Spatial Transformations of Diffusion Tensor Magnetic Resonance Images

D. C. Alexander*, *Member, IEEE*, C. Pierpaoli, P. J. Basser, and J. C. Gee

Abstract—We address the problem of applying spatial transformations (or "image warps") to diffusion tensor magnetic resonance images. The orientational information that these images contain must be handled appropriately when they are transformed spatially during image registration. We present solutions for global transformations of three-dimensional images up to 12-parameter affine complexity and indicate how our methods can be extended for higher order transformations. Several approaches are presented and tested using synthetic data. One method, the preservation of principal direction algorithm, which takes into account shearing, stretching and rigid rotation, is shown to be the most effective. Additional registration experiments are performed on human brain data obtained from a single subject, whose head was imaged in three different orientations within the scanner. All of our methods improve the consistency between registered and target images over naïve warping algorithms.

Index Terms—Diffusion tensor magnetic resonance imaging, DTI, DT-MRI, image processing, registration, spatial transformation, warp.

I. INTRODUCTION

DIFFUSION tensor magnetic resonance imaging (DT-MRI) is an MRI modality, recently introduced by Basser *et al.*, [1], [2], which combines a NMR (nuclear magnetic resonance) measurement of a diffusion tensor (DT) of water and conventional MRI. The measurement made at each voxel in a three-dimensional (3-D) DT-MR image is a symmetric second-order tensor, which describes the local diffusive behavior of water at the corresponding point in the imaged material. The DT may be represented by an ellipsoid describing the root-mean-squared displacement in each direction from the center of the voxel, [2], [3].

In this paper, we address the problem of applying spatial transformations to DT-MR images, which are required for registration, [4], [5], and spatial normalization of collections of images. Compared to the task of warping scalar images, the application of spatial transformations to DT-MR images is complicated by the fact that DTs contain orientational information, which is affected by the transformation. This effect must be accounted for in order to ensure the anatomical correctness of the transformed image. In this paper, we present and test various methods for applying affine transformations to these images. We go on to show that, once a suitable method has been devised for applying affine transformations to DT-MR images, the extension to higher order transformations is straightforward. Some preliminary work on this topic is described in [6] and [7], in which an elastic matching algorithm, [8]–[10], was used for inter-subject registration of DT-MR images of different subjects. The adaptation of such algorithms to work with DT-MR data remains the major motivation for the work described in this paper, although the techniques described should be relevant to any application in which tensor or vector images need to be processed.

A brief introduction to the properties of DT-MR images is provided in Section II. Then we discuss the problem of applying transformations to DT-MR images in Section III. Several methods are proposed and we justify theoretically the use of the PPD algorithm. Experiments to compare and validate the proposed methods have been performed over both real and synthetic DT-MR data sets. The design of the synthetic data sets, which is described in full in [11], is reviewed in Section IV and the experiments on this synthetic data and results obtained from them are presented. In addition, three DT-MR data sets have been obtained from a single human subject in different positions. An affine registration of the three images is obtained using the automated image registration (AIR) program of Woods *et al.*, [12], [13]. The human data sets and experiments on them are described in Section V. Finally, some concluding remarks are made in Section VI.

II. FEATURES OF THE DIFFUSION TENSOR

The primary features of a DT that we need to account for under a spatial image transformation of an image are the size, shape, and orientation of its ellipsoid. These can be characterized by the eigenvalues, λ_1, λ_2, and λ_3 and their corresponding eigenvectors, $\underline{e}_1, \underline{e}_2$, and $\underline{e}_3$, from the DT, which can be written as a 3×3 symmetric matrix, D. Henceforyh, we will assume that these eigenvalues are sorted so that $\lambda_1 > \lambda_2 > \lambda_3$; $\underline{e}_1$ is the major eigenvector of the DT (the principal direction) and $\underline{e}_3$ is the minor eigenvector. The eigenvectors describe the orientation of the DT and the eigenvalues describe its size and shape. The most common scalar indexes that are used to describe the size and shape of the DT are $\mathrm{Trace}(D)$, the anisotropy and the skewness. These correspond to the first, second, and third moments, respectively, of the distribution of eigenvalues of D, and
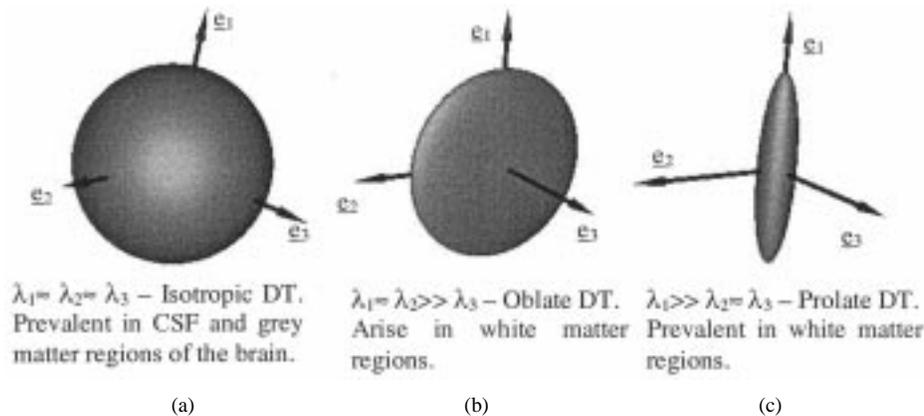
Fig. 1. Illustration of the different shapes of DT ellipsoids and the corresponding distribution of eigenvalues. (a) Isotropic DT is represented by a spherical ellipsoid; these measurements arise when diffusion is unhindered in all directions, as in regions of cerebro-spinal fluid (CSF) in the brain, or when it is hindered equally in all directions, as in grey matter, [14]. (b) Oblate DT is represented by a pancake shaped ellipsoid; this type of measurement can arise when tissue structure is planar, such as when white matter fibers intersect within a voxel, [14]. (c) Prolate DT is represented by a cigar shaped ellipsoid; this type of measurement arises when diffusion is hindered in two perpendicular directions, such as in white matter fiber tracts, [14].
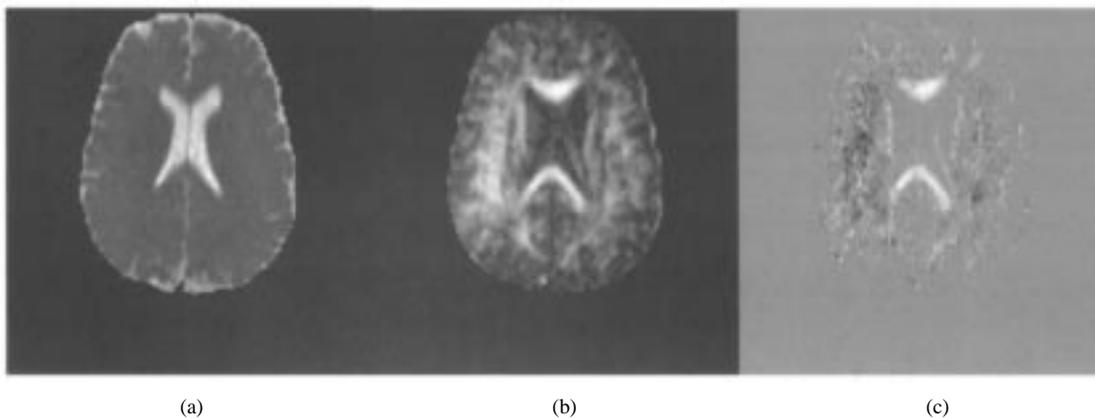

Fig. 2. (a) $\mathrm{Trace}(D)$, (b) RA, and (c) skewness maps from the one axial slice of a DT MR image acquired from the brain of a healthy, human volunteer. The $\mathrm{Trace}(D)$ map is hyperintense in regions of CSF where the total amount of diffusion is largest. The RA map is hyperintense in regions of white matter. The skewness map is hyperintense in regions of *prolate* anisotropic diffusion and hypointense in regions of *oblate* diffusion, both of which occur in white matter regions.

are independent of orientation, [16], [17]. $\mathrm{Trace}(D)$ is the sum of the eigenvalues and describes the size of the DT; a related index, which we will use later, is $\langle D \rangle$, which is the average of the three eigenvalues, equal to $\mathrm{Trace}(D)/3$.

The anisotropy is a measure of the variance of the eigenvalues, which measures the degree of directional bias in the diffusion profile. Many measures of anisotropy have been proposed and a review is provided in [15]. Here, we use the relative anisotropy (RA), which is given by $(D - \langle D \rangle I) : (D - \langle D \rangle I)/\langle D \rangle^2$, [16], where ":" denotes the tensor scalar product and I is the identity matrix. DTs with high anisotropy can take on a range of shapes. In general, their ellipsoids lie between two extreme cases: *prolate*, or cigar-shaped, ellipsoids, which have one large eigenvalue and two small ones ($\lambda_1 \gg \lambda_2 \approx \lambda_3$), and *oblate*, or pancake-shaped, ellipsoids, which have two large eigenvalues and one small one ($\lambda_1 \approx \lambda_2 \gg \lambda_3$). Different shaped DT ellipsoids are illustrated in Fig. 1. The skewness of the eigenvalues can be used to distinguish between these two cases—prolate DTs have positive skewness, while oblate DTs have negative skewness, [17]. Fig. 2 shows $\mathrm{Trace}(D)$, RA and skewness maps derived from an axial slice of a DT-MR image acquired from a healthy human volunteer.

III. SPATIAL TRANSFORMATIONS

In this section, we discuss the application of spatial transformations to DT-MR images. In order to preserve the integrity of an image undergoing such a transformation, each DT must be reoriented in such a way as to remain consistent with the surrounding anatomical structure of the image. We begin with the relatively straightforward problem of DT reorientation during rigid transformations. We then discuss the more complex problem of computing the appropriate reorientation under affine transformations of the image volume.

Three *reorientation strategies* are proposed, which provide estimates of the transformation that must be applied to each DT. We finish this section by showing how the methods described for affine transformations can be extended for higher order image transformations.

A. Rigid Transformations

The problem of applying transformations to tensor images is illustrated in Fig. 3, where a rigid, $45°$ rotation is applied to a single axial slice of a DT-MR image of the human brain. Fig. 3(a) shows a vector field derived from such an image slice.
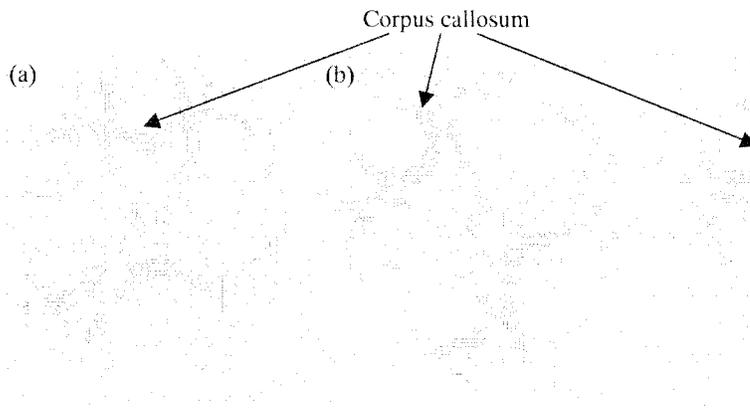
Fig. 3. Illustration of the requirement for DT reorientation. A vector field indicating the direction of $\underline{e}_1$ in anisotropic regions is shown for one particular slice of a DT-MR image. (a) Shows the original slice. (b) Shows the slice after a 45° rotation is applied, but no DT reorientation is performed. (c) Shows the slice after a 45° rotation is applied with DT reorientation.

If we apply the rotation naively to the image and simply copy the value at each pixel in the transformed image from the corresponding position in the original image, as we would for a scalar image, we obtain the result shown in Fig. 3(b). The lines within the corpus callosum indicated on the image no longer point along the pathway, as they did in the original image and as they should in the transformed image if the DTs are to be consistent with the anatomical structure. Clearly, the DTs themselves need to be rotated similarly and this can be achieved by applying the same rigid rotation matrix to each DT in the image via a similarity transform. Thus, if R is the rotation matrix representing the image transformation, each DT, D, is replaced by D', where

$$D' = RDR^T. \tag{1}$$

In Fig. 3(c), (1) has been applied at every voxel and the DT orientations are now consistent with the anatomy. Note that (1) does not change the size or shape of the DT, i.e., the eigenvalues are preserved, and only the orientation (eigenvectors) is affected.

B. Affine Transformations

When the transformation applied to the image is rigid, as in the example shown in Fig. 3, it is straightforward to determine the required reorientation of the DTs, as described above. In general, however, we need to extend the method to cope with higher order transformations. In this section, we discuss the problem of applying affine transformations and consider three potential methods for computing an appropriate reorientation for each DT.

An affine transformation consists of a vector, $\underline{t}$, which describes a rigid translation, and a linear transformation matrix, F. The image of a DT, D, under F is given by $\mathrm{F}D\mathrm{F}^T$. Unlike the case of rigid rotation, however, we cannot simply replace each D by the transformed matrix, $D' = \mathrm{F}D\mathrm{F}^T$. In general, D' has different size and shape (eigenvalues) than D. When a transformation is applied to a DT-MR image, we expect the shape of regions in the image to change, but the underlying tissue microstructure in those regions does not. Only the orientation of the tissue microstructure can change. Thus, we wish to preserve the size and shape of the DTs in the image, which reflect the properties of the tissue microstructure, but reorient them in a way consistent with the reorientation of the tissue caused by the transformation. Thus, we seek a rigid rotation matrix, R, at each point in the image, which reflects the local reorientation of the image that occurs as a consequence of F. R can then be used to reorient the DT at that point via the similarity transform given in (1).

1) Strategy 1—No Reorientation (NR): The simplest possible strategy is to ignore the problem of reorientation by applying no rotation to the DTs in the image, or equivalently to set R equal to the identity at every voxel. The value at each point in the transformed image is copied directly from the corresponding position in the untransformed image, as we would do for a scalar image.

This reorientation strategy, the NR strategy, is included here as a control or yardstick against which to measure the effectiveness of the less naïve warping strategies described below.

2) Strategy 2—Finite Strain (FS): Any nonsingular F can be decomposed into a rigid rotation component, R, and a deformation component, U, [19], where

$$F = UR. \tag{2}$$

One strategy for DT reorientation in an image undergoing transformation F, is to extract the rigid rotation component of the transformation, R in (2), and apply it to each DT in the image using (1). The rotation matrix R is straightforward to extract from F and is given by [19]

$$R = \left(FF^T\right)^{\frac{-1}{2}} F. \tag{3}$$

R needs to be computed only once and is constant over the entire image. We call this reorientation strategy the FS strategy, because the decomposition given by (2) holds for any affine transformation inducing arbitrary, but finite, strain on the material to which it is applied, [19]. When the amount of strain induced by the transformation is small, there is a simple linear approximation to the FS decomposition that can be used to speed up the computation of R, see [7], [19].
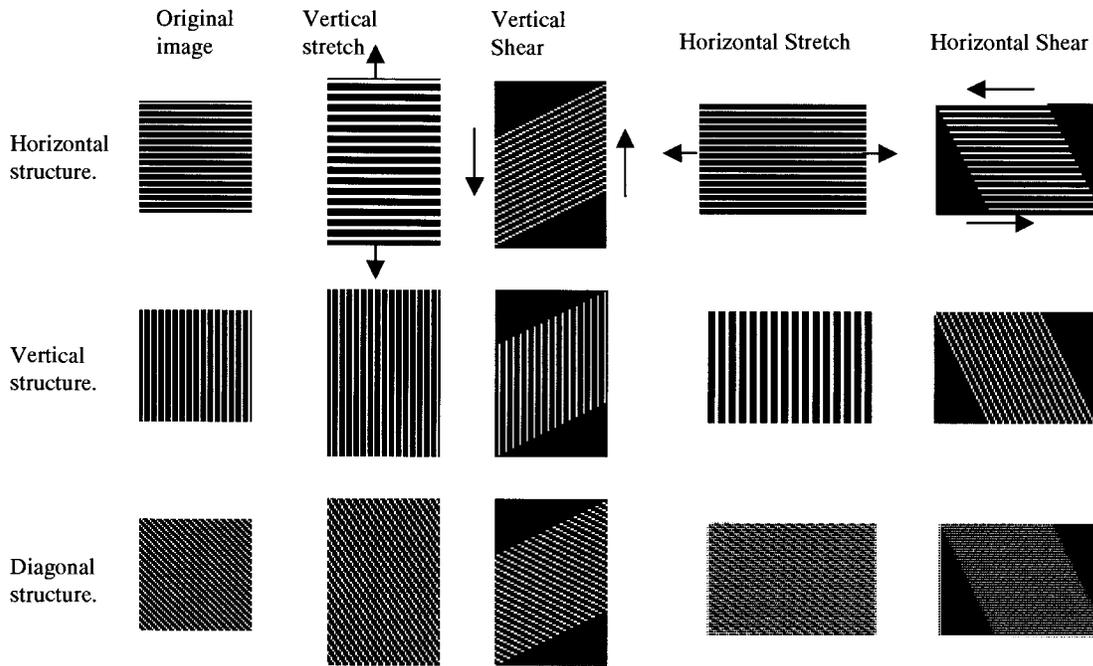
Fig. 4. Illustration of the dependency on image structure direction of the amount of reorientation due to deformative transformations. Vertical and horizontal stretch and shear are applied to three images with similar structure but different orientation. The resulting structure is similar but with different orientation and the amount of change in orientation is dependent on the original orientation of the image structure. For example, the vertical shear changes the direction of the horizontal fibers, but not the direction of the vertical ones.

3) Strategy 3—Preservation of Principal Direction: A drawback of the FS reorientation strategy is that the deformation component of the affine transformation is discarded and does not contribute to the estimated reorientation. Deformation includes transformations such as shearing and nonuniform scaling or stretching, which also affect the orientation of the underlying image (or tissue) structure, but in a more complex way than rigid rotation. In general, the reorientation that occurs as a consequence of these types of transformation depends on the original orientation of this structure. This is illustrated in Fig. 4, where three images, with strongly directional structure reflecting a portion of an image through which parallel tissue fibers are running, are stretched and sheared. The resulting change in orientation of the fibers can be seen to be dependent on their original direction.

An alternative reorientation strategy is proposed in this section, in which we directly examine the effects of the transformation on the eigenvectors of the DT at each point. We adopt the assumption that the directionality of the tissue structure corresponds to the direction of the eigenvectors of the DT [20], [21].

In prolate regions, the direction of the tissue fibers is characterized by $\underline{e}_1$. The direction of the fibers after the transformation can be found by applying the transformation directly to $\underline{e}_1$ and renormalizing to obtain a new unit vector, $\underline{n}_1$

$$\underline{n}_1 = \frac{F\underline{e}_1}{|F\underline{e}_1|}. \tag{4}$$

So, we require a rotation to apply to D that maps $\underline{e}_1$ to $\underline{n}_1$.

In oblate regions, the plane of tissue structure is characterized by $\underline{e}_1$ and $\underline{e}_2$. The affine transformation maps this plane to a new plane containing $F\underline{e}_1$ and $F\underline{e}_2$. So we require a rotation for D

that ensures that the new $\underline{e}_1$ and $\underline{e}_2$ span the same plane as $F\underline{e}_1$ and $F\underline{e}_2$.

In fact, we can take care of both these requirements in a single method, which we call the preservation of principal direction (PPD) reorientation strategy. A step-by-step outline of the PPD method is given in the Appendix. Briefly, we choose the rotation matrix, R, which maps $\underline{e}_1$ to $\underline{n}_1$ and $\underline{e}_2$ to a unit vector perpendicular to $\underline{n}_1$ in the plane spanned by $\underline{n}_1$ and $\underline{n}_2$, where $\underline{n}_2$ is defined in a similar way to $\underline{n}_1$ by replacing $\underline{e}_1$ with $\underline{e}_2$ in (4).

The PPD algorithm preserves the principal direction of the DT through the transformation as well as the plane of the first two eigenvectors and so is effective for both prolate and oblate DTs as well as for intermediate DTs for which $\lambda_1 > \lambda_2 > \lambda_3$. For prolate DTs, the mapping of $\underline{e}_2$ is unimportant, since $\underline{e}_2$ is poorly defined. Similarly, for oblate DTs, it is not essential that the new $\underline{e}_1$ lies exactly along $\underline{n}_1$ so long as it and the new $\underline{e}_1$ lie in the plane spanned by $\underline{n}_1$ and $\underline{n}_2$. For intermediate DTs, however, both parts of the transformation are important.

Note that in this scheme, the DT reorientation is not constant over the image for general affine transformations, unlike the FS strategy, and a separate R must be computed at each voxel.

C. Higher Order Transformations

Once an effective reorientation strategy has been established for affine transformations, the extension to higher order transformations is straightforward. Any image transformation, T, can be expressed as a displacement field, $\underline{u}(\underline{x})$, over the image, so that for each position $\underline{x}$ in the original image, $T(\underline{x}) = \underline{x} + \underline{u}(\underline{x})$. If T is an affine transformation, then we also have that $T(\underline{x}) = F\underline{x} + \underline{t}$, as described in Section III-B. If we differentiate these two expressions, with respect to $\underline{x}$, we can see that

$F = I + J_{\underline{u}}$, where I is the identity matrix and $J_{\underline{u}}$ is the Jacobian of the displacement field $\underline{u}$. From this, we obtain a local affine model of the more complex transformation by taking $F_{\underline{x}}$ to be equal to $I + J_{\underline{u}}$ at each point, $\underline{x}$. An appropriate reorientation matrix can then be computed separately at each $\underline{x}$, from $F_{\underline{x}}$, using the methods described earlier.

IV. SYNTHETIC DATA EXPERIMENTS

In this section, we describe a set of experiments performed using synthetic data, which test the validity of the approaches described in Section III. The synthetic data sets are 3-D DT arrays derived from a numerical phantom containing various objects that represent different anatomical features of the human brain. The phantom is rendered to a discrete 3-D DT data set by dividing up the space in which it is defined into a discrete voxel grid, assigning membership of each voxel to a particular object of the phantom and, thus choosing an appropriate DT to put in the voxel. No partial volume effects are introduced during rendering.

All of the objects in the phantom are constructed from points, lines, planes and ellipsoids, all of which are invariant under affine transformations, i.e., lines are mapped to lines, planes to planes, and ellipsoids to ellipsoids. Given such a numerical phantom and an affine transformation, there are two ways that we can obtain a data set containing a transformed version of the phantom.

- We can render the phantom to a 3-D DT array and then apply the affine transformation, using one of the reorientation strategies described in Section III.
- We can apply the transformation to the numerical phantom and then render the transformed phantom to a DT array.

Ideally, both methods should produce identical results. The second approach can be used to obtain "gold standard" data sets describing the transformed object. Data sets obtained using the first approach can be compared with the gold standard in order to test the effectiveness of the reorientation strategy.

In this section, we first review the construction of our numerical phantom and the synthetic DT data sets. Then, we describe the measures that we use to compare a transformed data set to the gold standard in order to obtain a performance index for a particular reorientation strategy. Finally, we describe the set of experiments that has been performed and present the results.

A. Synthetic Data Sets

A full description of the numerical phantom and the construction of synthetic data sets can be found in [11]. These data sets are 3-D DT arrays with an in-plane matrix size 128×128 and 37 slices. The voxel dimensions are assumed to be $1.7 \times 1.7 \times 3.5$ mm^3, which is the same as the voxels in the human data sets described in Section V.

There are four different types of tissue represented in the phantom, which are contained in separate regions of the array. Some slices through one of these data sets are shown in Fig. 5, which contains labels indicating the tissue types contained in each region. The DTs selected for different compartments within the phantom are representative of one of the four tissue
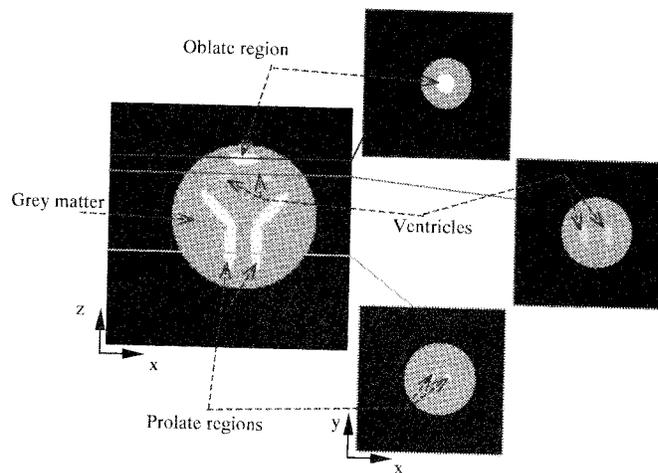

Fig. 5. One coronal (left) and three axial (right) slices through a rendered version of the numerical phantom. The upper axial slice passes through the ellipsoidal region that represents oblate white matter. The middle axial slice passes through the two ellipsoidal regions that represent the ventricles and contain large isotropic DTs. The bottom axial slice passes through the lower portion of the Y-shaped structure that contains prolate DTs and represents white matter fibers, such as the cortico-spinal tract.

types. They are drawn from statistical distributions based on the measurements made by Pierpaoli *et al.*, [14], on DT-MR images acquired from human subjects in a MR scanner. The distributions used are summarized in Table I and described in detail in [11].

In Table I, the major eigenvector, $\underline{e}_1$, in prolate regions is labeled "fixed," because it is constant rather than being drawn from a statistical distribution. The prolate region of the synthetic data sets consists of the Y shaped structure in the lower portion of the volume. This region is comprised of two pairs of cylindrical regions. Each pair of cylinders has one end point in common. The major eigenvector of every DT contained in a particular cylinder is parallel to the line joining the two endpoints of that cylinder. The minor eigenvector, $\underline{e}_3$, of DTs in the oblate region is also "fixed." The underlying tissue structure for this region is planar and a unit vector, $\underline{n}$, is defined in connection with this region that is the normal to the plane of tissue structure. Every $\underline{e}_3$ in the oblate region is set equal to $\underline{n}$.

B. Comparative Measures

In order to assess the reorientation strategies, we look at the relative orientations of DTs in corresponding locations in the transformed data set and the gold standard. In particular, we look at the angular separation of the principal axes of DTs in corresponding locations. To measure the angular separation of the major axes of two DTs, for example, we compute the cosine of the angle between the two $\underline{e}_1$s by taking the absolute value (since we are dealing with axes) of their scalar product. The cosine is inverted to give an angle between zero and $\pi/2$.

The eigenvectors are less well defined in more isotropic DTs, which causes the angular separation to be a less meaningful measure. For this reason, when the average angular separation is calculated over a particular region, we down-weight the contribution of more isotropic DTs. This is achieved by weighting the contribution to the average of each pair of DTs by the geometric

TABLE I

SUMMARY OF THE PARAMETERS USED TO SELECT DTS FOR EACH VOXEL IN THE SYNTHETIC DATA SETS. EACH EIGENVALUE AND EIGENVECTOR IS DRAWN FROM A STATISTICAL DISTRIBUTION DEPENDING ON THE REGION TO WHICH THE DT BELONGS. $N(a, b)$ INDICATES A NORMAL DISTRIBUTION WITH MEAN a AND STANDARD DEVIATION b. U_s IS A UNIFORM DISTRIBUTION ON THE SURFACE OF THE UNIT SPHERE IN THREE DIMENSIONS. $U_\perp(\underline{n})$ IS A UNIFORM DISTRIBUTION OF DIRECTIONS PERPENDICULAR TO UNIT VECTOR $\underline{n}$

Tissue type	Grey Matter	Ventricles (CSF)	White Matter	White Matter
Diffusion profile	Isotropic	Isotropic	Prolate	Oblate
λ_1	N(650, 20)(μm^2/s)	N(3200, 100)(μm^2/s)	N(1700, 150)(μm^2/s)	N(1000, 100)(μm^2/s)
λ_2	N(650, 20)(μm^2/s)	N(3200, 100)(μm^2/s)	N(300, 70)(μm^2/s)	N(1000, 100)(μm^2/s)
λ_3	N(650, 20)(μm^2/s)	N(3200, 100)(μm^2/s)	N(300, 70)(μm^2/s)	N(250, 40)(μm^2/s)
$\underline{e}_1$	U_s	U_s	fixed	$U_\perp(\underline{e}_3)$
$\underline{e}_2$	$U_\perp(\underline{e}_1)$	$U_\perp(\underline{e}_1)$	$U_\perp(\underline{e}_1)$	$\underline{e}_3 \times \underline{e}_1$
$\underline{e}_3$	$\underline{e}_1 \times \underline{e}_2$	$\underline{e}_1 \times \underline{e}_2$	$\underline{e}_1 \times \underline{e}_2$	fixed

mean of their RA values. Thus, over some set, I, of image locations the average angular separation with respect to $\underline{e}_i$, $E(\underline{e}_i)$, is given by

$$E(\underline{e}_i) = \frac{\sum_I (\sqrt{v_1 v_2} \times \cos^{-1} |\underline{e}_{i1} \cdot \underline{e}_{i2}|)}{\sum_I \sqrt{v_1 v_2}}. \qquad (5)$$

In (5), v_1 and v_2 denote the RA of the two DTs at a particular location; $\underline{e}_{i1}$ and $\underline{e}_{i2}$ denote the ith principal axes of the two DTs. The downweighting of more isotropic data is not crucial in the synthetic data experiments, where the anisotropy is quite consistent within separate regions, but is important in the human data experiments described in the next section.

Here, we compute both $E(\underline{e}_1)$, which is meaningful in regions of prolate DTs, and $E(\underline{e}_3)$, which is meaningful in regions of oblate DTs.

C. Experiments

A series of experiments was performed on a single numerical phantom. Random affine transformations were generated and applied to the phantom both before and after rendering. For each affine transformation, a gold standard of the transformed phantom is obtained, as well as a transformed data set of the original phantom using each reorientation strategy. Each transformed data set is then compared to the gold standard in order to assess the performance of the different reorientation strategies.

In order to generate random affine transformations, we parametrise the group as follows:

$$F = MS(\theta_x, \theta_y, \theta_z)R(\theta_r, \underline{r}). \qquad (6)$$

In (6), M is a nonuniform scaling matrix, given by $M = \text{diag}(m_x, m_y, m_z)$, S is a shearing transformation,

which shears by different amounts in three perpendicular directions and is given by the matrix

$$S(\theta_x, \theta_y, \theta_z) = \begin{bmatrix} 1 & \tan\theta_z & \tan\theta_y \\ 0 & 1 & \tan\theta_x \\ 0 & 0 & 1 \end{bmatrix}. \qquad (7)$$

and R is a rotation through angle θ_r about axis $\underline{r}$. Equation (6) is equivalent to the affine decomposition in (2) with MS equal to U. m_x, m_y, and m_z are drawn from $U(0.6, 1.6)$, a uniform distribution on the interval $(0.6, 1.6)$; $\underline{r}$ is drawn from a uniform distribution on the unit sphere (see [11] and [18]); $\theta_x, \theta_y, \theta_z$, and θ_r are all drawn from $U(-\pi/6, \pi/6)$.

Ten random affine transformations were generated and applied to the phantom in this way. The transformed data sets are resampled using a nearest neighbor approach (in order to avoid artifacts that might be introduced by using a less basic tensor interpolation) to select the value at each voxel. Each transformed data set was compared to the corresponding gold standard; $E(\underline{e}_1)$ and $E(\underline{e}_3)$ were computed for each reorientation strategy and averaged over those ten transformations. The averaged measures are plotted in Fig. 6. In Fig. 6, an additional column (labeled Ex. in the figures) is included for each of the individual regions, which indicates the expected value of each measure. These expected values were computed empirically by drawing pairs of DTs from the distributions corresponding to each separate region, see Table I and [11], and taking the average similarity over a large number of pairs.

D. Discussion

The plotted results show that for these data sets the PPD warping strategy is consistently the most effective and produces
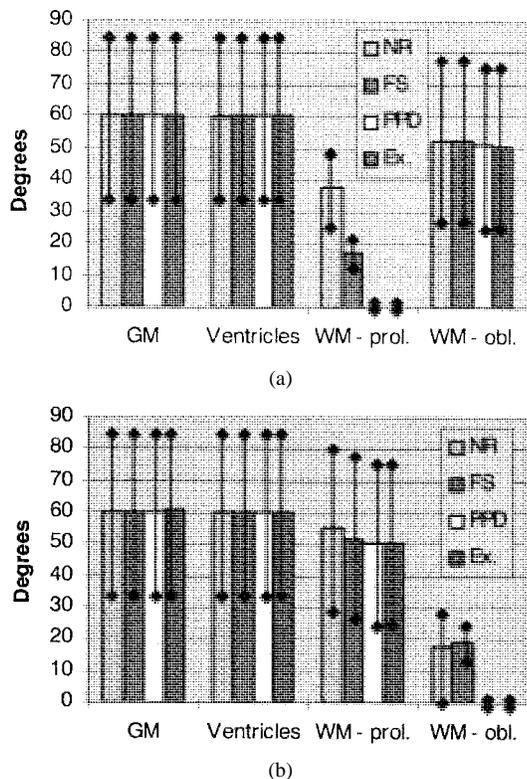
Fig. 6. Plot of (a) $E(\underline{e}_1)$ and (b) $E(\underline{e}_3)$ averaged over various regions of the numerical phantom images within ten pairs of transformed data sets. From left to right: "GM" refers to the average over "grey matter" regions, i.e., the interior of the brain volume excluding the "ventricles" and regions of oblate and prolate diffusion. "Ventricles" refers to the average over the two ellipsoidal regions of large isotropic diffusion. "WM-prol" refers to the average within the Y-shaped region of prolate diffusion. "WM-obl" refers to the average within the ellipsoidal region of oblate diffusion.

similarity measures almost identical to the ideal (expected) values in all cases.

The ideal value of $E(\underline{e}_1)$ is zero in the prolate region, as is that of $E(\underline{e}_3)$ in the oblate region. Furthermore, these measures should have zero variance in those regions, since no noise is added to $\underline{e}_1$ in the prolate region or $\underline{e}_3$ in the oblate region, which allows them to be matched perfectly. In the isotropic, CSF, and grey matter regions, the ideal value of $E(\underline{e}_1)$ is the expected angle between two uniformly randomly chosen axes in three dimensions. This value has not been computed analytically, but the empirically determined value is around $60°$. Similarly, for $E(\underline{e}_1)$ in the region of oblate DTs and $E(\underline{e}_3)$ in the regions of prolate DTs, the ideal value is equal to the expected angle between two uniformly randomly chosen axes in two dimensions. The empirically determined value for this expected value is around $50°$. Large variances of both measures are expected in isotropic regions, since $\underline{e}_1$ and $\underline{e}_3$ are uniformly distributed on the sphere. Similarly, we can expect large variance of $E(\underline{e}_3)$ in the prolate region and $E(\underline{e}_1)$ in the oblate region.

As expected, the NR strategy produces transformed data sets in which the DT orientations are inconsistent with the gold standard in anisotropic (prolate and oblate) regions. The variance of the significant measures in anisotropic regions is increased, because the different transformations cause different amounts of discrepancy in orientation. Reorientation has little effect in isotropic regions, so the comparative measures are close to the ideal values for all strategies.

The FS strategy also yields inconsistency in DT orientations. In the prolate region, it improves on the NR strategy but $E(\underline{e}_1)$ does not meet the ideal value of zero, indicating residual misalignment of DTs. Moreover, FS in fact performs worse on average than NR in the oblate region, in terms of the significant measure, $E(\underline{e}_3)$, although the variance is lower. FS can have a detrimental effect when, for example, we apply a shearing transformation along the structural planes in the oblate region. A shear perpendicular to the z axis (e.g., in (7), $\theta_x = 30°$, $\theta_y = \theta_z = 0°$), has no effect on the orientation of the tissue structure plane, which is originally aligned perpendicular to the z axis. NR is, thus, correct in this example, whereas FS rotates all DTs through approximately $16°$ about the x axis.

Only the PPD strategy manages to transform all the DTs so that their orientation is consistent with those in the gold standard. The variance of the significant measures in anisotropic regions is zero, which is consistent with the ideal value, since there is no variance in the significant eigenvector directions in these regions.

V. HUMAN DATA EXPERIMENTS

In order to test our reorientation strategies on DT-MR data acquired from a scanner, three data sets were acquired from a healthy 42-year-old male volunteer. These data sets and the experiments performed on them are described in this section.

A. Human DT-MRIs

The coordinate system of these acquisitions reflects the one used for the synthetic data sets: the x axis increases from the right to left sides of the head, the y axis increases back to front and the z axis increases bottom to top. One data set was acquired with the head positioned normally in the scanner so that the inter-hemispheric fissure is parallel to the yz plane and the anterior-posterior commissural line is almost aligned with the y axis. A second data set was acquired with the head in axially rotated position—rotated through approximately $25°$ about the z axis, and a third with the head in extended position—rotated through approximately $10°$ about the x axis.

The methodology for image acquisition and DT processing is similar to that reported in [22]. Briefly, images of 42 contiguous axial slices were acquired using a 1.5-T GE Signa Horizon EchoSpeed spectrometer. Slice thickness was 3.5 mm, field of view 220×165 mm, and in-plane matrix size 128×96, which gives rise to an in-plane voxel size of 1.7×1.7 mm. A total of 21 images were acquired per slice including three diffusion weighted images per direction with gradients applied in six different directions, [14], and three images with no diffusion weighting. The DT was computed in each voxel according to the method proposed by Basser *et al.*, [1]. Acquisition of each data set took roughly 35 min.

B. Experiments

Affine transformations of the axially rotated and extended position images were computed in order to register them with the image in normal position. These affine transformations were computed by applying Woods' AIR program, [12], [13], to RA maps from each of the three images. The two images were then
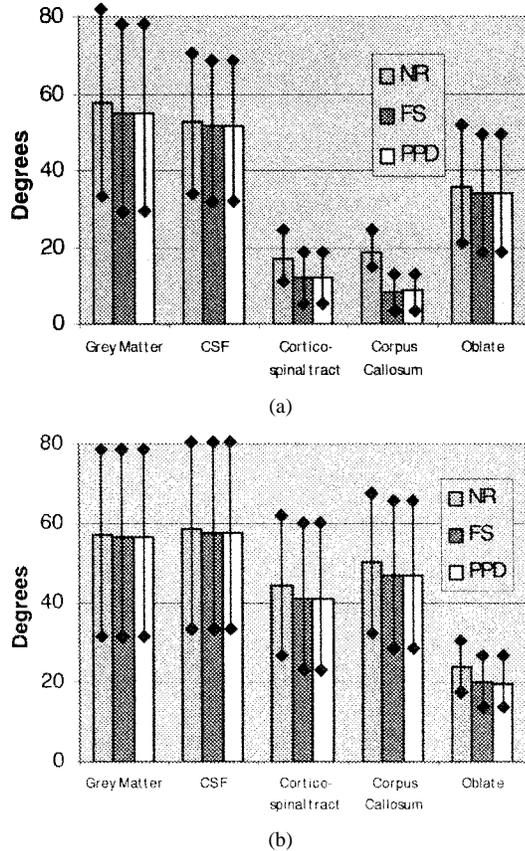
Fig. 7. Plots of (a) $E(\underline{e}_1)$ and (b) $E(\underline{e}_3)$ averaged over various anatomical regions (see text) of the human data sets after registration of the three images.

transformed into alignment with the normal position image using each reorientation strategy described in Section III. As in the synthetic data experiments, nearest neighbor resampling was used in each case.

Five regions of interest (ROIs) were outlined on the normal position image. Each is a subregion of a particular anatomic region of the image, drawn carefully so that *only* pixels from that anatomic region are contained in the set of points comprising that region of interest. The anatomic regions used are the following:

- ventricles (CSF);
- subcortical nuclei (grey matter);
- cortico-spinal tract (prolate white matter);
- corpus callosum (prolate white matter);
- subcortical white matter (oblate white matter).

The region of oblate diffusion was obtained by identifying regions on the skewness map of the normal position image, which contain predominantly DTs with negative skewness. The slice containing the largest of these regions is the slice shown in Fig. 2, where the large region of oblate diffusion can be seen on the left-hand side of the image.

Within each of the hand drawn ROIs, the DTs at corresponding positions were compared using $E(\underline{e}_1)$ and $E(\underline{e}_3)$, see Section IV-B, as in the synthetic data experiments. Averages were computed over each region for images transformed using each separate reorientation strategy. These average measures are plotted in Fig. 7.

C. Discussion

In the regions of isotropic diffusion, similar values of $E(\underline{e}_1)$ and $E(\underline{e}_3)$ are observed as were found in the synthetic data experiments. Values of both measures are slightly reduced, which suggests that there may be some coherence of DT orientation in these regions. This idea is reinforced by the fact that FS and PPD further reduce the measures in comparison to NR. It is unclear whether this directional coherence is due to structural coherence in these regions of the brain, or whether this is some artifact of the imaging process.

Both FS and PPD improve $E(\underline{e}_1)$ in the regions of prolate diffusion and the improvement is particularly marked in the corpus callosum. However, for these data sets, the results appear to be almost identical for these two strategies. Examination of the affine transformation matrices computed for registration of the images reveals that the largest deformative component in either transformation has a reorientational effect of less than $0.1°$. This means that the transformations between the images are almost perfectly rigid and so the advantages of PPD over FS are not revealed. A further consequence of this is that the variances of $E(\underline{e}_1)$ and $E(\underline{e}_3)$ are independent of reorientation strategy and are purely dependent on the variances of $\underline{e}_1$ and $\underline{e}_3$ in the original data. Notice that there is some improvement in $E(\underline{e}_3)$ in both prolate regions when reorientation is performed. DTs in these regions are not perfectly prolate, so $\lambda_1 > \lambda_2 > \lambda_3$, and there is some consistency in $\underline{e}_3$ as well as $\underline{e}_1$.

Finally, in the region of oblate diffusion, we observe some improvement in $E(\underline{e}_3)$ through use of reorientation. The improvement is less marked than in the prolate regions, because the anisotropy of these oblate DTs is much lower than those in the prolate regions. The eigenvector directions are therefore less clearly defined and so exhibit less consistency between corresponding regions in the images.

VI. Conclusion

When transformations are applied to a DT-MR image, accompanying reorientation of the data is required in order to preserve the consistency of the data with respect to the anatomic structure of the image. For rigid rotation and/or uniform scaling, the transformation matrix, normalized for scaling, can be applied directly to each DT in the image, via a similarity transform, in order to perform the necessary reorientation.

For more general affine transformations, an appropriate rotation has to be derived from the affine transformation to be applied to each DT. We have presented two possible methods for finding such a rotation and discussed their relative merits. One method, referred to as the FS reorientation strategy, decomposes the transformation into a rigid rotation and a pure deformation operations and uses the rigid rotation component for reorientation. The other, called the PPD is a reorientation strategy that takes into account and compensates for the additional reorientation caused by image deformation.

We have shown that, over the synthetic data sets described in Section IV, the latter, PPD method is an effective reorientation strategy. However, there are problems associated with FS, because it does not allow the amount of reorientation to vary over

the image. The reorientation strategies were also tested within an intra-subject human study. In this study, however, the transformation between different images is very nearly rigid and so, although both our methods are shown to be effective, the benefits of the PPD strategy, which is designed to accommodate the reorientational effects of the nonrigid part of the transformation, are not observed. However, these experiments verify that it is essential to employ a reorientation strategy, such as those proposed, in order to ensure the anatomical correctness of the transformed image.

For other tasks, such as registration of images from different subjects or deformable tissue, such as the heart or the spinal cord, higher order transformations are necessary. In such situations, we do not expect the FS strategy to be effective and PPD, or some similar technique, must be used. In addition to affine transformations, we have shown how the same methods can be used in conjunction with higher order transformations, such as those obtained from, for example, an elastic or nonlinear polynomial registration algorithm. Experiments to validate the use of the PPD method within such transformations are the focus of current investigations.

Alternatives to the PPD strategy can be envisaged, in particular, we note the connection between the estimation of appropriate reorientation transformations and the weighted orthogonal Procrustes problem, [23]. There are existing solutions to this problem, which might be used to compute reorientation matrices, but it is not clear how the relative importance of the axes should be weighted in noisy DT-MR data. Such possibilities will be investigated in future work. We note that our FS method is equivalent to the standard, singular value decomposition solution to the (unweighted) orthogonal Procrustes problem.

APPENDIX

Given a linear transformation matrix, F, and a DT, D, the PPD method proceeds as follows:

- Compute unit eigenvectors, $\underline{e}_1, \underline{e}_2, \underline{e}_3$, of D.
- Compute unit vectors $\underline{n}_1$ and $\underline{n}_2$ in the directions of $F\underline{e}_1$ and $F\underline{e}_2$.
- Compute a rotation matrix, R_1, that maps $\underline{e}_1$ onto $\underline{n}_1$. The axis, $\underline{r}$, and angle, θ, of this rotation are obtained from the vector and scalar products of $\underline{e}_1$ and $\underline{n}_1$.

A secondary rotation, about $\underline{n}_1$, is required to map $\underline{e}_2$ from its position *after* the first rotation, $R_1\underline{e}_2$, to the $\underline{n}_1$–$\underline{n}_2$ plane.

- Find the projection, $P(\underline{n}_2)$, of $\underline{n}_2$ onto a plane perpendicular to $R_1\underline{e}_1$—the new major eigenvector, $\underline{n}_1$. Note that:
  a) $R_1\underline{e}_2$ already lies in this plane, since $\underline{e}_1$ and $\underline{e}_2$ are orthogonal.
  b) $P(\underline{n}_2) = \underline{n}_2 - (\underline{n}_2 \cdot \underline{n}_1) \cdot \underline{n}_1$, lies in the plane spanned by $\underline{n}_1$ and $\underline{n}_2$.
- Compute a second rotation, R_2, that rotates $R_1\underline{e}_2$ onto a unit vector in the direction of $P(\underline{n}_2)$. The axis of this rotation is $R_1\underline{e}_1$ and the angle is obtained from the dot product of $R_1\underline{e}_2$ with $P(\underline{n}_2)/|P(\underline{n}_2)|$.
- Set $R = R_2R_1$ and reorient D using (1).

REFERENCES

[1] P. J. Basser, J. Mattiello, and D. Le Bihan, "Estimation of the effective self-diffusion tensor from the NMR spin echo," *J. Magn. Reson., B*, vol. 103, pp. 247–254, 1994.

[2] ——, "MR diffusion tensor spectroscopy and imaging," *Biophys. J.*, vol. 66, pp. 259–267, 1994.

[3] J. Crank, *The Mathematics of Diffusion*, Oxford, U.K.: Oxford Univ. Press, 1975.

[4] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.*, vol. 2, pp. 1–31, 1998.

[5] H. Lester and S. R. Arridge, "A survey of hierarchical nonlinear medical image registration," *Pattern Recogn.*, vol. 32, pp. 129–149, 1999.

[6] D. C. Alexander and J. C. Gee, "Elastic matching of diffusion tensor images," *Comput. Vision Image Understanding—Special Issue Anal. Volumetric Images*, vol. 7, no. 2, pp. 233–250, 2000.

[7] D. C. Alexander, J. C. Gee, and R. K. Bajcsy, "Strategies for data reorientation during nonrigid warps of diffusion tensor images," in *Proc. MICCAI'99*, 1999, pp. 463–472.

[8] R. Bajcsy and S. Kovacic, "Multi-resolution elastic matching," *Comput. Vision Graph. Image Processing*, vol. 46, pp. 1–21, 1989.

[9] J. C. Gee and D. R. Haynor, "Numerical methods for high dimensional warps," in *Brain Warping*, A. W. Toga, Ed. New York: Academic, 1998.

[10] J. C. Gee and R. K. Bajcsy, "Elastic matching: Continuum mechanical and probabilistic analysis," in *Brain Warping*, A. W. Toga, Ed. New York: Academic, 1998.

[11] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee, "Techniques for spatial normalization of diffusion tensor images," in *Proc. SPIE Medical Imaging, 2000*, vol. 3979, K. M. Hanson, Ed., 2000, pp. 470–481.

[12] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta, "Automated image registration: I general methods and intra-subject intra-modality validation," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 141–154, 1998.

[13] R. P. Woods, S. T. Grafton, J. D. G. Watson, N. L. Sicotte, and J. C. Mazziotta, "Automated image registration: II inter-subject validation of linear and nonlinear models," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 155–165, 1998.

[14] C. Pierpaoli, P. Jezzard, P. J. Basser, A. Barnett, and G. Di Chiro, "Diffusion tensor MR imaging of the human brain," *Radiology*, vol. 201, no. 3, pp. 637–648, 1996.

[15] C. Pierpaoli and P. J. Basser, "Toward a quantitative assessment of diffusion anisotropy," *Magn. Reson. Med.*, vol. 36, pp. 893–906, 1996.

[16] P. J. Basser and C. Pierpaoli, "Microstructural and physiological features of tissues elucidated by quantitative diffusion tensor MRI," *J. Magn. Reson., B*, vol. 111, pp. 209–219, 1996.

[17] T. E. Conturo, R. C. McKinstry, E. Akbudak, and B. H. Robinson, "Encoding of anisotropic diffusion with tetrahedral gradients: A general mathematical diffusion formalism and experimental results," *Magn. Reson. Med.*, vol. 35, pp. 399–412, 1996.

[18] N. I. Fisher, T. Lewis, and E. J. J. Embleton, *Statistical Analysis of Spherical Data*, Cambridge, U.K.: Cambridge Univ. Press, 1987.

[19] L. E. Malvern, *Introduction to the Mechanics of a Continuous Medium*. Englewood Cliffs, NJ: Prentice-Hall, 1969.

[20] P. J. Basser, "New histological and physical stains derived from diffusion tensor MR imaging," *Ann. NY Acad. Sci.*, vol. 820, pp. 123–128, 1997.

[21] E. W. Hsu, A. L. Muzicant, S. A. Matulevicius, R. C. Penland, and C. S. Henriquez, "Magnetic resonance myocardial fiber-orientation mapping with direct histological correlation," *Amer. J. Physiol.*, vol. 274, pp. H1627–H1634, 1998.

[22] A. Virta, A. L. Barnett, and C. Pierpaoli, "Visualizing and characterizing white matter fiber structure and architecture in the human pyramidal tract using diffusion tensor MRI," *Magn. Reson. Imag.*, vol. 17, no. 8, pp. 1121–1133, 1999.

[23] G. H. Golub and C. F. van Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins Univ. Press, 1996.