε_1, ε_2, and ε_3.

Caveats and Clarifications

Some precautions must be taken in using this approach. First, the overlap measure should be applied to voxels within an ROI in which the diffusion tensor field is approximately homogeneous. This condition can be tested by using maps of Trace(**D**), of the organizational (1,20) and lattice anisotropy indices (6), and of the fiber direction field within the ROI (1,20–22). By homogeneity of diffusion properties within an ROI, we do not assume that fibers within the voxels are necessarily continuous or connected to each other, just that their diffusion properties or characteristics (3,4) are similar. In fact, we can consider ROIs containing tissue in different regions.

NUMERICAL RESULTS

Monte Carlo simulations of DT-MRI experiments described previously in Ref. 6 were performed using MR parameters given in Ref. 13. Simulations were carried out for ROIs assumed to contain homogeneous tissue with diffusion properties representative of brain white matter, gray matter, and CSF obtained in DT-MRI studies of normal human brain (13). Human brain parenchyma was assumed to have Trace(**D**) = $\lambda_1 + \lambda_2 + \lambda_3 = 2100$ μm^2/sec (13).

In Fig. 1, the ROI-averaged λ_1, λ_2, and λ_3, and their standard errors are given vs. the signal-to-noise ratio (SNR). Here, 5000 Monte Carlo repetitions were performed to obtain precise estimates of the means and standard deviations. Orange, Yellow, and Purple lines indicate ROI-averaged eigenvalues sorted by magnitude; Red, Green, and Blue lines indicate ROI-averaged eigenvalues sorted using this new method. In Fig. 1a, results are shown for isotropic gray matter using true values, $\lambda_1 = \lambda_2 = \lambda_3 = 700$ μm^2/sec; in Fig. 1b, results are shown for fully anisotropic tissue satisfying $\lambda_1 > \lambda_2 > \lambda_3$ (i.e., all three eigenvalues distinct) with $\lambda_1/\lambda_2 = 1.5$ and $\lambda_2/\lambda_3 = 1.5$; in Fig. 1c, results are shown for tissue with prolate or cigarlike diffusion properties (i.e., true diffusivities satisfying $\lambda_1 > \lambda_2 = \lambda_3$), with $\lambda_1/\lambda_2 = 1.5$; and in Fig. 1d, results are shown for tissue with oblate or pancakelike diffusion properties (i.e., true diffusivities satisfying $\lambda_1 = \lambda_2 > \lambda_3$), with $\lambda_2/\lambda_3 = 1.5$. In all cases, the ROI-averaged λ_1 sorted by magnitude is significantly more biased above the true value than the ROI-averaged λ_1 sorted by the dyadic method, whereas the ROI-averaged λ_3 sorted by magnitude is significantly more biased below the true value than the ROI-averaged λ_3 sorted by the dyadic method for all SNR. Moreover, the sample standard deviations for all ROI-averaged eigenvalues sorted by magnitude are significantly smaller than those for the ROI-averaged eigenvalues sorted by the dyadic method for all SNR (see Tables 1 and 2). These tables also show a significant reduction in the bias and an increase in the variance of all the ROI-averaged

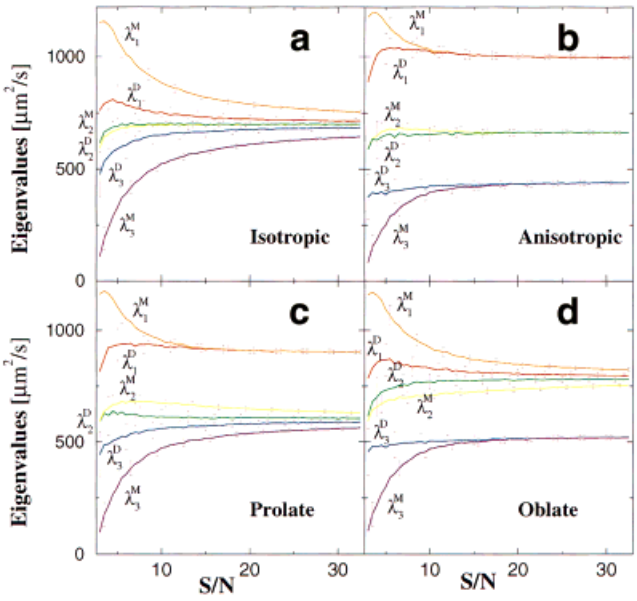

FIG. 1. Monte Carlo simulations of DT-MRI experiments for a 5 × 5 ROI. The ROI-averaged λ_1, λ_2, and λ_3, and their standard errors are given vs. the signal-to-noise ratio (SNR). Orange, Yellow, and Purple lines indicate ROI-averaged eigenvalues sorted by magnitude; Red, Green, and Blue lines indicate ROI-averaged eigenvalues sorted using the new dyadic tensor formalism. **a:** Isotropic gray matter with true values $\lambda_1 = \lambda_2 = \lambda_3 = 700$ μm^2/sec. **b:** Fully anisotropic tissue (i.e., three distinct eigenvalues with $\lambda_1 > \lambda_2 > \lambda_3$). **c:** Tissue with prolate or cigarlike diffusion properties (i.e., $\lambda_1 > \lambda_2 = \lambda_3$). **d:** Tissue with oblate or pancakelike diffusion properties (i.e., $\lambda_1 = \lambda_2 > \lambda_3$). The ROI-averaged λ_1 with magnitude sorting is significantly more biased above its true value than with dyadic sorting, whereas the ROI-averaged λ_3 with magnitude sorting is significantly more biased below its true value than with dyadic sorting for all SNR. Moreover, the sample standard deviations for all ROI-averaged eigenvalues are significantly smaller when sorted by magnitude than when sorted by the dyadic method for all SNR.

Table 1
Summary of Monte Carlo DT-MRI Simulation Results for Four Paradigmatic Tissue Types Showing Bias and Variances for Both Eigenvalues and Eigenvectors of the Effective Diffusion Tensor at SNR = 10*

Anisotropic (1.2:1.2)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	10.6	0.5	−15.5	2.76	2.72	3.67
Dyadic	3.1	−0.3	−3.8	3.92	4.73	5.65
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	9.1	23.4	10.4	0.54	0.73	0.57
Dyadic	7.7	9.6	7.5	0.35	0.39	0.35
Cigar (1.2:1)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	14.0	5.8	−22.4	2.81	2.78	2.98
Dyadic	2.8	2.8	−5.9	4.29	5.24	5.09
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	13.1	50.1	45.7	0.8	0.7	0.6
Dyadic	9.5	46.3	45.7	0.4	0.4	0.4
Pancake (1:1.2)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	20.6	−5.8	−17.8	2.90	2.45	3.28
Dyadic	5.4	−2.6	−3.4	4.60	4.60	5.51
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	46.8	51.6	13.3	0.7	0.8	0.6
Dyadic	46.2	46.6	9.3	0.4	0.4	0.4
Isotropic (1:1)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	26.4	−0.6	−25.8	1.15	1.05	1.57
Dyadic	7.1	−0.1	−7.0	1.70	1.71	1.93
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	22.4	51.9	24.1	0.8	0.8	0.8
Dyadic	8.1	8.6	8.3	0.4	0.4	0.4

*Trace(**D**) for all simulations is 2100 μm^2/sec. For the isotropic case, $\lambda_1 = \lambda_2 = \lambda_3 = 700$ μm^2/sec. For the anisotropic case, $\lambda_1/\lambda_2 = 1.2$ and $\lambda_2/\lambda_3 = 1.2$. For the prolate or cigarlike case, $\lambda_1/\lambda_2 = 1.2$. For the oblate or pancakelike case, $\lambda_2/\lambda_3 = 1.2$.

eigenvalues at SNRs of 10 and 25 for the four paradigmatic cases of diffusion anisotropy.

Using the same tissue properties as before, we also investigated the effect that missorting eigenvalues has on the distribution of their corresponding eigenvectors. To do this, we introduce a new way to represent the three dyadics in each voxel and to visualize their distribution within an ROI. In Fig. 2, the pairs $\pm\sqrt{\lambda_1}\varepsilon_1$, $\pm\sqrt{\lambda_2}\varepsilon_2$, and $\pm\sqrt{\lambda_3}\varepsilon_3$ are displayed as principal axes of a diffusion ellipsoid, and assigned Red, Green, and Blue, respectively, after sorting by either method. For all voxels within the ROI, these objects are collated and displayed together. As a visual aid, an ellipsoid with principal axes $\pm 0.9\sqrt{\lambda_1^{True}}\varepsilon_1^{True}$, $\pm 0.9\sqrt{\lambda_2^{True}}\varepsilon_2^{True}$, and $\pm 0.9\sqrt{\lambda_3^{True}}\varepsilon_3^{True}$ is constructed to show the size, shape, and orientation of the underlying root-mean-square (rms) displacement ellipsoid.

First, in Fig. 2a and b we assume the same isotropic diffusion properties as in Fig. 1a. In Fig. 2a, eigenvalues are sorted by magnitude. The distribution of the eigenvectors appears uniform, but one can see the concentric spherical shells containing Red, Green, and Blue points. In Fig. 2b, the eigenvalue–eigenvector pairs are sorted using the new method. Here, the three eigenvectors are localized within different sectors of the spherical ellipsoid, but colored points are no longer concentrically distributed. In Fig. 2c and d we use the diffusion properties as in Fig. 1d for an oblate or pancakelike tissue. In Fig. 2c, in the case of magnitude sorting, there are clusters of the Blue points around the axis of symmetry with concentric bands of uniformly distributed Red and Green points distributed around the equatorial plane. In Fig. 2d, in the case of dyadic sorting, the Blue points

Table 2
Summary of Monte Carlo DT-MRI Simulation Results for Four Paradigmatic Tissue Types Showing Bias and Variances for Both Eigenvalues and Eigenvectors of the Effective Diffusion Tensor at SNR = 25*

Anisotropic (1.2:1.2)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	1.5	0.2	−2.5	1.23	1.33	1.76
Dyadic	1.2	−0.2	−1.5	1.32	1.63	2.01
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	3.0	4.7	3.4	0.2	0.4	0.3
Dyadic	3.1	4.4	3.3	0.2	0.3	0.1
Cigar (1.2:1)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	2.3	4.9	−7.8	1.23	1.22	1.29
Dyadic	1.4	0.7	−2.5	1.44	1.93	1.89
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	4.2	46.2	45.3	0.3	0.7	0.6
Dyadic	4.3	45.4	45.3	0.3	0.4	0.4
Pancake (1:1.2)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	6.9	−4.4	−3.1	1.15	1.05	1.57
Dyadic	2.3	−0.8	−1.7	1.70	1.71	1.93
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	45.9	45.5	4.5	0.6	0.7	0.4
Dyadic	45.4	45.5	4.5	0.4	0.4	0.3
Isotropic (1:1)	$\%\left(\frac{\lambda_1^{ROI}-\lambda_1^{True}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\lambda_2^{ROI}-\lambda_2^{True}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\lambda_3^{ROI}-\lambda_3^{True}}{\lambda_3^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_1^{ROI}}{\lambda_1^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_2^{ROI}}{\lambda_2^{True}}\right)$	$\%\left(\frac{\bar{\sigma}_3^{ROI}}{\lambda_3^{True}}\right)$
Magnitude	10.3	−0.1	−10.2	1.15	1.01	1.10
Dyadic	2.8	−0.07	−2.8	1.96	1.97	1.94
	$\Delta\theta_1$[deg]	$\Delta\theta_2$[deg]	$\Delta\theta_3$[deg]	$\sqrt{\frac{\beta_2^1+\beta_3^1}{2\beta_1^1}}$	$\sqrt{\frac{\beta_2^2+\beta_3^2}{2\beta_1^2}}$	$\sqrt{\frac{\beta_2^3+\beta_3^3}{2\beta_1^3}}$
Magnitude	23.1	52.6	23.9	0.8	0.8	0.8
Dyadic	8.2	8.7	8.2	0.4	0.4	0.4

*Trace(**D**) for all simulations is 2100 μm^2/sec. For the isotropic case, $\lambda_1 = \lambda_2 = \lambda_3 = 700$ μm^2/sec. For the anisotropic case, $\lambda_1/\lambda_2 = 1.2$ and $\lambda_2/\lambda_3 = 1.2$. For the prolate or cigarlike case, $\lambda_1/\lambda_2 = 1.2$. For the oblate or pancakelike case, $\lambda_2/\lambda_3 = 1.2$.

still clustered around the axis of symmetry, but fewer Green and Red points contaminate this region. Moreover, the Red and Green points no longer lie in different "orbits" around the equatorial plane; Green and Red clusters are now clearly visible. In Fig. 2e and f we use the diffusion properties for an anisotropic or "asymmetric" tissue as in Fig. 1b. In Fig. 2e, magnitude-sorting results in a separation of Red, Green, and Blue points, but each colored cluster is clearly contaminated by points having other colors. In Fig. 2f, using dyadic sorting, the distribution of the colored points is more localized near their respective "true" principal axes, and there is concomitantly less color contamination within each cluster. The fact that there are relatively sharp boundaries between the color clusters indicates that measures in Eqs. [11] and [12] yield similar results in homogeneous ROIs.

Tables 1 and 2 also illustrate features of the eigenvector distribution. They show a significant reduction in both the angular bias $\Delta\Theta$ and in the angular dispersion at SNRs of 10 and 25. Note that the reduction of angular dispersion in the degenerate cases, in which two or three eigenvalues are equal, is artificial and we do not assign any significance to it.

Magnitude and dyadic sorting are also compared in Fig. 3. ROI-averaged λ_1, λ_2, and λ_3, and their standard errors are given as a function of the number of voxels within the ROI. Depending on the size of the ROI, 1000 to 10,000 Monte Carlo repetitions were performed to obtain precise estimates of ROI-mean and the standard errors of these eigenvalues. Results are shown for isotropic and anisotropic cases. The long-dashed brown lines indicate the eigenvalues of the ROI-averaged tensor. Note, that the dyadically sorted eigenvalues never converge to the true values, even for infinitely large ROIs.

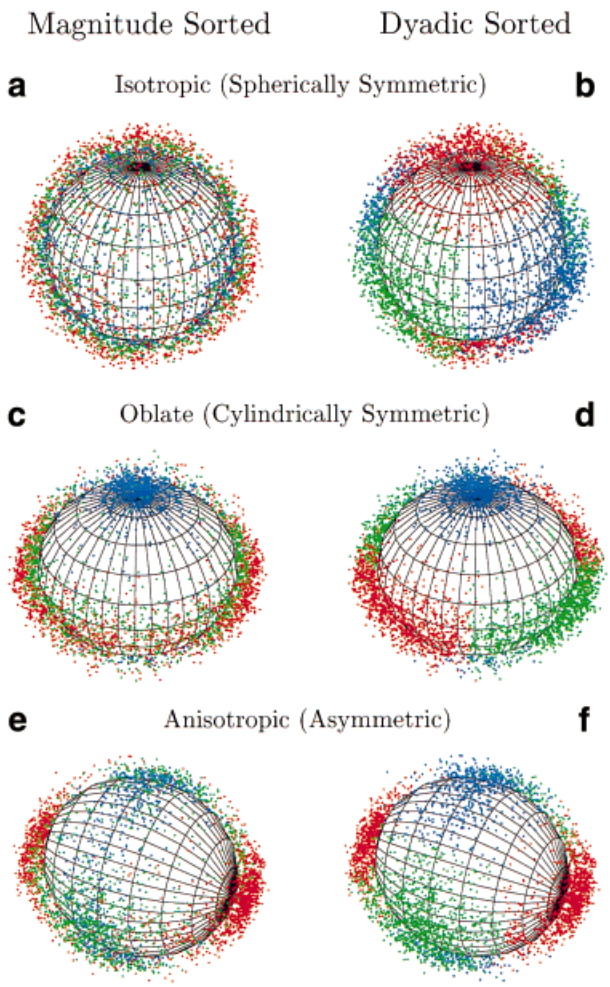

FIG. 2. Spatial distribution of the eigenvectors within an ROI weighted by their respective eigenvalues when sorted by magnitude and by the dyadic tensor method. The pairs $\pm\sqrt{\lambda_1}\varepsilon_1$, $\pm\sqrt{\lambda_2}\varepsilon_2$, and $\pm\sqrt{\lambda_3}\varepsilon_3$, are assigned Red, Green, and Blue points, respectively. The points are plotted with respect to a rms displacement ellipsoid whose principal axes are $\pm 0.9\sqrt{\lambda_1^{\text{True}}}\varepsilon_1^{\text{True}}$, $\pm 0.9\sqrt{\lambda_2^{\text{True}}}\varepsilon_2^{\text{True}}$, and $\pm 0.9\sqrt{\lambda_3^{\text{True}}}\varepsilon_3^{\text{True}}$. **a,b:** Isotropic diffusion properties are those used in Fig. 1a. **c,d:** Diffusion properties are assumed for an oblate or pancakelike tissue as in Fig. 1d. **e,f:** Diffusion properties are assumed for an anisotropic or "asymmetric" tissue as in Fig. 1b.

DISCUSSION

When the eigenvalues are sorted by magnitude, the means and variances of their ROI-averaged eigenvalues are biased (e.g., see Fig. 1). There is an artifactual bias in the sample mean eigenvalues and an artifactual reduction in their variance. This combination could lead one to the erroneous conclusion that differences between these eigenvalues are statistically significant, particularly in isotropic media (see Fig. 1a) where, in the absence of background noise, all "true" eigenvalues are equal. Figure 1 also shows that this new sorting method significantly reduces the bias in the estimate of the mean and increases the estimated variance of each ROI-averaged eigenvalue at all SNR, in all Monte Carlo simulations. These two effects ensure greater overlap of the distributions of the ROI-averaged eigenvalues. Moreover, the same improvement is seen in the simulation of anisotropic media, including fully anisotropic (Fig. 1b), prolate or cigarlike (Fig. 1c), and oblate or pancakelike (Fig. 1d), and in the distributions of eigenvalues over the entire range of SNRs. It should also be noted in Fig. 1 that at SNR of 15 and greater there is no sorting bias for anisotropic tissue. This, of course, is dependent on the degree of anisotropy. For less-anisotropic structures sorting bias can be significant even at much larger SNRs.

Figure 2 presents a new way to illustrate graphically the misclassification of magnitude-sorting eigenvalue–eigenvector pairs. In Fig. 2e, there are many more Green points in the Red region, and many more Blue points in the Green region than we see in Fig. 2f, where the colored clusters are more homogeneous. This indicates that after dyadic sorting, there are fewer misclassifications. We also see from Fig. 2 that properly classifying eigenvalues within an ROI decreases the variance of their corresponding eigenvectors (i.e., scatter about the mean eigenvector). This is reflected in the clouds of colored points being more dispersed in magnitude sorting than in dyadic sorting. The reduction in the dispersion of eigenvectors is also demonstrated in Tables 1 and 2.

This reduction in the variance of the eigenvectors, accompanied by an increase in the variance of the distribution of the eigenvalues, appears counterintuitive. This paradox can be understood by recognizing that 1) eigenvalues and their corresponding eigenvectors are always calculated and sorted in pairs, and that 2) the three eigenvectors in each voxel are mutually orthogonal. From 1) we see that when eigenvalues are misclassified, so are their corresponding eigenvectors. Thus, from 2) we conclude that swapping eigenvalues always leads to swapping orthogonal eigenvectors, leading to a broadening of the eigenvec-

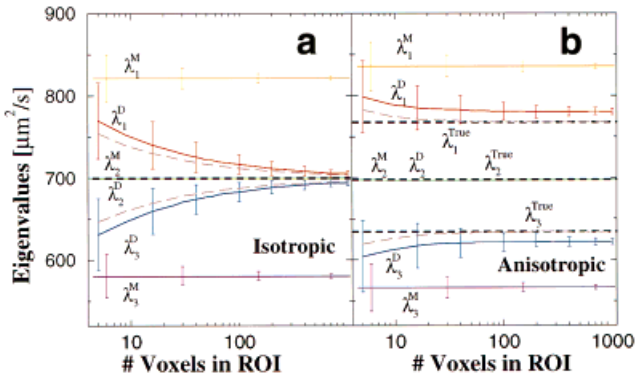

FIG. 3. ROI-averaged λ_1, λ_2, and λ_3, and their standard errors given as a function of the number of voxels within the ROI, at SNR = 15 for two cases (**a**) isotropic diffusion tensor with $\lambda_1 = \lambda_2 = \lambda_3 = 700$ μm^2/sec, and (**b**) anisotropic diffusion tensor with $\lambda_1 = 768$, $\lambda_2 = 698$, and $\lambda_3 = 634$ μm^2/sec, respectively. Orange, Yellow, and Purple lines indicate ROI-averaged eigenvalues sorted by magnitude. Bias is high and persists as ROI size increases in magnitude sorting. Red, Green, and Blue lines indicate ROI-averaged eigenvalues sorted using the new dyadic tensor formalism. Bias is initially lower and drops monotonically as a function of ROI size in dyadic sorting, but there is a persistent although small residual bias in dyadic sorting. Brown long-dashed lines indicate the eigenvalues of the averaged tensor within the ROI. Black dashed lines indicate the true eigenvalues.

tor distribution. We also see how magnitude sorting introduces a bias in the ROI-averaged eigenvectors. In simulations performed on anisotropic white matter shown in Fig. 2e and f, missorting the eigenvalues takes place more frequently between λ_1 and λ_2 (i.e., between Red and Green) and between λ_3 and λ_2 (i.e., between Green and Blue) than between λ_1 and λ_3 (i.e., between Red and Blue). Thus, ε_1 is swapped with ε_2 and ε_3 is swapped with ε_2 more frequently than ε_1 is swapped with ε_3. This asymmetric exchange process causes both the ROI-averaged eigenvectors $\langle\varepsilon_1\rangle$ and $\langle\varepsilon_3\rangle$ to be tipped toward $\langle\varepsilon_2\rangle$. This mixing results in a directional bias in and a loss of mutual orthogonality among the three ROI-averaged eigenvectors.

A related problem arises in comparing results from overlap measures Eqs. [11] and [12]. In principle, Eq. [12] possesses some advantages over the overlap measure we proposed earlier. First, when eigenvalues are negative, Eq. [12] does not behave pathologically. Second, when the difference between the true eigenvalues is comparable to or smaller than the sum of their standard deviations, weighting eigenvectors by their eigenvalues may result in erroneous transpositions, like those that already occur in magnitude sorting. However, at higher SNRs, and in more anisotropic regions, Eq. [11] should provide additional information with which to assign the eigenvalues correctly within each voxel. In general, the smaller the variability or uncertainty in the eigenvectors is compared with the variability in the eigenvalues, the better Eq. [12] will perform compared with Eq. [11]. Still, this discussion only highlights the large gap in our understanding of the relationship between uncertainties of the eigenvalues and of the eigenvectors of the diffusion tensor as measured in DT-MRI. Further study is required to be able to quantitatively assess the relative merits of these approaches.

Artifacts associated with sorting eigenvalues by magnitude have impeded the implementation of a number of promising color-based schemes to represent information contained in diffusion tensor data. For example, Latour (23) suggested a clever color imaging method to represent several features of anisotropic diffusion simultaneously within a single image by encoding the sorted eigenvalues in each voxel λ_1, λ_2, and λ_3 using Red, Green, and Blue intensities, respectively. According to the preceding arguments, however, bias in the magnitude-sorted eigenvalues would produce an image whose Red intensity is always artifactually high and whose Blue intensity is artifactually low, which is what we find experimentally. A scheme to sort eigenvalues reliably and robustly would ameliorate this chromatic aberration, particularly in the isotropic regions. Similarly, Pierpaoli and Pajevic (24,25) recently proposed a novel method to display the direction of fiber tracts within each voxel. In its simplest form, it entails encoding the three components of the eigenvector associated with the largest eigenvalue in each voxel λ_1, ($|\varepsilon_{1x}|$, $|\varepsilon_{2y}|$, $|\varepsilon_{3z}|$), by using Red, Green, and Blue intensities, respectively. According to the previous arguments, missorting the largest eigenvalue in a voxel will cause the colors to be incorrectly assigned. In particular, it would cause an observer to infer that the fiber tract is perpendicular to its actual orientation. This artifact would also reduce the reliability of fiber tractography (21), a new method to follow nerve and other fiber tracts using DT-MRI data. If any of the ε_1 is missorted along a tract trajectory, then the tract following routine would inadvertently take a right-angle turn there, causing the computed path to deviate from the true path of the tract.

Systematic artifacts in DWIs, such as ghosting due to motion, eddy current distortion, and so forth, should all be corrected before trying to sort eigenvalues and eigenvectors within ROIs. For example, motion artifacts that make an isotropic medium (such as CSF or gray matter) appear organized and anisotropic will not be corrected by dyadic sorting. They introduce an apparent anisotropy in the medium that will only be further enhanced by our sorting procedure. Thus, we must be mindful of these and other systematic artifacts, particularly when attempting more challenging tasks using DT-MRI data, such as establishing whether a fiber tract is continuous, ascertaining whether there are long-range connections between different parts of the brain (or between other tissues), or measuring whether the degree of coherence of fiber-tract trajectories in vivo varies in development or disease.

In Fig. 3, bias is high in magnitude sorting (Orange and Purple lines) and remains high even as the size of the ROI increases. Bias is initially lower in dyadic sorting (Red and Blue lines) and drops monotonically as a function of ROI size. Within the homogeneous ROI, eigenvalues of the averaged tensor (Brown long-dashed lines) represent a lower bound on the bias of any eigenvalue sorting scheme. When comparing this lower bound with the dyadically sorted values we see that even for an infinitely large ROI there is still some persistent bias whose magnitude depends on the nature of the noise in the experimental tensor data. If the nature of noise is such that much of the variability in tensor data comes from perturbations of its eigenvalues (shape of the diffusion ellipsoid changes but not its orientation) then there will be no persistent bias. In the opposite case, where variability comes from perturbations in orientation only, then the primary assumption of dyadic sorting is violated, and the scheme actually introduces bias. As we have shown in our simulations, in DT-MRI the nature of noise is such that the dyadic-sorting schemes can significantly reduce bias. However, the experimental variability in orientation and eigenvalue magnitudes are not decoupled. Thus, there is a persistent bias that is a function of SNR and diminishes for large SNR.

Finally, this new method is fundamentally different from one we proposed recently to eliminate artifacts in the estimate of ROI-averaged eigenvalues, even within heterogeneous anisotropic regions, such as the corpus callosum (11). This method involved 1) calculating the three scalar invariants in each voxel I_1, I_2, and I_3 of the diffusion tensor (11), which are insensitive to the order of the eigenvalues within a voxel; 2) averaging the scalar invariants over an entire ROI; 3) using these averages as coefficients of a characteristic equation; 4) solving this equation for three ROI-averaged eigenvalues; and finally 5) sorting these three eigenvalues by magnitude. In this way, magnitude sorting, which introduces the unwanted bias, is deferred to the final step, where its deleterious effects were expected to be mitigated. However, Monte Carlo simulations performed to test this method showed only a marginal improvement in the bias of the ROI-averaged eigenvalues when compared with sorting them by magnitude

within each voxel (11) in ROIs with a small number of voxels. It appears now that, whereas I_1 = Trace(**D**) (11) is Gaussian distributed (26) and unbiased at most SNRs, I_2 and I_3 are not Gaussian distributed. Moreover, the three eigenvalues λ_1, λ_2, and λ_3 are coupled to I_1, I_2, and I_3 through the nonlinear characteristic equation, which complicates the determination of a parametric distribution of the eigenvalues.

Statistical artifacts resulting from noise in DWIs have complicated the characterization of diffusion properties in normal, diseased, developing, and aging tissues. In particular, difficulties arise in determining the distribution of eigenvalues, in measuring diffusion anisotropy, and in determining fiber-tract direction fields and trajectories. Bias in estimates of the mean of the eigenvalues (principal diffusivities) and of eigenvectors (principal directions) of the diffusion tensor, as measured by MRI (DT-MRI), which result from sorting eigenvalues by magnitude within each voxel, are elucidated here. A new framework for ordering eigenvalues and eigenvectors in tandem results in an increase in the variance of the distribution of sample eigenvalues and a concomitant decrease in the dispersion or variance of the distribution of eigenvectors in Monte Carlo–simulated diffusion tensor MRI data in media assumed to have properties similar to gray and white matter in human brain. This new methodology to sort eigenvalues and eigenvectors together within an ROI should further improve the quality and utility of DT-MRI data in characterizing tissue microstructure and architecture.

We see that using dyadic tensors to represent eigenvalue–eigenvector pairs overcomes many conceptual and technical problems encountered when representing them as vectors. We have shown how to use dyadic tensors to calculate the first and second moments of the eigenvector distribution within an ROI, which had not been possible previously. Moreover, we have introduced a novel color-based graphical means to display the distribution of eigenvalue–eigenvector pairs and to visualize their misclassification when eigenvalues are sorted by magnitude. Dyadic tensors allow us to measure "overlap" in a rigorous rather than ad hoc way, and provide a natural geometric interpretation of an eigenvalue–eigenvector pair as a principal axis of a diffusion ellipsoid. Still, this work highlights the need to develop a comprehensive understanding of the distribution of eigenvalues and eigenvectors obtained from diffusion tensor MRI data.

APPENDIX

Properties of the Dyadic Tensor $\lambda_i\varepsilon_i\varepsilon_i^T$

This dyadic tensor representation of eigenvalue–eigenvectors pairs possesses a number of desirable properties. First, by construction, corresponding eigenvectors and eigenvalues are always paired or grouped together. Second, although the dyadic tensor $\varepsilon_i\varepsilon_i^T$ is a three-dimensional object, the information is essentially one-dimensional. Specifically, $\varepsilon_i\varepsilon_i^T$ has only one nonzero eigenvalue whose value is 1. Its corresponding eigenvector is ε_i (or $-\varepsilon_i$). To see this, note that $(\varepsilon_i\varepsilon_i^T)\varepsilon_i = 1\ (\varepsilon_i)$ and $(-\varepsilon_i)\ (-\varepsilon_i^T)(-\varepsilon_i) = 1\ (-\varepsilon_i)$. The other two eigenvectors of $\varepsilon_i\varepsilon_i^T$ lie in a plane perpendicular to ε_i, both having eigenvalues of 0. Thus, the Range of $\varepsilon_i\varepsilon_i^T$ is one-dimensional (or linelike) and is spanned by ε_i, whereas the Null space of $\varepsilon_i\varepsilon_i^T$ is two-dimensional (or planelike) and is spanned by the two remaining orthogonal eigenvectors whose corresponding eigenvalues are zero. Third, unlike the vector ε_i, which encodes the direction of ε_i, the dyadic tensor $\varepsilon_i\varepsilon_i^T$ encodes the orientation of ε_i. It is easy to see that $\varepsilon_i\varepsilon_i^T$ is insensitive to the sign of ε_i, since $(+\varepsilon_i)(+\varepsilon_i^T) = (-\varepsilon_i)(-\varepsilon_i^T)$. Thus, geometrically, $\varepsilon_i\varepsilon_i^T$ possesses reflectional symmetry about the origin or antipodal symmetry. Fourth, the length of $\lambda_i\varepsilon_i\varepsilon_i^T$ equals λ_i. Specifically, it is given by

$$|\lambda_i\,\varepsilon_i\,\varepsilon_i^T|^2 = \lambda_i\,\varepsilon_i\,\varepsilon_i^T : \lambda_i\,\varepsilon_i\,\varepsilon_i^T = \lambda_i^2\ \mathrm{Trace}(\varepsilon_i\,\varepsilon_i^T\,\varepsilon_i\,\varepsilon_i^T) = \lambda_i^2\ \mathrm{Trace}(\varepsilon_i\,\varepsilon_i^T) = \lambda_i^2. \quad [A1]$$

Therefore, since λ_i is always positive for a diffusion tensor

$$|\lambda_i\,\varepsilon_i\,\varepsilon_i^T| = \lambda_i. \quad [A2]$$

Moreover, when $\lambda_i = 1$, the size or length of this dyadic tensor is unity. Thus, one can think of $\varepsilon_i\varepsilon_i^T$ as a symmetric object, like a baton, having a unit length whose orientation is coincident with that of ε_i, but which has no direction (i.e., no discernible "tip" or "tail"); $\lambda_i\varepsilon_i\varepsilon_i^T$ is a batonlike object with length λ_i, whose orientation is coincident with that of ε_i.

This geometric picture of the dyadic tensor is also useful when considering the decomposition of the diffusion tensor as a sum of its three orthogonal dyadic tensors:

$$\mathbf{D} = \lambda_1\varepsilon_1\varepsilon_1^T + \lambda_2\varepsilon_2\varepsilon_2^T + \lambda_3\varepsilon_3\varepsilon_3^T. \quad [A3]$$

Each dyadic $\lambda_i\varepsilon_i\varepsilon_i^T$ represents a component of **D** lying along one of the three principal axes of a diffusion ellipsoid. The surface of constant probability that defines the diffusion ellipsoid can then be viewed as resulting from the superposition of contributions from each of the three independent components of the diffusion tensor. Moreover, when we take a projection of the diffusion tensor along a particular direction (e.g., given by the unit vector **r**), we obtain a scalar apparent diffusion coefficient (ADC) along **r**,

$$\mathrm{ADC}_r = \mathbf{r}^T\mathbf{D}\mathbf{r} = \lambda_1(\mathbf{r}\cdot\varepsilon_1)^2 + \lambda_2(\mathbf{r}\cdot\varepsilon_2)^2 + \lambda_3(\mathbf{r}\cdot\varepsilon_3)^2. \quad [A4]$$

Using this decomposition, we see how each of the principal diffusivities is clearly weighted according to the square of the cosine of the angle that each principal axis makes with **r**. Thus, the dyadic representation again makes it clear why the principal axes of the diffusion ellipsoid have no inherent direction, only an orientation, since the preceding expressions are insensitive to the sign of **r** and to the sign of each of the eigenvectors of **D**.

REFERENCES

1. Basser PJ, Mattiello J, LeBihan D. MR diffusion tensor spectroscopy and imaging. Biophys J 1994;66:259–267.
2. Basser PJ, Mattiello J, LeBihan D. Estimation of the effective self-diffusion tensor from the NMR spin echo. J Magn Reson B 1994;103: 247–254.

3. Basser PJ. Inferring microstructural features and the physiological state of tissues from diffusion-weighted images. NMR Biomed 1995;8:333–344.
4. Basser PJ, Pierpaoli C. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. J Magn Reson B 1996;111:209–219.
5. Pierpaoli C, Basser PJ. New invariant "lattice" index achieves significant noise reduction in measuring diffusion anisotropy. In: Proceedings of the 4th Annual Meeting of ISMRM, New York, 1996. p 1326.
6. Pierpaoli C, Basser PJ. Toward a quantitative assessment of diffusion anisotropy [published erratum appears in Magn Reson Med 1997; 37(Jun):972]. Magn Reson Med 1996;36:893–906.
7. 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;32:194–200.
8. Garrido L, Wedeen VJ, Kwong KK, Spencer UM, Kantor HL. Anisotropy of water diffusion in the myocardium of the rat. Circ Res 1994;74:789–793.
9. Pierpaoli C, Basser PJ. Method to significantly reduce bias and variance of diffusion anisotropy measurements. US Patent: NIH Licensee, issued on Oct. 19, 1999.
10. Basser PJ, Pajevic S. Method to reduce eigenvalue sorting bias in DT-MRI. In: Proceedings of the 7th Annual Meeting of ISMRM, Philadelphia, PA, 1999. p 1788.
11. Basser PJ, Pierpaoli C. Estimating the principal diffusivities (eigenvalues) of the effective diffusion tensor. In: Proceedings of the 5th Annual Meeting of ISMRM, Vancouver, Canada, 1997. p 1739.
12. van Gelderen P, d. Vleeschouwer MHM, DesPres D, Pekar J, van Zijl PCM, Moonen CTW. Water diffusion and acute stroke. Magn Reson Med 1994;31:154–163.
13. Pierpaoli C, Jezzard P, Basser PJ, Barnett A, Di Chiro G. Diffusion tensor MR imaging of the human brain. Radiology 1996;201:637–648.
14. Bingham C. An antipodally symmetric distribution on the sphere. Ann Stat 1974;2:1201–1225.
15. Huppi PS, Maier SE, Peled S, Zientara GP, Barnes PD, Jolesz FA, Volpe JJ. Microstructural development of human newborn cerebral white matter assessed in vivo by diffusion tensor magnetic resonance imaging. Pediatr Res 1998;44:584–590.
16. Basser PJ. Quantifying errors in fiber direction and diffusion tensor field maps resulting from MR noise. In: Proceedings of the 5th Annual Meeting of ISMRM, Vancouver, Canada, 1997. p 1740.
17. Crank J. The mathematics of diffusion. Oxford, UK: Oxford University Press; 1975.
18. Morse PM, Feschbach H. Methods of theoretical physics. New York: McGraw-Hill; 1953.
19. Basser PJ. New historical and physiological stains derived from diffusion-tensor MR images. Ann NY Acad Sci 1997;820:123–138.
20. Basser PJ, LeBihan D. Fiber orientation mapping in an anisotropic medium with NMR diffusion spectroscopy. In: Proceedings of the Society for Magnetic Resonance in Medicine 11th Annual Meeting, Berlin, Germany, 1992. p 1221.
21. Basser PJ. Fiber-tractography via diffusion tensor MRI (DT-MRI). In: Proceedings of the 6th Annual Meeting of ISMRM, Sydney, Australia, 1998. p 1226.
22. Mori S, Crain BJ, Chacko VP, van Zijl PC. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. Ann Neurol 1999;45:265–269.
23. Latour L. Personal communication. London, UK: London Workshop of Diffusion Imaging; 1995.
24. Pajevic S, Pierpaoli C. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. Magn Reson Med 1999;42:526–540.
25. Pierpaoli C. Oh no! One more method for color mapping of fiber tract direction using diffusion MR imaging data. In: Proceedings of the 5th Annual Meeting of ISMRM, Vancouver, Canada, 1997. p 1741.
26. Basser PJ, Pajevic S. Quantitative statistical tests for assessing changes in the trace of the diffusion tensor: clinical and biological implications. In: Proceedings of the 7th Annual Meeting of ISMRM, Philadelphia, PA, 1999. p 1789.