Seamless Warping of Diffusion Tensor Fields

Dongrong Xu*, Xuejun Hao, Ravi Bansal, Kerstin J. Plessen, and Bradley S. Peterson

Abstract—To warp diffusion tensor fields accurately, tensors must be reoriented in the space to which the tensors are warped based on both the local deformation field and the orientation of the underlying fibers in the original image. Existing algorithms for warping tensors typically use forward mapping deformations in an attempt to ensure that the local deformations in the warped image remains true to the orientation of the underlying fibers; forward mapping, however, can also create "seams" or gaps and consequently artifacts in the warped image by failing to define accurately the voxels in the template space where the magnitude of the deformation is large (e.g., $|\text{Jacobian}| > 1$). Backward mapping, in contrast, defines voxels in the template space by mapping them back to locations in the original imaging space. Backward mapping allows every voxel in the template space to be defined without the creation of seams, including voxels in which the deformation is extensive. Backward mapping, however, cannot reorient tensors in the template space because information about the directional orientation of fiber tracts is contained in the original, unwarped imaging space only, and backward mapping alone cannot transfer that information to the template space. To combine the advantages of forward and backward mapping, we propose a novel method for the spatial normalization of diffusion tensor (DT) fields that uses a bijection (a bidirectional mapping with one-to-one correspondences between image spaces) to warp DT datasets seamlessly from one imaging space to another. Once the bijection has been achieved and tensors have been correctly relocated to the template space, we can appropriately reorient tensors in the template space using a warping method based on Procrustean estimation.

Index Terms—Bijection, diffusion tensor image (DTI), DTI warping, procrustean estimation.

I. INTRODUCTION

DIFFUSION tensor imaging (DTI) quantifies the diffusion of water molecules in living tissues. In traditional anatomical magnetic resonance (MR) images of the human brain, both white matter and gray matter appear as homogeneous, internally undifferentiated tissues [1]–[3]. However, because water molecules diffuse preferentially along nerve fibers, fiber tracts can be reconstructed using the diffusion tensors (DTs) estimated in DT datasets. Thus, DTI can be used as a tool for mapping anatomical connectivity within the human brain.

To compare how brains differ across individuals or groups of individuals, brain images from any magnetic resonance imaging (MRI) modality are first coregistered into a common coordinate (or "template") space, a process termed "spatial normalization." Normalizing DT datasets, however, is particularly complex. A tensor geometrically represents a probabilistic measure of the local diffusion properties of water molecules, and each voxel of a DT dataset contains a second-order tensor, expressed as a 3×3 symmetric mathematical matrix, that is positive definite and varies in orientation and shape across the brain. Thus, spatial normalization of DT datasets requires both that voxels be displaced to the correct location in the template space (typically a canonical brain to which the brains of individuals in a study population are coregistered on a voxel-by-voxel basis) and that the directional orientation of tensors within the template space be adjusted to reflect accurately the modified orientation in the template space of the corresponding fibers in the original image. In addition, the general shape of each tensor must be maintained throughout the warping process to avoid erroneously altering local measurements of diffusivity.

Several methods have been developed to address these complexities in the warping and spatial normalization of DT datasets. The finite strain (FS) [4] method calculates the QR-decomposition (the orthogonal matrix triangularization) [5], [6] of a locally linear approximation of the displacement field while maintaining the rotational component that is needed to reorient the tensor. Another method, preservation of principal direction (PPD) [4], estimates the reorientation of the tensor at a particular voxel by examining the spatial transformation of the tensor along one or two of its principal eigenvectors. The multiple channel normalization method (MCN) [7], a variant of the PPD method, simultaneously uses all components of the tensor while successively updating tensor orientation in the registration process. Another approach, tensor normalization [8], employs the AIR coregistration package [9], [10] together with the PPD tensor reorientation strategy [4]. In addition, a technique for the nonrigid registration of 3-D tensors [11] reorients tensors using a Jacobian matrix extracted from the local deformation field (DF). Finally, a unique normalization algorithm for DT datasets based on Procrustean Estimation determines a reorientation matrix statistically for each tensor by considering the fiber orientation and DF in each voxel within a small neighborhood around it [12].

All of these methods for spatially normalizing DT datasets across individuals are based on a forward mapping strategy, in which point-to-point correspondences across subjects are defined at each grid point (the center of a voxel, where the tensor is
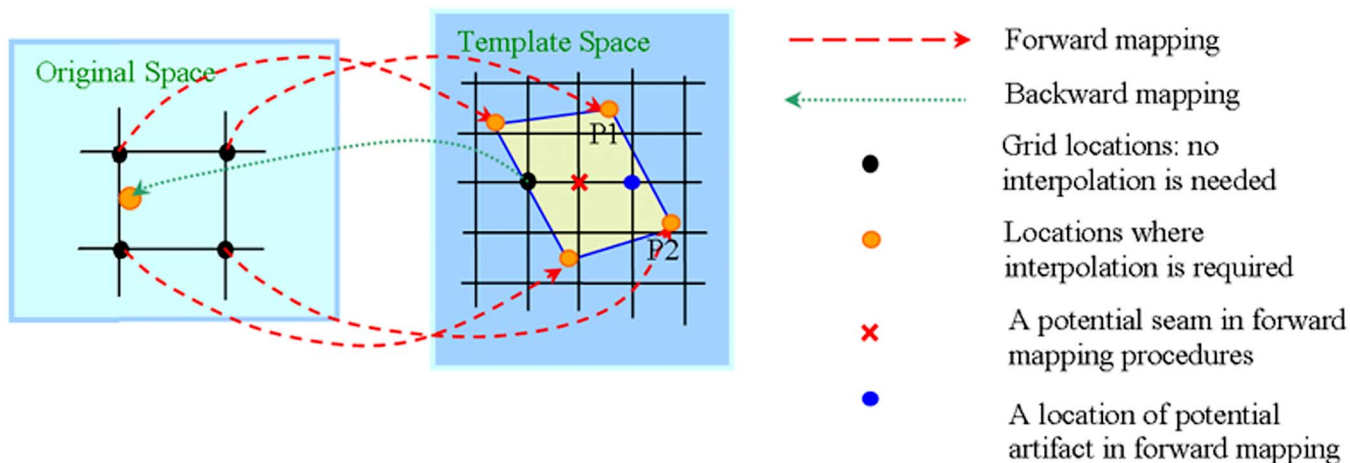
Fig. 1. Forward and backward mappings. In forward mapping procedures, seams (depicted in the template space with a red "×") will occur when pixels in the original imaging space (black dots) are distributed only to their immediately neighboring grid points in the template space (orange dots). Consequently, artifacts (an example is represented by a blue dot in the figure) will be introduced at locations of a partial seam (e.g., location P1 generates a seam at the grid point of the blue dot but location P2 does not). In contrast, if backward mapping is used, every discrete location (voxel) in the template space will receive its corresponding value from a point in the original imaging space. However, this description of backward mapping applies only to scalar images. In DTI, backward mapping alone is insufficient for warping of tensor maps because of the additional need to reorient the tensors (Fig. 2).

located) by mapping locations from the original imaging space to locations in a template space, usually no longer at grid points (Fig. 1). However, forward mapping can generate seams in the tensor dataset—i.e., areas in the template space where tensor data are absent because the DF that is used to warp the original to the template space has enlarged a region locally within the original imaging space. Such seams can significantly impair the quality of DTI data by introducing artifacts, as we will demonstrate. Backward mapping obviates the problem of introducing seams in the template space by mapping each grid point in the template space back to a point in the original imaging space, thereby defining every grid point in the template space using information interpolated from the original imaging space. This backward mapping is complete and by definition will introduce no seams into the warped dataset. If the same mapping could thus be used to warp a DT dataset in the original imaging space into the template space, seams would be eliminated from the template space. Backward mapping, however, merely defines correspondences between points in the template space with those in the original space; it cannot determine how tensors in the original space should be reoriented in template space, because that reorientation requires information about the directional orientation of tensors in the DT dataset existing within the original imaging space only. By definition, backward mapping cannot transfer that information from a DT dataset in the original imaging space to the template before the normalization procedure is completed. Thus, any procedure that uses backward mapping for warping or normalizing DT datasets must include an additional procedure that adjusts the orientation of tensors in the template space based on the information about tensor orientation in the original imaging space.

We propose a method of spatial normalization that seamlessly warps DT datasets using a *bijection*, a combination of forward and backward mappings in which the correspondence between locations in the original imaging space and the template space is one-to-one. In combination with seamless warping, our method

also uses Procrustean estimation to reorient tensors appropriately in the template space based on local deformation and orientation of the corresponding tensor in the original imaging space. We demonstrate the effectiveness of our procedure in four experiments using both simulated and *in vivo* data.

II. BACKGROUND

A. Forward and Backward Mapping in DTI Datasets

In any coregistration of images, correspondences are established between the pixels of an original imaging space and those of a template imaging space, typically using techniques of either forward or backward mapping. Backward mapping is usually the standard operation in coregistering images, because of the significant and inherent limitations of forward mapping.

Forward mapping requires that values in the original imaging space are interpolated to be distributed to grid points in the template space; this distribution may be problematic, however, when the original and template images are not of similar resolution or scale, or when the deformation across the imaging spaces varies nonlinearly. In these instances, seams will occur if the value of the relocated voxel from the original imaging space is distributed only to immediately neighboring voxels in the template space, leaving some voxels without data (Fig. 1).

Seams can be avoided in forward mapping procedures only through the application of mapping algorithms that are complex, time-consuming, and relatively difficult to implement. For example, a mapping algorithm can remedy seams by interpolating to every voxel covered by a deformed voxel in the target space the value of this voxel carried from the original imaging space. However, such mapping algorithms may fail to produce smooth transitions between neighboring voxels and potentially may create artifacts. Forward mapping is currently only used in real-time applications when many texels (small regions of adjacent pixels or voxels in an image) are likely to map to a single pixel or voxel of the template space, as is the case, for
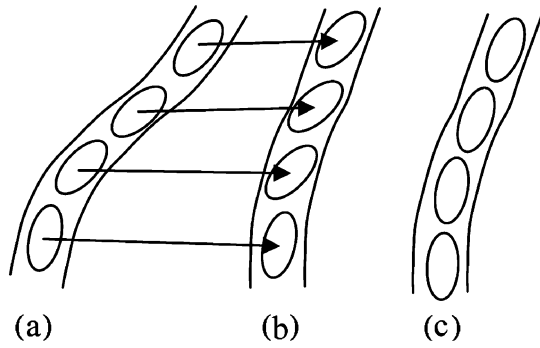
Fig. 2. Reorientation of a tensor field during warping. (a) Original tissue. (b) Incorrect warping of tensors (they are merely relocated but not reoriented). (c) Correct warping of the tensors (they are also appropriately reoriented along the direction of the nerve fiber).
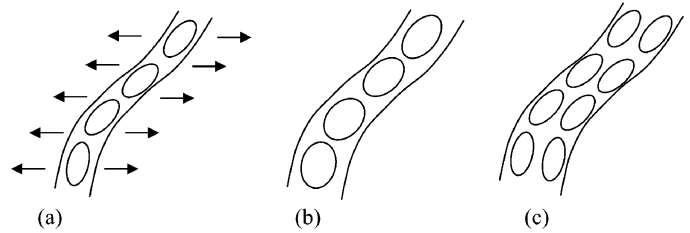

Fig. 3. Effect of spatial transformation on tensors. This figure illustrates the importance of removing the scaling component of a displacement field during DTI warping in order to preserve the shape of the tensor. (a) Original tissue. (b) Incorrect warping, in which the overall shape of the tensors is altered because of the presence of the scaling component within the displacement field. (c) Correct warping, in which the shape of the tensors has been preserved.

example, when playing digital video files on a computer screen [13]. Splatting[1] can compensate for seams in this type of image that are caused by mismatches in sampling densities; splatting, however, requires so-called "footprint"[2] calculations [14], which are time-demanding and labor-intensive. In contrast, backward mapping generates warped images by finding a value in the original imaging space for each voxel (i.e., at each grid point) of the template space, thereby avoiding the creation of seams when the original image is warped to the template. Thus, backward mapping has become the preferred method for image coregistration and for almost all types of deformation of scalar images (e.g., warping, morphing, and texture mapping).

Unlike scalar images, however, DT datasets contain in their tensors high-dimensional information about the diffusivity of water. Thus, the warping of DT datasets requires not only determining the correct correspondence between voxels in the original and template imaging spaces, but also the avoiding of seams and the geometrically appropriate reorientation of tensors at each voxel in the template space, a reorientation that is based on the geometric orientation of tensors locally within the original imaging space (Fig. 2). Reorienting tensors in this way will permit more accurate and more valid fiber tracking within the warped image in the template space. Backward mapping allows for the seamless displacement of voxels, but it does not in itself provide a means for reorienting tensors appropriately in the template space, as we discuss in more detail below. To meet the requirements for both seamless warping and the appropriate reorientation of tensors, we therefore combine backward and forward mapping to create a bijection, ensuring that the correspondence between voxels across the images is constrained to be one-to-one in both directions. Thus, the backward mapping aspect of the bijection provides a seamless displacement of tensors, whereas its forward mapping aspect provides information necessary for an accurate reorientation of tensors in the tem-

plate space according to their orientation in the original imaging space.

B. Reorienting Tensors

The correct reorientation of tensors in warped DT datasets, which is crucially important to the accurate reconstruction of fiber tracts in the template space, must take into account the orientation of tensors in the original image together with the local deformation that is induced through the warping procedure. When normalizing scalar images, the scalar value for each voxel of the image (e.g., a gray-scale intensity) is simply transferred to the template space according to a displacement vector, thereby forming the warped image in the template space. This approach, however, cannot be used to warp tensor fields (at least, not without considerable modification) because the application of displacement vectors alone will erroneously alter the position of tensors relative to one another within the warped image [Fig. 2(a)], with individual tensors no longer reflecting the correct geometric orientations in local neighborhoods within fiber tracts [Fig. 2(b)]. Instead, tensors should ideally continue to be reoriented along the direction of the fiber after warping within the template space [Fig. 2(c)].

In addition, warping tensors according to the spatial transformation that is calculated between the original and template spaces will also produce inaccurate results by altering the shape of individual tensors, thereby also altering measurements of local diffusivity and distorting inferences concerning tissue microstructure that are based on those measures of local diffusivity. Consider a simple case of spatial transformation in which a zoom-in scaling effect visually widens a reconstructed fiber bundle [Fig. 3(a)]. The tensors will be incorrectly stretched by direct application of a spatial transformation alone to the image [Fig. 3(b)], when in fact the actual width of the fiber bundle (and consequently the local diffusion characteristics and the tensor's shape) should not change [Fig. 3(c)].

Moreover, shearing effects during the warping procedure will introduce a rotational component to a tensor's warping [4], [15]. This rotational component will appropriately tilt tensors that are not originally parallel to the applied force (top two rows of tensors in Fig. 4); however, it will leave unchanged those tensors that are originally aligned parallel to the shearing force (lower two rows of tensors in Fig. 4). Thus, any strategy used to reorient tensor fields must account for both the orientation of underlying fibers in the original imaging space as well as for the effects of

---

[1]Splatting: A method of image rendering in which the final image is generated by computing for each voxel in the volume dataset its contribution to the final image. The algorithm works by virtually "throwing" the voxels onto the imaging plane or volume. In this process, every voxel in the object space leaves a "footprint" in the imaging space that will represent the object.

[2]Footprint: In a splatting procedure, the footprint defines a scope within which the effect of a thrown voxel contributes to the resulting image.
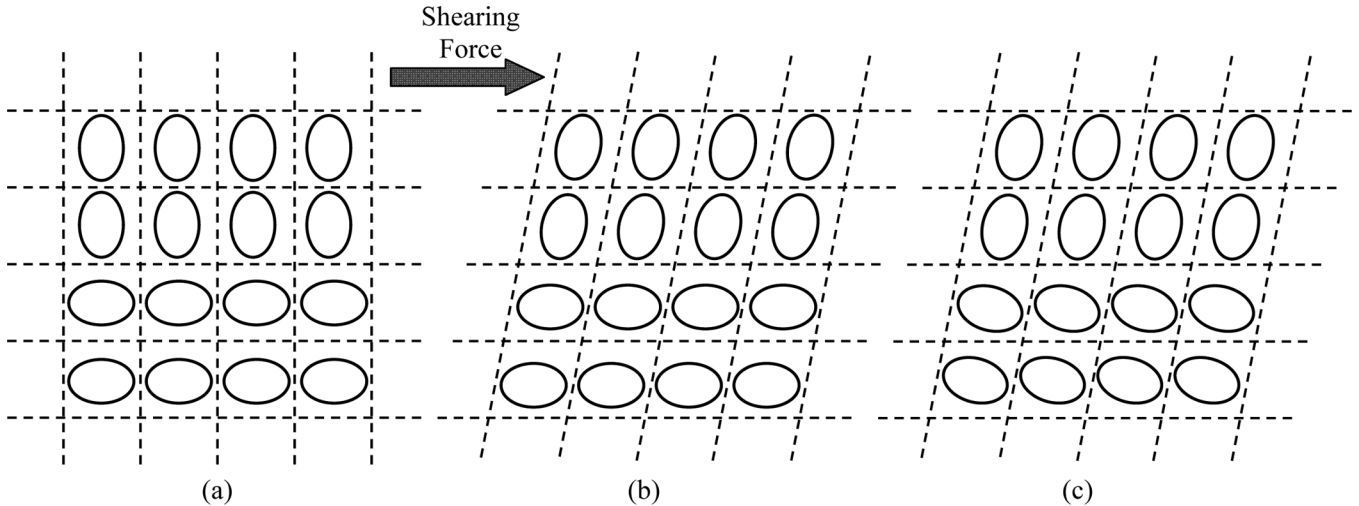
Fig. 4.   Differing tensor orientations prescribe differing rotational effects within the same spatial transformation. Tensor reorientation depends on knowledge both of the DF that imposes the spatial transformation and the tensor's orientation in the original imaging space. Shearing force should introduce a rotation that will affect only tensors that are not parallel to the direction of the shearing force (tensors in the top two rows in this figure). However, applying the rotational component to all tensors directly, without considering a tensor's original orientation, will result in tensors being reoriented incorrectly, including those tensors that are parallel to the shearing force (tensors in the bottom two rows in this figure). (a) Tissue structure in the original imaging space. (b) Correct reorientation during horizontal shearing. (c) Incorrect reorientation caused by applying the rotational component to all tensors indiscriminately.

local deformation on the orientation. Moreover, using only the local deformation, without considering the original orientation of underlying fiber tracts, may generate an inaccurate reorientation of tensors [Fig. 4(c)]. For example, in the nonrigid registration technique [11], calculating the Jacobian matrix (which is irrelevant to the orientation of local underlying fibers) and applying it to the local tensor will reorient the tensor, even if the tensor should not be reoriented.

One common solution to the problem of tensor reorientation during spatial normalization is to assume that the orientation of fibers in the original imaging space runs along the same direction as the primary eigenvector of the tensor at each voxel, as would be the case if DT datasets were noise free. This in fact is the assumption made when using the PPD method of fiber reorientation during forward mapping [4]: an affine transformation F is estimated from the DF, and then tensors are reoriented according to the rotation produced by the spatial transformation when applied to the primary eigenvector. The PPD method requires inclusion of the second eigenvector to handle reorientation of oblate tensors—i.e., tensors in which the first and second eigenvalues are similar to one another but considerably dissimilar to the third eigenvalue. Although this approach generates a reasonable first approximation to the correct reorientation of tensors, it is sensitive to errors in estimating F and to noise-induced errors in estimating the primary or secondary eigenvectors of a tensor field.

To address the problem of accurately reorienting tensors during the spatial normalization of DT datasets, we reorient tensors using a method based on Procrustean Estimation [12]. We apply it in combination with a bijection mapping, which determines a complete set of point-to-point correspondences between the original and template imaging spaces. Procrustean Estimation accounts for noise by assuming that the true orientation of underlying fiber tracts follows a statistical distribution that can be estimated from tensor measurements

in the neighborhood of the voxel under consideration. Accordingly, reorientation of the tensor is determined statistically by an estimator that identifies the optimal rotation of the tensor during spatial transformation.

III. METHODS FOR SEAMLESS WARPING OF DT DATASETS

A. Overview

The seamless warping of a DTI dataset using bijection mapping comprises two stages of image processing: 1) generation of a bijection map, which ensures that the bidirectional correspondence between the original and template imaging spaces is constrained to a one-to-one correspondence between locations in both imaging spaces and in both mapping directions and 2) the warping of one DTI dataset into the space of another, which includes both the relocation and appropriate reorientation of tensors from the original space to the template space.

The general procedures that are involved in bijection mapping are as follows. We first establish a voxel-wise one-directional correspondence between the original and template imaging spaces by coregistering the fractional anisotropy (FA) maps [16]–[18] for DT datasets in both the original and template imaging spaces (Fig. 5, Steps 1 and 2). FA maps are scalar images derived from DTI datasets that represent general indices of directionally constrained diffusion. We chose to use FA-maps to establish voxel-wise correspondences between the original and template imaging spaces because of the wide availability of algorithms for coregistering scalar-based images. Representative algorithms include HAMMER [19] and algorithms that warp scalar images based on the dynamics of fluid-flow [20], [21]. Moreover, an FA-map, unlike a conventional T_1-weighted image, always coregisters perfectly with its corresponding DT dataset, because both the FA map and DT dataset contain identical geometric distortions of brain morphology; thus a displacement field derived from a scalar FA-map can be used
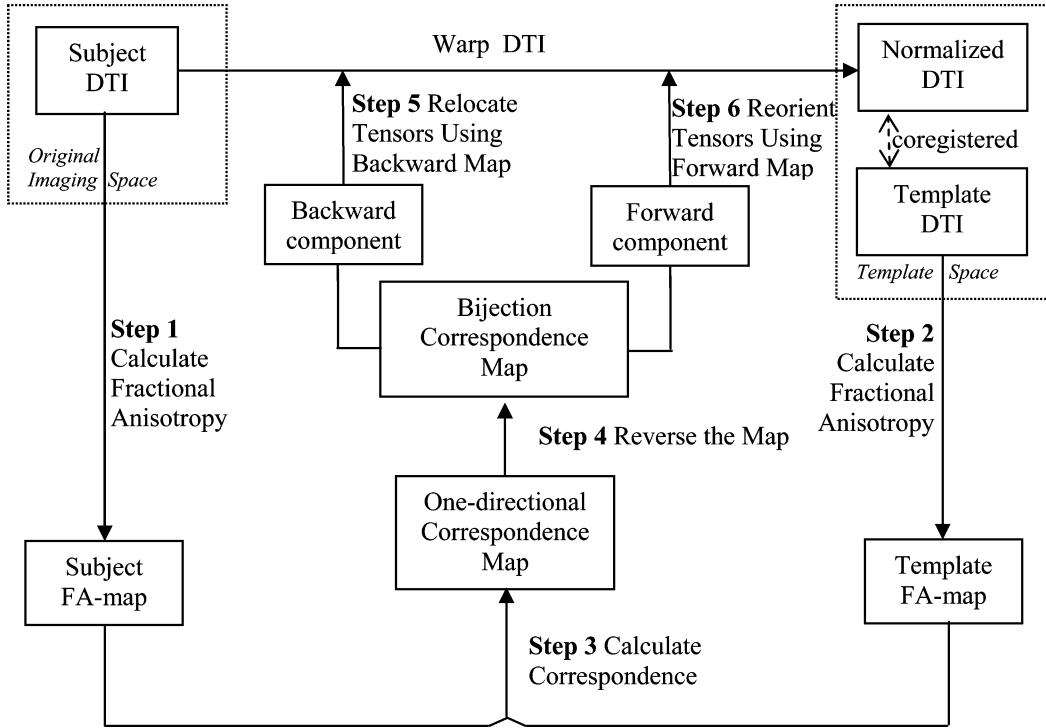
Fig. 5. Overview of the seamless warping procedure for DTI Datasets. A one-directional mapping of correspondences between an original and a template imaging space (i.e., a DF in either forward or backward direction) is generated based on the FA-maps associated with the DT images in the two imaging spaces. Bijection map of these correspondences is then generated by reversing the one-directional map. Backward component of the bijection map is subsequently used to determine the location of the tensors from the original imaging space within the template space. Forward component of the bijection map, in contrast, is used to estimate the proper directional orientation of the tensor in the template space based on its orientation in the original imaging space. Steps 1–4 constitute the first stage of our procedure (see Section III), and Steps 5–6 constitute its second stage.

to deform its corresponding DT dataset with high fidelity. The correspondence map generated by coregistering the FA-maps of the DT datasets in the original and template spaces therefore defines a displacement vector at each voxel in one imaging space that points to a corresponding voxel in the other imaging space, thus constituting the first DF of the bijection mapping (Fig. 5, Step 3). We then generate the second DF (DF_R) that is a reverse mapping of DF (Fig. 5, Step 4). A bijection is thus created in which the two deformation fields constrain each other to one-to-one correspondences consistently between voxels in each of the imaging spaces. The consistency is defined as $\mathrm{DF}_R (\mathrm{DF}(\boldsymbol{x})) = \mathbf{x}$, where $\mathbf{x}$ is a location coordinate, for example, in the original imaging space.

Finally, we seamlessly warp DT datasets using this derived bijection. The seamless warping itself entails two substeps. 1) We use the backward mapping component of the bijection mapping to find a tensor in the original space that corresponds with each voxel in the template space. We then relocate all tensors in the original space to the corresponding destinations in the template space (Fig. 5, Step 5). 2) We use Procrustean Estimation to reorient the tensors within the template space based on the known orientation information of tensors in the original imaging space and the local DF contained in the forward mapping component of the bijection mapping (Fig. 5, Step 6).

B. Algorithm for Generating a Bijection Displacement Field

A bijection displacement field combines the displacement fields of forward and backward mapping. To generate DF_R

based on DF, we need to estimate a backward displacement vector at each grid point in the template space by interpolating the correspondences defined by DF to all points on the grid. Linear and *B*-spline interpolations are well-known methods for the interpolation of deformation fields. However, linear interpolation is inappropriate for the highly nonlinear process of deforming a DT dataset from one imaging space to another. Moreover, *B*-spline interpolation is usually most suitable for interpolating data that are regularly spaced; the sample data available in the target space that we use to calculate DF_R are not regularly spaced, however, because those sample data are defined by DF, which is usually highly nonlinear. We have therefore developed an alternative method for nonlinear interpolation, which we will now describe.

Suppose we want to calculate the mapping V_b from imaging space B to imaging space A at a location b in space B. If we have a group of mapping samples V_p at locations p in the neighborhood around b and if the probability density function (pdf) of V_p is $f(V_p)$, then V_b can be computed as follows:

$$V_b = \int_{p \in \mathrm{Neighbor}(b)} f(V_p) \cdot V_p \cdot dp. \qquad (1)$$

Given that a one-directional DF $(x): A \rightarrow B$ maps locations from imaging space A to imaging space B, then for any location $\mathbf{a}$ in A, $\mathrm{DF}(\mathbf{a})$ represents a location in B that can serve as a mapping sample for locations within the neighborhood of $\mathrm{DF}(\mathbf{a})$. The DF is generally highly nonlinear; therefore, $\mathrm{DF}(\mathbf{a})$ is usually not regularly spaced. Because we want the bijection cor-

respondence to be smooth, and because a neighboring sample will resemble the true mapping better than will a remote sample, we estimate the pdf based on the Euclidean distance between the sample and the location under consideration, with the assumption that the distribution is Gaussian. Thus, for any location $\mathbf{b} \in B$, the reverse mapping can be obtained as

$$VF_R(b) = \int_{x \in A} f(\mathrm{DF}(x), b) \cdot x \cdot dx$$

$$f(a,b) = \frac{1}{\sqrt{2\pi} \cdot \sigma} e^{\frac{-\mathrm{dist}(a,b) \times \mathrm{dist}(a,b)}{2 \times \sigma \times \sigma}} \qquad (2)$$

where $\mathbf{a}, \mathbf{b}$, and $\mathbf{x}$ are locations, and $\mathrm{f}(\cdot)$ is the pdf. Notice that $\mathrm{DF}(\mathbf{x})$ and $\mathbf{b}$ are in space B.

In real digital images, we have only a discrete version of DF, in which all DF vectors defined on grid points in A map to non-grid points in B (Fig. 1); however, a reverse DF_R must be defined in B on grid points as well. Therefore, we compute the discrete version of the reverse DF_R as follows:

$$VF_R(b) = \sum_{x \in A} [f(\mathrm{DF}(x), b) \times x]$$

$$f(a,b) = \exp\left(\frac{-\mathrm{dist}(a,b) \times dist(a,b)}{2 \times \sigma \times \sigma}\right)$$
$$\div \sum_{x \in A} \exp\left(\frac{-\mathrm{dist}(\mathrm{DF}(x), b) \times \mathrm{dist}(\mathrm{DF}(x), b)}{2 \times \sigma \times \sigma}\right). \qquad (3)$$

Because DF is highly nonlinear, the neighborhood around some grid points in space B may not receive an adequate number of samples and thus lack data. We, therefore, use super-sampling (which will produce a finer image resolution, in contrast to sub-sampling) to generate closer and more influential samples. The larger the variation σ is, the greater the number of voxels included in calculating the reverse mapping field. The more voxels included in the calculation, the smoother the reverse mapping.

C. Procrustean Tensor Reorientation (PTR)

To reorient tensors properly in the template space, we first must determine the orientation of the underlying fibers in the original imaging space. We therefore approximate this orientation statistically based on the primary eigenvectors of tensors in the vicinity of the voxel of interest. The shape of this neighborhood is also estimated to follow the principal direction of the fiber tract. This procedure is expressed in pseudo-code as follows.

PTR_1. Find an image warping transformation for FA or coregistered T_1-weighted images, using an available method for warping scalar images, such as HAMMER [19], fluid-flow based nonlinear deformation [20], [21], SPM [22], [23], or AIR [24], [25]. Then, to reorient the diffusion tensor properly at any given voxel, do the following.

PTR_2. Draw random samples of the fiber direction from a neighborhood around the given location of interest. These samples are assumed to follow a pdf that is related to the underlying fiber direction; they are also assumed to be the primary eigenvectors of the diffusion tensors in the vicinity of the voxel of interest.

PTR_3. Adapt the shape of the neighborhood so that it becomes elongated along the direction of the samples drawn in Step **PTR_2**.

PTR_4. Upon convergence of the shape of the neighborhood, proceed to Step **PTR_5**; otherwise, go to Step **PTR_2**.

PTR_5. Collect all random samples from the neighborhood and find their new directions after image warping via the transformation in Step **PTR_1**.

PTR_6. Estimate the rotation matrix that minimizes the mean square error between the pairs of directions before and after image warping in Step **PTR_5**.

PTR_7. Reorient the diffusion tensor of the voxel of interest according to the estimated rotation matrix in Step **PTR_6**.

We now detail the crucial mathematical components of our procedure for tensor reorientation, including neighborhood estimation, calculation of the required rotation for proper directional reorientation, and sample weighting.

Estimating the Rotation Matrix: If we assume the absence of noise in the tensor measurements, finding a rotation matrix $\mathbf{U}$ that maps the principal direction (PD) of the tensor at a particular location to its deformed configuration by the displacement field is relatively straightforward. Let $\mathbf{x}$ be the coordinates of a given voxel. Let $\mathbf{D}$ be the corresponding tensor that is measured at that location. Assume that the orientation of the underlying fiber that passes through $\mathbf{x}$ is known and is represented by a unit vector $\mathbf{v}$. By applying the warping transformation, the vector $\mathbf{v}$ is mapped to a new vector, which is then normalized to be a vector $\mathbf{v'}$ with unit magnitude. The rotation of $\mathbf{v}$ to $\mathbf{v'}$ is represented by a unitary matrix $\mathbf{U}(\mathbf{v})$ that is then applied to $\mathbf{D}$, reorienting it to $\mathbf{U}(\mathbf{v}) \cdot \mathbf{D} \cdot \mathbf{U}^T(\mathbf{v})$. Note that the shape of the tensor does not change; rather, its eigenvectors are simply rotated by $\mathbf{U}(\mathbf{v})$.

In real imaging datasets, of course, we do not know the true orientation of the underlying fibers. Rather, the PD of a tensor is a noisy, probabilistic observation of a fiber's direction. Assume temporarily, however, that we do know the pdf of the orientation of the fiber tract in the neighborhood of a particular voxel. Then we can perform random samplings of this distribution, each time generating a vector $\mathbf{v}$. For each such vector $\mathbf{v}$ at location $\mathbf{x}$, its deformed configuration $\mathbf{v'}$ is determined by a straightforward application of the displacement field to the vector and then normalization of the resulting vector $\mathbf{F}(\mathbf{x}+k \cdot \mathbf{v}/2) - \mathbf{F}(\mathbf{x}-k \cdot \mathbf{v}/2)$, where $\mathbf{F}$ is the displacement field obtained through **Step PTR_1**, and k is a small constant that enforces local sampling. This procedure yields a number of orientations and their respective deformed configurations. The rotation matrix that best fits this set of measurements is then determined.

Next, suppose that we know the pdf of the vector $\mathbf{v}$, which is denoted by $f(\mathbf{v})$. Then we seek the matrix $\mathbf{U}$ that minimizes the following expression:

$$\varepsilon\{\|v' - U \cdot v\|^2\} = \min_U \left(\int_v f(v)\|v' - U \cdot v\|^2 dv\right) \qquad (4)$$

where $\varepsilon\{\cdot\}$ denotes the expected value of its argument. An estimator for the rotation matrix (ERM) $\mathbf{U}$ at location $\mathbf{x}$ is determined as follows.

ERM_1: Perform samplings of the pdf f$(\cdot)$, thereby generating a number of unit vectors $\boldsymbol{v}$; arrange the vectors to form the columns of a matrix $\boldsymbol{B} = (\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_n)$. (N.B. The true pdf is actually unknown. See subsection "neighborhood sampling strategy" below for details on processing the pdf and resampling).

ERM_2: Find the configuration for each sample, $\boldsymbol{v}$, after applying the displacement field, and normalize it to a unit length vector, $\boldsymbol{v}'$. Use these vectors to form the columns of another matrix $\boldsymbol{A} = (\boldsymbol{v}'_1, \boldsymbol{v}'_2, \ldots, \boldsymbol{v}'_n)$.

ERM_3: Compute the transformation matrix U so that U minimizes

$$(\|\boldsymbol{A} - \boldsymbol{U} \cdot B\|_2)^2.$$

The well-known solution to this problem is called Procrustean Estimation [5], a term originating from the Procrustes fitting procedure, which minimizes differences in size, location, and orientation across structured entities. Here, $\mathbf{U}$ can be obtained from operations on two matrices, $\mathbf{V}$ and Σ, that are components of the singular value decomposition (SVD) of $\mathbf{A} \cdot B^T$

$$\mathbf{U} = \mathbf{V} \cdot \mathbf{W}^T \tag{5}$$

$$\mathbf{A} \cdot B^T = \mathbf{V} \cdot \Sigma \cdot \mathbf{W}^T. \tag{6}$$

This expression minimizes the following:

$$\|\mathbf{A} - \mathbf{U} \cdot \mathbf{B}\|_2^2 = \|\mathbf{A}\|_2^2 + \|\mathbf{B}\|_2^2 + 2\sum_i \sigma_i(\mathbf{A} \cdot \mathbf{B}^T) \tag{7}$$

where $\sigma_i\left(\mathbf{A} \cdot \mathbf{B}^T\right)$ is the singular value of matrix $\mathbf{A} \cdot \mathbf{B}^T$. The effect of Procrustean Estimation can be interpreted as the approximate rotational transformation that minimizes the differences between point coordinate-pairs in $\mathbf{A}$ and $\mathbf{U} \cdot \mathbf{B}$. When we use normalized vectors to form $\mathbf{A}$ and $\mathbf{B}$, all samples contribute equally in Procrustean Estimation. This estimation always produces a pure rotation. Thus, in summary, if a statistical representation of the orientation of an underlying fiber tract is known through $f\left(\cdot\right)$, then the matrix $\mathbf{U}$ can be determined statistically. $\mathbf{U} = \mathbf{V} \cdot \mathbf{W}^{\mathbf{T}}$ is the matrix that best approximates the rotational component of the displacement field along the direction implied by $f\left(\cdot\right)$.

Neighborhood Sampling Strategy: Because in practice the pdf of tensor measurements at each voxel is unknown, we consider a small neighborhood $N(\mathbf{x})$ around location $\mathbf{x}$ from which we draw samples of the vector $\mathbf{v}$. We assume that if the neighborhood $N(\mathbf{x})$ is relatively small around $\mathbf{x}$, and if the direction of the underlying fiber remains approximately constant, then any variations of the samples in the direction of the primary eigenvector measured at points within $N(\mathbf{x})$ reflect the variations implied by $f\left(\cdot\right)$. As one option for generating random samples of $\mathbf{v}$ drawn from $f\left(\cdot\right)$, we could attempt to estimate $f\left(\cdot\right)$ from samples obtained from $N(\mathbf{x})$ by assuming the presence of a Gaussian or similar stochastic distribution, and then applying a random sampling procedure to generate a large number of representative samples, as described in **Step ERM_1**. This method, however, depends on an arbitrary and likely invalid assumption of the distribution of vectors in the neighborhood of $\mathbf{x}$. Random sampling that involves a large number of voxels across an image would also be excessively time-consuming.

As an alternative, we can directly use the samples in $\boldsymbol{N}(\mathbf{x})$ in **Step ERM_1**, thus eliminating the weakness of assuming a specific statistical distribution of the samples. In this case, the samples from $\boldsymbol{N}(\mathbf{x})$ effectively replace the *pdf f* $(\mathbf{v})$; (4) now becomes

$$\varepsilon\{\|v' - U \cdot v\|^2\} = \min_U \left(\int_{y \in N(x)} \|v'(y) - U \cdot v(y)\|^2 dy \right). \tag{8}$$

We have assumed that the samples in $\boldsymbol{N}(\mathbf{x})$ follow $f\left(\cdot\right)$; therefore, $\boldsymbol{N}(\mathbf{x})$ should not include heterogeneous regions or differently oriented fibers, because this would violate our assumption that $f\left(\cdot\right)$ is approximately constant within $\boldsymbol{N}(\mathbf{x})$. Because white matter fiber bundles are physically thin, elongated structures, we propose using an ellipsoidal neighborhood whose precise shape is dictated by the shape of the tensor field in the neighborhood around $\mathbf{x}$. The neighborhood for sampling (NFS) is determined according to the following iterative procedure:

NFS_1: Select a value, V, for the volume of the neighborhood from which samples are drawn.

NFS_2: Determine the ellipsoidal neighborhood, $\boldsymbol{N}(\boldsymbol{x})$, around $\boldsymbol{x}$ that has the same volume V, axis ratio, and orientation as the tensor $\boldsymbol{D}$ at $\boldsymbol{x}$.

NFS_3: Calculate the average tensor, $\bar{\mathrm{D}}$, of the tensors within $\boldsymbol{N}(\boldsymbol{x})$.

NFS_4: Keeping the same volume V, update the ellipsoidal neighborhood, $\boldsymbol{N}(\boldsymbol{x})$, at $\boldsymbol{x}$, with the same axis ratio, and orientation as those of the tensor $\bar{\mathrm{D}}$.

NFS_5: If the change in shape of the $\boldsymbol{N}(\boldsymbol{x})$ is lower than a previously specified threshold, terminate the iteration; otherwise, go to **Step NFS_3**.

The volume V of the ellipsoidal neighborhood and the resampling density are held constant during this procedure to ensure that the same number of samples of the vector $\mathbf{v}$ is used for each voxel in the image.

Weighted Procrustean Estimation: So far we have assumed that all samples drawn from $\boldsymbol{N}(\cdot)$ are weighted equally, and we have used equally weighted samples as a first approximation in the experiments reported in this paper. However, using a weighted sum (i.e., an integral) in (8) may be desirable for two reasons. First, vectors $\mathbf{v}$ that are drawn from a location farther away from $\mathbf{x}$ are more likely to belong to a slightly differing distribution $f\left(\cdot\right)$, depending on the variability of curvature in the fiber tract. Second, the ratios of the eigenvalues associated with the measured vectors $\mathbf{v}$ represent confidence in the directionality of the vectors, and vice versa. Voxels that have higher diffusion anisotropy support a more reliable estimation of the orientation of the major tensor; for example, samples drawn from thin and elongated ellipsoids should be weighted more than samples drawn from relatively shorter and rounder ellipsoids.

Extending the algorithm using weighted sums is straightforward, in that it changes (8) to

$$\min_U \left(\int_{N(x)} w^2(y)\|v'(y) - U \cdot v(y)\|^2 dy \right)$$

$$= \min_U \left(\int_{N(x)} \|w(y) \cdot v'(y) - U \cdot w(y) \cdot v(y)\|^2 dy \right)$$

$$w = \frac{\sqrt{FA}}{1 + \sqrt{d}} \tag{9}$$

where FA is fractional anisotropy and d is the distance from the sample to the location under consideration. In this case, $w \cdot \mathbf{v}$ rather than $\mathbf{v}$ forms the columns of matrix $\mathbf{B}$, and $w \cdot \mathbf{v}'$ rather than $\mathbf{v}'$ forms the columns of matrix $\mathbf{A}$. As long as the vectors of the corresponding columns in both matrices are of the same weighted length, this extended Procrustean Estimation produces a rotation $\mathbf{U}$ that operates the same as discussed previously. Longer vectors, however, will influence the estimate to a relatively greater degree.

D. Warping the DT Dataset and Reorienting the Second PD of a DT

We first use backward mapping to determine for each voxel v_T in the template space a corresponding voxel v_O in the space of the original image. However, v_O is usually not located at a grid location, and therefore to calculate tensor D_O at v_O, we either interpolate neighboring tensors or we interpolate the raw diffusion-weighted images from which the tensors are estimated. The displacement field defines where the tensor should be relocated, and our reorientation algorithm then decides how to rotate the tensor correctly, based on the tensor's original orientation using the forward mapping component of the bijection.

Until now, our discussion of tensor reorientation has focused on the first principal direction, PD_1, of a given tensor. However, reorienting the PD_1 alone may not produce a correct reorientation of the second PD (PD_2). Thus, the reorientation for PD_2 must be estimated as well. Because the three PDs of a tensor span an orthogonal space, we calculate the reorientation of the PD_2 only on a plane α that is perpendicular to the reoriented PD_1. Suppose U_1 is the rotation matrix corresponding to PD_1. We denote PD_{2p} as the projection of vector $(U_1 \cdot \text{PD}_2)$ on α, and PD_{2p}' the projection of the displaced PD_2 on α. For computing the rotation matrix U_2 corresponding to PD_2, we use PD_{2p} and PD_{2p}' as the components to form the columns of $\mathbf{B}$ and $\mathbf{A}$ in Steps **ERM_1** and **ERM_2**, respectively. Thus, U_2 rotates the tensor along the axis $(U_1 \cdot \text{PD}_1)$, in which $U = U_2 \cdot U_1$ is the reorientation matrix that we are trying to identify. In this way, both the first and second PDs, and consequently the third PD, are appropriately reoriented. This strategy is similar to that used in the PPD method previously discussed [4], except that Procrustean Estimation is employed.

IV. EXPERIMENTS AND RESULTS

We designed four experiments to assess the effectiveness of our algorithms for seamless warping. In the first experiment, we measured the accuracy of our algorithm for generating bijection displacement fields. In the second experiment, we used simulated datasets to show that our DTI warping procedure seamlessly warps tensor fields while appropriately positioning and reorienting all tensors. In addition, we showed that artifacts normally caused by seams were likewise removed. In the third experiment, we compared results from the warping of images of phantoms using our bijectional mapping algorithm with results using forward mapping alone. In the fourth experiment, we used *in vivo* datasets to assess whether our seamless DTI warping algorithm has any advantages over the DTI warping algorithms

that use only forward mapping. In this head-to-head comparison, we conclude that our algorithm fulfills the goals of seamless warping and produces a better image quality by eliminating artifacts caused by seams.

A. Experiment 1

Experimental Design: Using a collection of 54 one-directional deformation fields (dimensions $= 256 \times 256 \times 80$, resolutions $= 1 \times 1 \times 1 \text{ mm}^3$), we first constructed 54 bijections and then tested the consistency of the two-directional components (both forward and backward) of the bijection. In this experiment, we were only concerned with assessing whether our algorithm could generate bijections based on one-directional mappings. We were not concerned with how the deformation fields were generated or on what subject populations they were tested. We therefore selected the largest dataset of deformation fields that we had collected to date to provide the greatest statistical power for testing the mentioned consistency. These 54 deformation fields were generated using a nonlinear deformation algorithm based on fluid flow [20], [21], described in a previous study [26]; the entire sample contained 55 subjects, with the dataset from one healthy subject selected randomly as the template for spatial normalization. The first stage of the proposed seamless warping algorithm [i.e., (3)], was applied to generate the bijection version of the mapping fields. To measure the consistency of correspondence of these bidirectional maps, we compared the coordinates of each voxel $\mathbf{x}$ with $\text{DF}_{\text{B} \rightarrow \text{A}} (\text{DF}_{\text{A} \rightarrow \text{B}}(\mathbf{x}))$ and calculated the average errors and standard deviations of spatial correspondence based on: $\text{distance}(\mathbf{x}, \text{DF}_{\text{B} \rightarrow \text{A}}(\text{DF}_{\text{A} \rightarrow \text{B}}(\mathbf{x})))$. To simplify the procedure of generating bijection for applications in this real-world dataset, we used $\sigma = 1 \text{ mm}$ and repeated the procedure for generating bidirectional maps at 1, 2, and 3 levels of super-sampling, corresponding to the insertion of 0, 1, and, 2 samples between neighboring displacement vectors in DF at the original resolution.

Results: The average error of our bidirectional map was approximately 0.07 mm (Fig. 6). Our results show that super-sampling increases the accuracy and decreases the variability in estimating the bijection, especially when estimating at locations where the magnitude of deformations has been extreme. It is important to note, however, that although the average error and variability decreased slightly at higher levels of super-sampling, the computational cost of increasing the level of super-sampling increased exponentially. In our dataset, we found that the cost-benefit trade-off of accuracy versus computational demands was optimal at a super-sampling level of 2. Running on a SUN Solaris workstation with one 900 MHZ CPU, the generation of a bijection for a DF having dimensions of $256 \times 256 \times 80$ voxels required approximately 45 min at a super-sampling level of 1 and 185 min at a level of 2.

B. Experiment 2

Experimental Design: We used simulated datasets to demonstrate the effectiveness of our algorithm for producing seamless and correctly reoriented tensor datasets. We tested its ability to produce artifact-free results compared with DTI warpings that
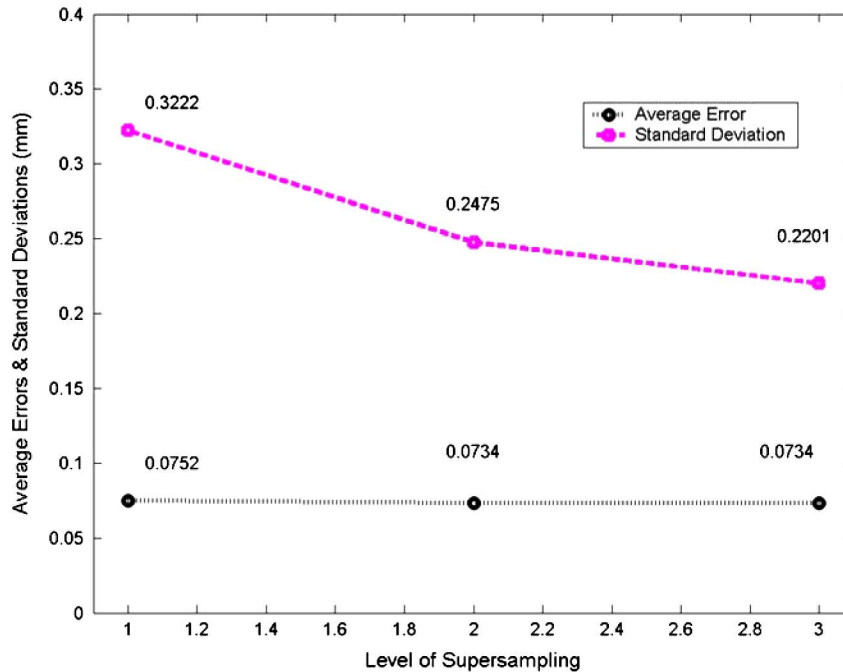
Fig. 6. Measurement of the consistency of bidirectional mapping. This figure shows the average error and standard deviation of our algorithm for generating bijections. Level of super-sampling refers to the number of new samples to be inserted between two neighboring mapping vectors in the DF at the original level of resolution.

use forward mapping. We synthesized a 3-D DT dataset containing four bundles of fibers running horizontally, vertically, or perpendicular to the field of view [Fig. 7(a)]. Tensors were defined as having eigenvalues 0.5, 0.3, 0.05, and consequently FA $= 0.6673$ [16], [17]. The first and second principal directions of the tensor field in the original space were color-coded using the common schema of mapping horizontal to red, vertical to green, and perpendicular to blue [Figs. 7(a) and (b)]. We also displayed the vector view of the principal directions of the tensor field [Fig. 7(c)]. We warped this tensor image using the displacement field of a vortex, in which the rotational displacement was least in the center of the image and progressively greater toward its perimeter [Fig. 7(d)]. We applied our algorithm of seamless warping using bijection mapping. Then we repeated these procedures using conventional DTI warping based on forward mapping.

We quantified differences in these warped datasets by comparing their signal-to-noise ratios (SNRs). We first generated and then warped the scalar FA-map of the original tensor field using the vortex DF. This warped FA-map FA_g was taken as ground truth for signal reference in the foreground and background of the dataset. To examine closely how the magnitude of the deformation affected signal-to-noise ratio (SNR), we calculated SNR in varying bands of the determinant value J of the Jacobian matrix for the DF by masking the image with differing thresholds in which J was constrained by $a \leqslant J \leqslant b$ in the DF. As is well known, J should always be positive for topology-preserved warping, which was the case in this simulation. $J < 1$ indicates local contraction; $J = 1$ indicates neither contraction nor expansion; and $J > 1$ indicates local expansion. *Results*: The direction of displacement in the vortex field obviously was not parallel to the orientation of the underlying tissue.

Therefore, virtually every tensor (except at the central point of the image) experienced a shearing force that was not parallel to its original orientation. The vortex DF should therefore have deformed each of the four fibers and appropriately reoriented their tensors. Color maps [Fig. 7(e) and (f)] demonstrated that the extracted PD_1 and PD_2 of each tensor in the warped field were appropriately reoriented using our method, in that the PD_1's followed the orientations of the deformed fibers and the PD_2's were simultaneously correctly rotated. In Fig. 7, the circumstances for Fiber 4 differed somewhat from those of the other fibers because this particular fiber received rotational forces but no shearing along its PD_1 vectors. Therefore, its PD_1 was rotated along its PD_1 and thus was not reoriented by the forces and by our algorithm, whereas its PD_2 was reoriented. In contrast, compared with the PD_1 of the DT dataset warped using forward mapping [Fig. 7(g)], the warped image contained seams in locations where the deformation was relatively extreme [Fig. 7(g) inset], even though the tensors were mostly relocated and reoriented properly. A vector view of the PD_1 of the deformed tensor field [Fig. 7(h)] shows more intuitively how the tensors were correctly repositioned and reoriented. In central portions of the image [Fig. 7(e) and (f)] where the displacement was minimal and of similar in magnitude across voxels, the structure of the warped tensor field changed only slightly from the original image. The thickness of the synthesized fibers changed appropriately according to the degree of local stretching that was imposed by the displacement field.

Whereas the SNR of the warped image generated with our bijectional map for simulated data was perfect ($\mathrm{SNR} = +\infty$), the SNR based on forward mapping alone dropped rapidly from 87.7 ($0.0 \leqslant J \leqslant 1.0$) to 3.0 (for $1.0 < J \leqslant 2.0$) and then monotonically to 0.6 (for $13.0 \leqslant J \leqslant 14.0$) (Fig. 8). The SNR
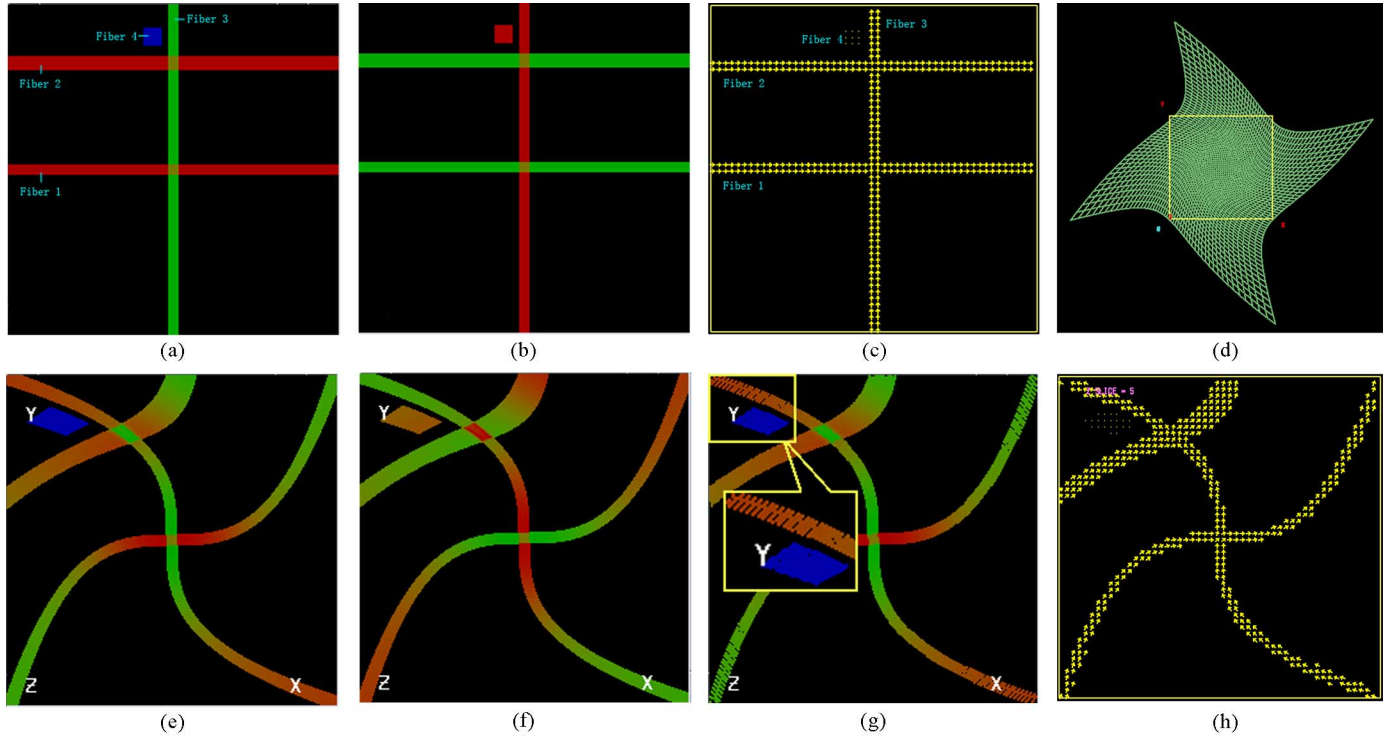
Fig. 7. Improvement in fiber reorientation and removal of seams in a simulated dataset. (a) and (b) Color maps of the first two PDs, PD1 and PD2, of tensors in the simulated dataset. Red = horizontal orientation; green = vertical orientation; blue = orientation perpendicular to the plane of viewing. (c) Vector view of the PD1 of the tensors. (d) Depicts the vortex-like displacement field that was applied to the simulated dataset; the yellow rectangle is a reference that represents the size of the simulated, original imaging space. (e) and (f) Color maps of PD1 and PD2 of the warped dataset in the template space that was constructed using our bijection seamless warping. Every voxel in the warped image is properly defined. (g) Counterpart of (e) that was produced by warping the DTI dataset using forward mapping. Here, seams appear (see inset) at locations of extreme local deformation. (h)Vector view of PD1 of the deformed DT field. (X, Y, and Z denote the coordinate system).

showed a significant transition from $J < 1$ to $J > 1$ and then kept declining as J grew larger. Because we used in our simulation a unified tensor with $FA = 0.6673$ and tested SNR on the resulting FA maps, the decrease in SNR indicated a decrease in average FA, suggesting the degradation of tensors when warping uses a forward mapping strategy.

C. Experiment 3

Experimental Design: We constructed a phantom to measure how accurately fiber tracts can be recovered after normalization using our seamless warping algorithm compared with a conventional forward-mapping one. We synthesized five datasets, each of which included a realistic 3-D tensor field and a real 3-D T_1-weighted image. The 3-D T_1-weighted images are only used as geometrical references for their realistic anatomy in the process of synthesizing the corresponding DT datasets. We first randomly selected from our larger morphometric database one T_1-weighted image that served as a template. We manually drew a number of curves in this image, simulating two major fiber bundles crossing through the corpus callosum to the opposite hemisphere. We then randomly selected from our database five additional brains with varying morphology. Using the elastic warping algorithm STAR [27], guided in part by a number of manually drawn sulcal curves [28], we warped the template to five different configurations reflecting the cortical and subcortical anatomy of the five brains. We also spa-

tially transformed the curves representing fibers of the corpus callosum using the same transformation. (We transformed the mathematical expression of the graphical curves but not the image voxels underlined by the curves, so that no seams would be generated in this preparatory step.) We thereby generated five synthetic brains with differing realistic anatomies and fiber tracts. These synthetic brains differed from one another in general shape as well as in local features, such as the morphology of their gyri, sulci, and cerebral ventricles.

Based on the five T_1-weighted brains and their coregistered, simulated fiber tracts, we subsequently synthesized five corresponding tensor fields in the presence of random noise. Because the direction of the synthesized fibers was known, the tensor field was generated by taking the known orientation of the fibers as one of the eigenvectors having a corresponding eigenvalue of 1.0 and randomly selecting the other two orthogonal directions that had less prominent eigenvalues of 0.2 and 0.1, corresponding to an FA value equal to 0.834 [16], [17].

Within the remaining volume of each brain, we defined tensors that were nearly isotropic ($\lambda_1 = 1.0, \lambda_2 = \lambda_3 = 0.98, FA = 0.0117$). For regions outside of the brain, tensors were defined as zero. FA values in both of these nonfiber regions were thus very close (0.0117) to or equal to 0. We then reversely calculated the attenuation ratio of the diffusion-weighted signal to simulate measurements obtained from living brains in an MRI scanner. Gaussian additive noise was then superimposed

and the tensors were reconstructed. The shapes of the noise-superimposed tensors thus could range from prolate to spherical, and consequently their FA values could range from 0.0 to 1.0.

Based on the theory of DTI [29], [30], we reversed the process of DTI reconstruction to simulate noise (SN) in raw imaging data and subsequently in tensors as follows.

SN_1 Select a set of k directions, $d_1, d_2, \ldots, d_k, d_i = (d_{i,x}, d_{i,y}, d_{i,z})$, that are evenly distributed in 3-D space. These will serve as the directions along which the diffusion-weighted (DW) images are measured.

SN_2 Recover the attenuation ratio of DW measurements along each selected direction $d_i (i = 1, \ldots, k)$.

For each noise-free tensor D_{original}, calculate every A_i along the ith direction.

$A_i = \exp(-bd_i^T D_{\text{original}} d_i)(0.0 < A_i \leqslant 1.0)$

where b is the b-value and d_i is the gradient direction vector.

SN_3 Add to A_i a noise component having a Gaussian distribution of $N(0, \sigma)$

$A_i' = A_i + N(0, \sigma)$.

SN_4 Reconstruct the A_i's $(i = 1, \ldots, k)$ to a noise-superimposed tensor D_{noise} using linear regression.

After noise was added, tensors in the brain that were originally nonzero became degraded (usually a tensor's FA was lowered and its orientation was modified), whereas regions outside the brain were replaced by noise-superimposed tensors that had random orientations and FA values. We then normalized the five simulated DT datasets to the selected template with displacement fields that were generated by HAMMER [19] based on the warping and coregistration of their corresponding T1-weighted images (HAMMER was used for this purpose simply because it is a readily available and accurate tool for the normalization of scalar images). By averaging across the corresponding tensors in each volume, we obtained the average tensor field. Note that here we used T1-weighted images instead of an FA-map for generating the displacement fields for spatial normalization, because the DT datasets were simulated based on the anatomy of T1-weighted images; they therefore were perfectly coregistered. Also note that we intentionally used STAR for data preparation and HAMMER for normalization, so that registration errors would not likely be covered by the same algorithm, thereby making our experiment more valid.

To measure the deviation of warped fibers, we calculated the voxel-wise difference in orientation of the PD between the template and the average PD across the five brains warped into the template space. Because PD and -PD represent the same orientation of the primary direction of the diffusion tensor, the deviation at each voxel was measured within the range $(0°, 90°)$. We repeated this calculation of PD differences using forward mapping procedures for the reorienting of tensors. We compared the differences in PDs of the tensors using the proposed seamless warping with the differences in PDs using forward warping alone in order to determine whether our method of reorientation offered an improvement over the conventional approach to DTI warping.

Results: When simulating noise, we used 6 spatial directions [31] evenly distributed in 3-D space (algorithm steps **SN_1** ∼ **SN_4**). Table I shows that in fiber bundles (excluding tensors

that lie on the boundary between heterogeneous tissues types), the average variance of the PD generated from forward mapping is $5.70°$. Although this value already is quite small, our seamless warping procedure managed to reduce PD variance by 26.0%. When taking into account the boundaries of warped fibers where partial volume and some interpolations between well-defined fibers and undefined regions already introduced errors, improvement was approximately 9.8%. The results show that our proposed seamless DTI warping algorithm substantially outperformed the forward mapping algorithm with Procrustean estimation, which itself was already quite robust.

D. Experiment 4

Experimental Design: We demonstrate here the effectiveness of our seamless warping algorithm using *in vivo* datasets from 13 healthy human subjects. After subjects provided written informed consent, they were scanned on a Phillips 3T MRI scanner using a single-shot EPI sequence with the SENSE parallel imaging scheme: SENSitivity Encoding, SENSE reduction factor $R = 2.5$, imaging matrix 96×96, FOV 240×240 mm, nominal resolution 2.5 mm, zero-filled to a 256×256; axial slices thickness 2.5 mm parallel to the anterior–posterior commissure line; total 50–55 slices whole brain coverage without gaps. Diffusion weighting was encoded along 30 independent directions and the b-value 700 s/mm^2. Five additional images with minimal diffusion weighting ($b = 33$ s/mm^2) were also acquired with gradient applied along slice-direction. The scanning time of a single dataset took 20–25 min.

The raw diffusion-weighted images were first realigned using AIR [32], [33] to remove motion that occurred during the scan. Subsequently, all individual images were visually inspected to discard slices that contained motion artifact, after which the remaining images were added for each slice. Because the DWI data have only minimal distortion around the sinuses [12], no distortion corrections were necessary. The pixel intensities of the multiple diffusion-weighted images were then fitted to obtain the six elements of the diffusion tensor.

We selected one subject randomly from this pool of 13 subjects as a template, and then calculated the DF from the dataset for each subject to the template space using HAMMER [19] applied to their FA-maps. We then used a forward mapping strategy and our bijection approach to warp data from each individual subject to the template brain. Thus, we obtained two sets of warping for all 13 subjects coregistered in the common template space.

Results Fig. 9 shows a typical slice of the warped dataset from one representative subject using the forward mapping strategy [Fig. 9(a)] or our bijection approach [Fig. 9(b)]. Although the two color maps are similar in general, artifacts caused by seams (described in Fig. 1) at several locations for forward mapping do not appear in the bijection-derived map. We attribute such artifacts mainly to the seams generated by forward mapping. Similar artifacts occurred throughout the imaging volume.

Fig. 10(a) shows the PD color map of the average of all 13 datasets normalized using our bijection warping strategy. The blue fiber bundles, which indicate fiber pathways running in the inferior-superior direction within the posterior limb of the
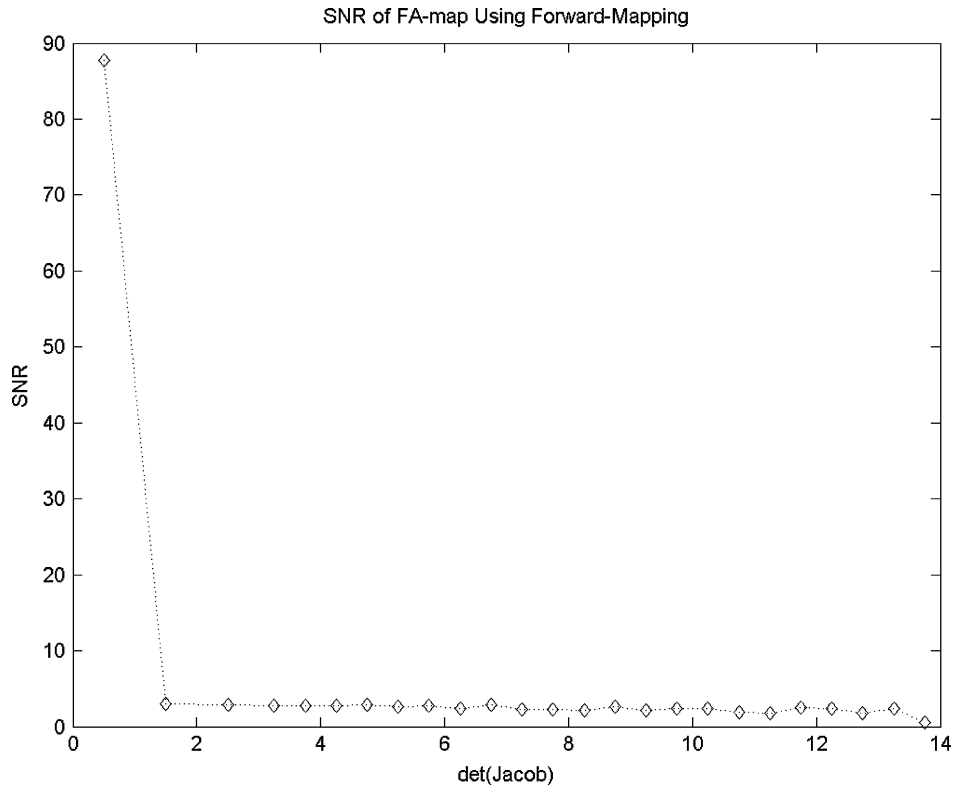
Fig. 8. SNR with forward mapping. SNR drops significantly at expanded locations, where the determinant of the Jacobian matrix of local deformation is greater than 1.0. In contrast, the corresponding SNR with bijection warping remained at $+8$ (not shown).

TABLE I
COMPARISON OF PD VARIANCE USING FORWARD AND
SEAMLESS BACKWARD DTI WARPING

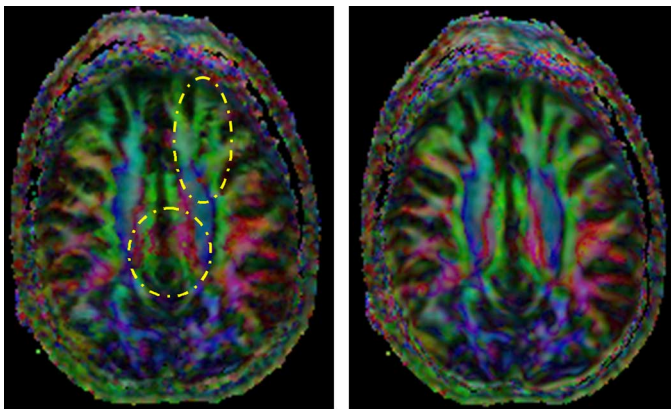|	Fiber Bundles Only	Fiber Bundles with Boundary
Forward Warping	5.70°	11.36°
Seamless Bijection Warping	4.22°	10.25°
Improvement	25.96%	9.77%

Fig. 9. Differences between forward and bijection mappings in *in vivo* datasets. Left: PD colormap of a warped DTI dataset, generated with forward mapping. Zig-zag shaped artifacts (in the yellow ellipse) and the blurring artifacts (in the yellow circle) are caused by seams. Right: PD color map of the same subject generated with bijection mapping. Note the absence of the artifacts.

internal capsule, are brighter and more clearly delineated than are those generated by forward mapping [Fig. 10(b)]. Consequently, fiber tracking on these structures was, therefore, easier and more robust [Fig. 10(c)]. Finally, the average error for the 12 bijections generated for this experiment was consistently 0.093–0.097 mm, with a standard deviation at 0.17 mm.

V. DISCUSSION

We have proposed a method for warping DT datasets that uses a bijection DF along with Procrustean Estimation to reorient tensors appropriately in a template imaging space. By combining forward and backward mapping under the constraint of consistency (see Section III-A), our method achieved seamless warping and proper geometric reorientation of tensors in the warped template space using information about the local deformation and geometric orientation of tensors in the original imaging space. Moreover, our method generated reverse deformation fields with a high degree of accuracy. We demonstrated that the bijection map generated using this algorithm is a consistent one-to-one map (as defined quantitatively in Section III), having an average error of less than 0.1 mm between the forward and backward mapping components of the bijection. We also found that using a level of detail of 2 in super-sampling is optimal in terms of the cost-benefit trade-off of accuracy versus computational demands, in that more extensive super-sampling provides only slightly better mapping accuracy but with exponentially increasing computational demands. Nevertheless, in areas of extreme deformation, a higher level of super-sampling may be needed to achieve better accuracy, and in these instances,
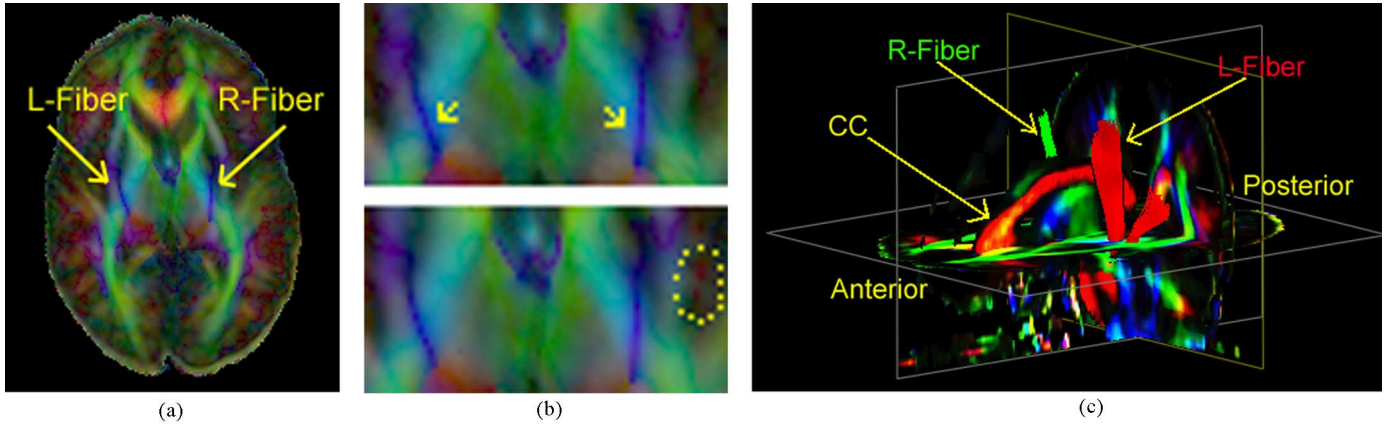
Fig. 10. Averaged DTI dataset from 13 subjects using bijection warping. (a) PD colormap (using the same color schema as in Fig. 7) of a single axial slice in the averaged tensor map. Blue fiber bundles (arrows) are exquisitely sharp, allowing easy fiber tracking. Fibers on the left (L-Fiber) or right (R-Fiber) side of the image are tracked in (c) using these two blue strips as seed points. (b) Close-up comparison of the blue strips in (a) for the averaged results using a warping based on forward mapping (upper panel) and a warping based on bijection mapping (lower panel). Although the two views look similar, subtle differences can be found: the blue strips at the arrows in the upper panel are thicker and fuzzier compared with the same blue strips in the lower panel, representing a less distinct delineation of fiber bundles; in addition, the boundary of the reddish region in the yellow dotted loop is sharper in the lower panel than is the same region in the upper panel, again showing a clearer delineation of tissue organization. Both indicate that the bijection warping method preserves signal better than forward mapping does. (c) Tracked fibers are embedded in a triplanar view of the color map of average PDs in the template space. Note that the colors of the two reconstructed fiber bundles (red and green) do not carry directional information, but are used purely for the purposes of identifying the two fiber bundles. This is a left dorsal anterior perspective (viewed from the left, above the forehead). The reconstructed fiber bundles track to the dorsum of the brain in the region of premotor and sensorimotor cortices. CC=corpus callosum; L-Fiber=fiber tracts in the left hemisphere; R-Fiber $=$ fiber tracts in the right hemisphere.

adaptive super-sampling that is based on the magnitude of deformation may be desirable for generating a more accurate bijection mapping. Finally, although throughout this paper we have presumed that the bijection is initiated using a forward mapping, the bijection can be initiated with a backward mapping as well.

Using FA measurements as indices of tensor morphology, we have demonstrated that seamless warping of DT datasets using bijection robustly preserves tensor morphology. In contrast, we also demonstrated quantitatively that when using forward mapping, the SNR of DT datasets decreased significantly at locations of extreme deformation ($J > 1$), thereby creating seams and consequently artifacts (Fig. 9).

We used Procrustean Estimation to estimate tensor rotation instead of directly averaging their PDs. Although Procrustean Estimation may appear more complex than other methods for estimating rotation, it is actually quite elegant mathematically and computationally efficient. By comparison, averaging tensor orientations based on their PDs to estimate tensor rotation is much more complicated, because many more orientations are considered within a given neighborhood. This is the case because both of the eigenvectors v and $-v$ represent the same orientation, and thus both $\pm(v_1 + v_2)$ and $\pm(v_1 - v_2)$ are reasonable and meaningful average orientations of v_1 and v_2. Procrustean Estimation, on the other hand, essentially calculates the rotation that minimizes, in terms of least-squares, the differences of corresponding coordinate pairs, and these differences are rotation-invariant. Therefore, whereas the coordinates entered in matrixes **A** and **B** (in algorithm ERM) are relative positions derived from PD vectors, Procrustean Estimation is free from the dilemma of interpolating orientations.

A central and long standing issue in the processing of DT datasets has been the technical difficulty of achieving an accurate normalization of images from this modality alone. Normalization of DT datasets before generating FA maps, fiber tracts,

PD color maps, and other measurements derived from DT data is more efficient than performing normalization on these measurements after these measures are generated because, without an initial normalization, each set of measurements must be normalized individually to the template space. In contrast, if all DT datasets are first normalized to a template space, all of these measurements can then be derived directly in the template space and therefore do not require spatial normalization individually. Thus, normalizing DT datasets first, before generating measurements derived from DT data, drastically reduces the number of times that normalization must be performed (as well as eliminating the duplication of data and need for additional storage space that are required in order to normalize each set of measurements separately).

Despite these considerable advantages of spatial normalization prior to calculation of derivative measures, investigators sometimes choose the less efficient of these two options because of the known technical difficulty of accurately normalizing datasets that contain high-dimensional information. Indeed a common practice is to normalize each set of derived measures (e.g., FA maps, fiber tracts) without normalizing the DT datasets themselves (although even this approach is challenging, in that the derived measurements may also contain high-dimensional information that makes them difficult to normalize). Our algorithm addresses these various problems and challenges by providing a means for normalizing DT datasets more accurately prior to calculation of derivative measures.

Our seamless warping procedure in principle makes possible the accurate interpolation of tensors because of its use of bijection (see Section III-D). In general, interpolating DWI data and then calculating a new tensor based on those data is the most precise way to interpolate a tensor. However, because this method requires the use of all raw DW data for interpolating tensors, raw DW data must either be saved or recovered, the latter being com-

putationally demanding (e.g., the Stejskal equation [34] can be applied in reverse to tensors, the calculated $b0$ baseline image, and gradient directions in order to recover DWI data, similar to our approach in the Procedure SN, as described in Section IV-C). Moreover, conventional DTI warping typically uses forward mapping procedures that require interpolation of tensors in the target space. DWI data, however, are usually not readily available for interpolation in the target space, the methods for interpolating tensors directly in the target space are computationally demanding, and their accuracy is questionable. Our method for seamless warping, in contrast, accurately interpolates tensors in the original space because the backward mapping component in the bijection allows convenient access to the raw DWI data in the original imaging space, thereby allowing for the efficient and accurate interpolation of DWI data.

In our experiments, we used trilinear interpolation of tensor components for tensor interpolation, because we found that the slight gain in accuracy obtained from interpolating DWI data in the original space did not justify its additional expense of data management and computation. Use of trilinear interpolation is justified because DTI data usually do not change significantly within a small neighborhood (particularly not within the distance of the single voxel over which we interpolated). Linear interpolation of tensors thus provides a very nice and efficient approximation of the true values of tensors reconstructed from interpolated DWI data, and for these reasons it is the method most commonly used for tensor interpolation.

VI. CONCLUSION

Techniques that have been widely used to normalize scalar images to a template have been difficult to extend to DTI datasets. The greatest impediments to spatial normalization of DTI datasets have been the two-fold challenge of avoiding the introduction of seams (and thereby artifacts) into the template space while reorienting tensors in a way that preserves both their morphology and the integrity of underlying fiber tracts as identified in the original imaging space. Our seamless warping algorithm using bijection combined with Procrustean Estimation has successfully overcome these difficulties. DT datasets contain high-dimensional measurements of diffusivity that must be preserved during the warping procedure because they reflect the local structure of biological tissues. Thus any technique for spatial normalization that is extended from the domain of traditional scalar imaging to DTI likely requires careful and substantial modification to ensure preservation of that biologically relevant information. Indeed, this same cautionary note applies to the extension of scalar techniques to any MRI modality other than anatomical imaging. Future work should develop similar valid and robust algorithms for the processing and analysis of data from multiple MRI modalities within a spatially normalized template.

ACKNOWLEDGMENT

REFERENCES

[1] K. Fredericksen, L. Amodei, W. E. Kaufmann, P. C. van Zijl, G. D. Pearlson, and S. Mori, "Automated white matter delineation tools based on diffusion tensor imaging," *Int. Soc. Magn. Reson. Med.*, p. 1532, 2001.

[2] M. Kubicki, C.-F. Westin, R. W. McCarley, and M. E. Shenton, "The application of DTI to investigate white matter abnormalities in schizophrenia," *Ann. NY Acad. Sci.*, vol. 1064, pp. 134–148, Dec. 2005.

[3] S. Mori and P. C. van Zijl, "Fiber tracking: Principles and strategies—A technical review," *NMR Biomed.*, vol. 15, pp. 468–480, 2002.

[4] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee, "Spatial transformations of diffusion tensor magnetic resonance images," *IEEE Trans. Med. Imag.*, vol. 20, no. 11, pp. 1131–1139, Nov. 2001.

[5] G. H. Golub and C. F. Van Loan, *Matrix Computation*. Baltimore, MD: Johns Hopkins Univ. Press, 1983.

[6] L. E. Malvern, *Introduction to the Mechanics of a Continuous Medium*. New York: Prentice-Hall, 1969.

[7] M. Guye, G. Parker, S. Symms, P. Boulby, C. Wheeler-Kingshott, A. Salek-Haddadi, G. J. Barker, and J. S. Duncan, "Combined functional MRI and tractography to demonstrate the connectivity of the human primary motor cortex in vivo," *NeuroImage*, vol. 19, pp. 1349–1360, 2003.

[8] D. K. Jones, L. D. Griffin, D. C. Alexander, M. Catani, M. A. Horsfield, R. Howard, and S. C. R. Willams, "Spatial normalization and averaging of diffusion tensor MRI data sets," *NeuroImage*, vol. 17, no. 2, pp. 592–617, Oct. 2002.

[9] R. P. Woods, S. T. Grafton, C. J. Holmes, S. R. Cherry, and J. C. Mazziotta, "Automated image registration—ii: Intersubject validation of linear and nonlinear models," *J. Comp. Assisted Tomogr.*, vol. 22, no. 1, pp. 153–165, 1998.

[10] R. Woods, S. Grafton, C. Holmes, S. Cherry, and J. Mazziotta, "Automated image registration: I. General methods and intrasubject, intramodality validation," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 139–152, 1998.

[11] J. Ruiz-Alzola, C. F. Westin, S. K. Warfield, C. Alberola, S. Maier, and R. Kikinis, "Nonrigid registration of 3-D tensor medical data," *Med. Image Anal.*, vol. 6, no. 2, pp. 143–161, 2002.

[12] D. Xu, S. Mori, D. Shen, P. van Zijl, and C. Davatzikos, "Spatial normalization of diffusion tensor fields," *Magn. Reson. Med.*, vol. 50, no. 1, pp. 175–182, Jul. 2003.

[13] H. Zhang, "Forward shadow mapping," *Proc. Europgraphics Rendering Workshop*, pp. 131–138, 1998.

[14] L. A. Westover, "SPLATTING: A parallel, feed-forward volume rendering algorithm," Ph.D. dissertation, Dept. Computer Sci., Univ. North Carolina, Chapel Hill, 1991.

[15] J. C. Gee, M. Reivich, and R. Bajcsy, "Elastically deforming a three-dimensional atlas to match anatomical brain images," *J. Comput. Assist. Tomogr.*, vol. 17, no. 2, pp. 225–236, 1993.

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

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

[18] K. M. Hasan, A. L. Alexander, and P. A. Narayana, "Does fractional anisotropy have better noise immunity characteristics than relative anisotropy in diffusion tensor MRI? An analytical approach," *Magn. Reson. Med.*, vol. 51, pp. 413–417, 2004.

[19] D. Shen and C. Davatzikos, "HAMMER: Hierarchical attribute matching mechanism for elastic registration," *IEEE Trans. Med. Imag.*, vol. 21, no. 11, pp. 1421–1439, Nov. 2002.

[20] R. Bansal, L. H. Staib, R. Whiteman, Y. M. Wang, and B. S. Peterson, "ROC-based assessments of 3-D cortical surface-matching algorithms," *NeuroImage*, vol. 24, no. 1, pp. 150–162, Jan. 2005.

[21] G. E. Christensen, S. C. Joshi, and M. I. Miller, "Volumetric transformation of brain anatomy," *IEEE Trans. Med. Imag.*, vol. 16, no. 6, pp. 864–877, Dec. 1997.

[22] K. Friston, A. Holmes, K. Worsley, J.-B. Poline, C. Frith, and R. Frackowiak, "Statistical parametric maps in functional imaging: A general linear approach," *Hum. Brain Mapp.*, vol. 2, pp. 189–210, 1995.

[23] K. J. Friston, J. Ashburner, C. D. Frith, J.-B. Poline, J. D. Heather, and R. S. J. Frackowiak, "Spatial registration and normalization of images," *Hum. Brain Mapp.*, vol. 3, pp. 165–189, 1995.

[24] R. Woods, S. Grafton, C. Holmes, S. Cherry, and J. Mazziotta, "Automated image registration: I. general methods and intrasubject, intramodality validation," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 139–152, 1998.

[25] 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.

[26] K. J. Plessen, R. Grüner, A. Lundervold, J. G. Hirsch, D. Xu, R. Bansal, Â. Hammar, A. J. Lundervold, T. Wentzel-Larsen, S. A. Lie, A. Gass, B. S. Peterson, and K. Hugdahl, "Reduced white matter connectivity in the corpus callosum of children with tourette syndrome," *J. Child Psychol. Psychiat.*, vol. 47, pp. 1013–1022, 2006.

[27] C. Davatzikos, "Spatial transformation and registration of brain images using elastically deformable models," *Comput. Vis. Image Understand.*, vol. 66, no. 2, pp. 207–222, 1997.

[28] C. Davatzikos, A. Genc, D. Xu, and S. M. Resnick, "Voxel-Based morphometry using the RAVENS maps: Methods and validation using simulated longitudinal atrophy," *NeuroImage*, vol. 14, no. 6, pp. 1361–1369, Dec. 2001.

[29] P. J. Basser and D. K. Jones, "Diffusion-tensor MRI: Theory, experimental design and data analysis—A technical review," *NMR Biomed.*, vol. 15, no. 7-8, pp. 456–467, Nov. 2002.

[30] D. Le Bihan *et al.*, "Diffusion tensor imaging: Concepts and applications," *J. Magn. Reson. Imag.*, vol. 13, pp. 534–546, 2001.

[31] S. Skare, M. Hedehus, M. E. Moseley, and T. Q. Li, "Condition number as a measure of noise performance of diffusion tensor data acquisition schemes with MRI," *J. Magn. Reson.*, vol. 147, pp. 340–352, Dec. 2000.

[32] R. Woods, S. Grafton, C. Holmes, S. Cherry, and J. Mazziotta, "Automated image registration: I. General methods and intrasubject, intramodality validation," *J. Comput. Assist. Tomogr.*, vol. 22, pp. 139–152, 1998.

[33] R. P. Woods, S. T. Grafton, J. D. G. Watson, N. L. Sicotte, and J. C. Mazziotta, "I: Intersubject validation of linear and nonlinear models," *J. Comput. Assist. Tomogr.*, vol. 22, no. 1, pp. 153–165, Jan. 1998.

[34] E. O. Stejskal and J. E. Tanner, "Spin diffusion measurements: Spin-echoes in the presence of a time dependent field gradient," *J. Chem. Phys.*, vol. 42, pp. 288–292, 1965.