

TIMER: Tensor Image Morphing for Elastic Registration

Pew-Thian Yap^a, Guorong Wu^a, Hongtu Zhu^b, Weili Lin^a, Dinggang Shen^{a,*}

^a Department of Radiology and BRIC, University of North Carolina at Chapel Hill, NC, USA

^b Department of Biostatistics and BRIC, University of North Carolina at Chapel Hill, NC, USA

ARTICLE INFO

Article history:

Received 29 January 2009

Revised 15 April 2009

Accepted 16 April 2009

Available online 3 May 2009

Keywords:

Diffusion tensor imaging

Elastic registration

Tensor regional distributions

Tensor boundaries

Log-Euclidean manifold

ABSTRACT

We propose a novel diffusion tensor imaging (DTI) registration algorithm, called *Tensor Image Morphing for Elastic Registration* (TIMER), which leverages the hierarchical guidance of regional distributions and local boundaries, both extracted directly from the tensors. Currently available DTI registration methods generally extract tensor scalar features from each tensor to construct scalar maps. Subsequently, regional integration and other operations such as edge detection are performed to extract more features to guide the registration. However, there are two major limitations with these approaches. First, the computed regional features might not reflect the actual regional tensor distributions. Second, by the same token, gradient maps calculated from the tensor-derived scalar feature maps might not represent the actual tissue tensor boundaries. To overcome these limitations, we propose a new approach which extracts regional and edge information directly from a tensor neighborhood. Regional tensor distribution information, such as mean and variance, is computed in a multiscale fashion directly from the tensors by taking into account the voxel neighborhood of different sizes, and hence capturing tensor information at different scales, which in turn can be employed to hierarchically guide the registration. Such multiscale scheme can help alleviate the problem of local minimum and is also more robust to noise since one can better determine the statistical properties of each voxel by taking into account the properties of its surrounding. Also incorporated in our method is edge information extracted directly from the tensors, which is crucial to facilitate registration of tissue boundaries. Experiments involving real subjects, simulated subjects, fiber tracking, and atrophy detection indicate that TIMER performs better than the other methods (Yang et al., 2008; Zhang et al., 2006).

© 2009 Elsevier Inc. All rights reserved.

Introduction

Diffusion tensor imaging (DTI) is capable of measuring in vivo tissue water diffusion non-invasively, and has been widely employed to delineate potential neurological disease related white matter abnormalities. To this end, accurate registration of diffusion tensor images across different subjects is a critical prerequisite for detailed voxel-by-voxel statistical analysis. However, spatial normalization of diffusion tensor images is challenging both technically and computationally given that tensor data representation is in nature high dimensional, and it is required that the tensors not only be spatially warped, but also be reoriented to appropriately reflect the anisotropic diffusivity.

In conventional MRI, such as T_1 - and T_2 - weighted images, white matter often appears as a homogeneous structure, even though it consists of many axonal bundles of various sizes and orientations. The advent of DTI makes a deeper investigation into these structures possible (Huang et al., 2006; Lee et al., 2008; Liu et al., 2007; Yang

et al., 2008b). It has been observed that water diffusivity is more restricted in the transverse direction of the axonal fibers in white matter tissue than along the axial direction. DTI takes advantage of this anisotropic diffusivity and measures a symmetric second order tensor at each voxel characterizing the anisotropy (Alexander, 2006). The tensor provides an estimate of the dominant orientation of microstructural fibers, and also rotationally invariant statistics of the anisotropy of the particle displacement probability density function. Assuming Gaussian distribution, the shape of the probability density function is determined by the eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ of the tensor \mathbf{D} . The Gaussian function has an ellipsoidal contour, relative lengths of the major axes of which have the same proportions as $(\lambda_i)^{\frac{1}{2}}$. Statistics of anisotropy come from the distribution of eigenvalues. The principal direction (PD) of the tensor is presumed to coincide with the underlying fiber orientation. With fiber orientation information, one can identify various axonal tracts within the homogeneous-looking white matter. Therefore, DTI can provide unique microstructural and physiological insight into the white matter of brains, which has not been found available in other MR modalities, and has thus become a powerful technique for analyzing the underlying white matter structures of brains (Mori et al., 2001; Pierpaoli et al., 1996). The rich microstructural information of white matter facilitates the study

* Corresponding author.

E-mail addresses: ptyap@med.unc.edu (P.-T. Yap), grwu@med.unc.edu (G. Wu), hzzhu@bios.unc.edu (H. Zhu), weili_lin@med.unc.edu (W. Lin), dgshen@med.unc.edu (D. Shen).

of development, aging, and disease on specific white matter regions of interest.

To quantify diffusion tensor abnormalities utilizing voxel-based statistical analysis, spatial normalization is necessary in order to minimize the anatomical variability between studied brain structures. Conventional methods generally extract tensor scalar features from each tensor individually, and by constructing scalar maps, regional integration and other operations such as edge detection can be performed to extract final features for correspondence matching. As an example, Yang et al. (2008a) capitalize on tensor structural geometry and orientation information for registration guidance. Specifically, geometric measures, describing the diffusion tensor geometrical resemblance to the generic structures of prolateness, oblateness, and sphericity (Westin et al., 1997), are computed for each voxel. The distributions of the measures in a spherical neighborhood of each voxel are then estimated using local histograms. These local histograms, together with boundary information extracted from the fractional anisotropy (FA) and apparent diffusion coefficient (ADC) maps are finally fed into a deformable correspondence matching mechanism (Shen, 2007) to align the images.

There are, however, two major limitations with these approaches. First, the computed regional features might not reflect the actual regional tensor distributions. Second, by the same token, gradient maps calculated from the tensor-derived scalar feature maps might not represent the actual tissue tensor boundaries. While it is possible to use the results of scalar structural image registration to establish correspondence matching, it is natural to expect that the inclusion of additional information within the diffusion tensor will lead to more accurate matching and hence more robust registration, particularly in white matter regions where the microstructural variation is not observable in conventional MRIs. In view of this, new approaches are proposed to work directly on the tensors. However, one should keep in mind that methods involving single voxel correspondence matching can be susceptible to the effect of noise and also spurious variation of anatomical structures. A more robust approach is to leverage information gathered from the voxel neighborhood to help estimate better voxel statistics, so as to make correspondence matching less affected by noise.

In this paper, we propose a novel diffusion tensor image registration algorithm, called *Tensor Image Morphing for Elastic Registration* (TIMER). Instead of working with the tensors in a voxel-wise fashion, we estimate the statistical properties, such as mean and variance, at a particular voxel location by taking into consideration the information furnished by the voxel neighborhood. Notably, these regional information are obtained directly from the tensors and not their scalar maps. Scalar maps do not retain full tensor information as some information is discarded in the process of their computation, and hence they do not reflect true tensor structures. As a simple example, a homogeneous region on a FA map might contain tensors of radically different orientations and there is no way to tell the difference by solely gauging from the FA map, simply because the FA map does not hold orientation information. Computing regional distribution directly using the tensors take all these information into account and yields a more precise estimation of the statistical property at a particular voxel location. However, measures derived from tensor regional distributions in a single resolution are very limited in their discriminating power as in reality anatomical details often appear in different scales. This problem can be overcome by simply combining a set of regional features at multiple scales. Tensor regional information from each scale captures different spatial information and when combined forms a rich discriminative feature vector. It can be observed from Fig. 1 that regional information of any single resolution is not adequate for differentiating different anatomical structures, but by simply combining regional information gained from three different resolutions, we can sufficiently achieve differentiation.

It is worth noting that features captured from regional distribution information are coarse features, thus the use of these features might not yield accurate correspondence detection. To ensure accurate alignment of tissue boundaries, local features such as SIFT (Lowe, 2004) and RIFT (Lazebnik et al., 2005) can be incorporated to aid the registration. For TIMER, we extend Canny edge detector (Canny, 1986) to work directly on the diffusion tensors. The extended Canny edge detector outputs a point-wise boundary map, with zero as non-boundary and other non-zero values indicating tensor boundary strength. Computed at a coarse scale, boundaries of major white matter tracts can be located and employed to yield an initial estimation of the alignment, and at the same time mitigate the matching ambiguity posed by the more detailed white matter tracts at a finer scale. More boundaries are included in a concerted effort to refine the initial alignment at later stages of the registration. Edge information, together with the above-mentioned regional information, is employed to form an attribute vector for each voxel, which serves as a morphological signature to reflect the anatomical context of the voxel at multiple scales. Attribute vectors are utilized to assist automated correspondence detection in a hierarchical and spatially deformable fashion (Shen and Davatzikos, 2002).

TIMER is essentially a feature-based registration algorithm which utilizes anatomical knowledge in determining point correspondences. During the initial stage of the registration, a small number of the most representative voxels are selected to guide the registration. This essentially mitigates the problem of local minima, which often happens when voxels with less discriminative attribute vectors are employed, resulting in ambiguity during correspondence matching. As the registration progresses, more and more voxels are included so that they can work in a concerted manner to refine the registration.

The rest of the paper is organized as follows: The Materials and methods section explains the working mechanism of TIMER. Constituents of the attribute vector of TIMER are detailed, along with its similarity measure. The Results section gives the results of simulated deformation field, fiber tracking and atrophy detection experiments, which were used to evaluate TIMER. We give further discussion and conclude the paper in The Discussion section.

Materials and methods

Dataset

The dataset consisted of diffusion tensor images of 22 subjects, acquired using a 1.5 T scanner. Each image was acquired using 30 gradient directions with the diffusion weighting of $b = 700 \text{ s/mm}^2$ ($NEX = 2$). The imaging dimension was 256×256 with a rectangular field of view (FOV) of $240 \times 240 \text{ mm}^2$ and image resolution of $0.9375 \text{ mm} \times 0.9375 \text{ mm} \times 2.5 \text{ mm}$. All of the diffusion tensor data, as well as the derived scalar maps, were skull-stripped to extract the brain parenchyma before they were used in the experiments.

Attribute vectors

In this paper, building upon the ideas in (Shen and Davatzikos, 2002), we formulate a deformable registration scheme which aims to alleviate the problem of local minima, which is critical since brain anatomy is inherently high dimensional, complex and ambiguous. This is accomplished by employing a hierarchical mechanism where initially only a small subset of voxels, which exhibit highly distinctive attribute vectors, are used to drive the volumetric deformation. This essentially means that we are now solving a lower-degrees-of-freedom approximation of the deformation field, and is hence less prone to local minima. In the later stages, as more voxels participate in driving the registration, the deformation field can be refined with increasing complexity. The driving voxels are selected hierarchically according to how uniquely their attribute vectors stand out among




Fig. 1. The principal diffusivities of the mean in the neighborhood of the point indicated in the top image is compared with those of the other points in the brain image volume. Dark red indicates high similarity and dark blue indicates otherwise. (a), (b) and (c) are the results obtained at fine, middle and coarse resolution representations of the diffusion tensor image, respectively. (d) is the result when all the individual resolutions in (a), (b) and (c) are combined. The small delineated dark red area in (d) is indicative that a correspondence can be correctly located as opposed to (a), (b) and (c) where the dark red regions are spread out with no clear clue of where the corresponding point is. The white cross in (d) corresponds with the red cross in the top image.

others in the image, reducing ambiguity and hence local minima. To make the attribute vectors more descriptive and hence more discriminative, they are computed at multiple scales so that anatomical structures at different resolutions are accounted for.

Fig. 2 illustrates how multiscale features are extracted from the image. Instead of varying the neighborhood size, the image is progressively downsampled and multiscale features are extracted from the fine, middle and coarse resolution images using a fixed neighborhood radius. The attribute vector in TIMER is designed to comprise three different types of features: 1) Regional features (means and variances), 2) Edge features (tensor edges and FA map

edges), and 3) Geometrical features (FA values and principal diffusivities). A number of other features can be used in place of these features. For regional features, possible candidates include fiber-tract organization measures (Basser and Pierpaoli, 1996), higher order moments (Mukundan and Ramakrishnan, 1998), and also other inter-tensor measures (Arsigny et al., 2006). For local boundary features, SIFT (Lowe, 2004) and RIFT (Lazebnik et al., 2005) are possible choices. And for tensor geometrical features, prolateness, oblateness, and sphericity measures (Westin et al., 1997) and orientation features (Yang et al., 2008a) can be employed. However, we find that the features currently used in TIMER are sufficiently distinctive to give




Fig. 2. Computation of multiscale features. Instead of changing neighborhood size, the image is progressively downsampled and multiscale features are computed using a fixed neighborhood radius as indicated by the circle. This approach cuts down on the computation time. The first three images on the left show the coarse, middle and fine resolution images, respectively. The next three images show the edges detected at different resolutions.

good registration accuracy. These multiscale features are grouped into an attribute vector for each voxel, i.e.:

$$\mathbf{a} = \left[\left(\mathbf{a}_F^{\text{Regional}}, \mathbf{a}_F^{\text{Edge}}, \mathbf{a}_F^{\text{Geometrical}} \right), \left(\mathbf{a}_M^{\text{Regional}}, \mathbf{a}_M^{\text{Edge}}, \mathbf{a}_M^{\text{Geometrical}} \right), \left(\mathbf{a}_C^{\text{Regional}}, \mathbf{a}_C^{\text{Edge}}, \mathbf{a}_C^{\text{Geometrical}} \right) \right] \quad (1)$$

where the subscripts denote the resolution (F: fine, M: middle, C: coarse) from which the features are derived. The similarity of two attribute vectors $\mathbf{a}(\mathbf{x}_1)$ and $\mathbf{a}(\mathbf{x}_2)$, after normalizing each of these attributes to be within the range of [0,1], is defined as:

$$m(\mathbf{a}(\mathbf{x}_1), \mathbf{a}(\mathbf{x}_2)) = \prod_{s \in \{L, M, C\}} \left[\prod_i \left(1 - |a_{s,i}^{\text{Regional}}(\mathbf{x}_1) - a_{s,i}^{\text{Regional}}(\mathbf{x}_2)| \right) \times \prod_j \left(1 - |a_{s,j}^{\text{Edge}}(\mathbf{x}_1) - a_{s,j}^{\text{Edge}}(\mathbf{x}_2)| \right) \times \prod_k \left(1 - |a_{s,k}^{\text{Geometrical}}(\mathbf{x}_1) - a_{s,k}^{\text{Geometrical}}(\mathbf{x}_2)| \right) \right] \quad (2)$$

where $a_{s,i}^{\text{Regional}}(\mathbf{x})$, $a_{s,j}^{\text{Edge}}(\mathbf{x})$, and $a_{s,k}^{\text{Geometrical}}(\mathbf{x})$ are the i -th, j -th and k -th element of $\mathbf{a}_s^{\text{Regional}}(\mathbf{x})$, $\mathbf{a}_s^{\text{Edge}}(\mathbf{x})$, and $\mathbf{a}_s^{\text{Geometrical}}(\mathbf{x})$ at scale $s \in \{F, M, C\}$, respectively. Fig. 3 shows that the attribute vectors are rich enough to warrant the differentiation of different anatomical structures and hence can be utilized to assist correspondence matching in the course




Fig. 3. Distinctiveness of attribute vector. The similarity map in the green box is magnified for a closer inspection. Dark red indicates high similarity and dark blue otherwise.

of registration. More detailed descriptions of the features used in TIMER is furnished in the following sections.

Regional statistical features

For each point in the diffusion tensor image, regional tensor distribution information can be extracted from its neighborhoods in a multiscale fashion. Specifically, for a given tensor $\mathbf{D}(\mathbf{x})$, we can extract the regional tensor distribution information from its neighborhood $\{\mathbf{D}(\mathbf{t}) | \mathbf{t} \in N(\mathbf{x})\}$, where $N(\mathbf{x})$ is the set of coordinates of voxels in the neighborhood of \mathbf{x} . By varying the size of neighborhood $N(\mathbf{x})$ or scaling the image, we can obtain rich multiscale regional tensor distribution information. For effective estimation of tensor distribution information, we first use logarithm to transform the original tensor to its log-space counterpart, i.e., $\log(\mathbf{D}(\mathbf{x}))$, and then utilize some conventional techniques to estimate the tensor distribution information. In TIMER, we have used two common statistical measures – mean and variance.

But it is not necessary to stop at the means and variances. From these means and variances, we can always further derive some other features. For example, we can first utilize the mean to measure the distribution of a neighborhood of voxels in the log-space. A matrix exponential operation can then be performed to transform the mean $\mathbf{M}(\mathbf{x})$ back to the original tensor space, obtaining $\mathbf{M}(\mathbf{x})$, and from this tensor we can compute a series of geometric features, which can include FA, ADC, and also prolateness, oblateness, and sphericity measures as suggested in (Westin et al., 1997, 2002).

Regional mean

Utilizing Log-Euclidean metrics, we can define the regional mean in a neighborhood $N(\mathbf{x})$ of voxel \mathbf{x} as:

$$\mathbf{M}(\mathbf{x}) = \exp \left[\frac{\sum_{\mathbf{t} \in N(\mathbf{x})} \log(\mathbf{D}(\mathbf{t}))}{|N(\mathbf{x})|} \right] \quad (3)$$

where $|N(\mathbf{x})|$ is the cardinality of set $N(\mathbf{x})$. From the above mean, we can compute the principal diffusivities, i.e., the eigenvalues, as:

$$\lambda_1^{(M)}(\mathbf{x}) \geq \lambda_2^{(M)}(\mathbf{x}) \geq \lambda_3^{(M)}(\mathbf{x}) \quad (4)$$

where $\lambda_k^{(M)}(\mathbf{x})$ represents the k -th largest eigenvalue of matrix $\mathbf{M}(\mathbf{x})$.

Regional variance

Similarly, we can define the regional variance as:

$$\underline{\mathbf{V}}(\mathbf{x}) = \left[\sum_{\mathbf{t} \in N(\mathbf{x})} \frac{[\log(D(\mathbf{t})) - \log(\mathbf{M}(\mathbf{x}))]^2}{|N(\mathbf{x})|} \right] \quad (5)$$

and the principal variabilities as:

$$\lambda_1^{(\underline{\mathbf{V}})}(\mathbf{x}) \geq \lambda_2^{(\underline{\mathbf{V}})}(\mathbf{x}) \geq \lambda_3^{(\underline{\mathbf{V}})}(\mathbf{x}). \quad (6)$$

We scale the eigenvalues above according to the equation below to yield their mean normalized values:

$$\tilde{\lambda}_i^{(\underline{\mathbf{V}})}(\mathbf{x}) = \lambda_i^{(\underline{\mathbf{V}})}(\mathbf{x}) / \sum_{k=1}^3 \lambda_k^{(\underline{\mathbf{M}^2)}(\mathbf{x}). \quad (7)$$

We have used $\underline{\mathbf{M}}^2$ instead of $\underline{\mathbf{M}}$ to match the dynamic range of $\underline{\mathbf{V}}$. Note that tensor mean and variance by themselves are not orientation invariant since, as tensors, they retain orientation information. But by taking only their eigenvalues, we are effectively only keeping orientation invariant information for correspondence matching.

Boundary features

Edge information extracted from tissue boundaries is crucial for facilitating tensor image registration especially when boundary accuracy is concerned. In TIMER, edge information from both tensors and FA map are incorporated.

Edge detection on tensors

To better extract tissue boundaries, we propose to extend Canny edge detector to cater for diffusion tensor images. Canny edge detector (Canny, 1986) is regarded as one of the best edge detectors developed in the field. It can be used to extract maximal image gradient boundaries, and is robust to noise due to the employment of Gaussian filter to smooth out noise prior to edge detection. For fast edge detection, 3D Gaussian-based image filtering is implemented using three subsequent steps of one-dimensional (1D) Gaussian filtering along the anterior–posterior, superior–inferior and left–right directions, followed by gradient maps computation. Using these steps, edge detection can be accomplished rapidly and robustly. Note that in TIMER, direct tensor edge detection is performed in the logarithmic space. Log-Euclidean metrics (Arsigny et al., 2006) make this especially convenient since the operations involved are as simple as that in the classical Euclidean framework, with addition/subtraction and scaling operations properly defined. Tensor convolution with function $g(\mathbf{x})$ is defined as:

$$[\log(\mathbf{D}) * g](\mathbf{x}) = \left\{ \left[[\log(\mathbf{D})]_{ij} * g \right](\mathbf{x}) \mid i, j = 1, 2, 3 \right\}. \quad (8)$$

With tensor convolution defined, operations such as Gaussian filtering and edge detection can be performed directly on the tensors. For each voxel in the volume, a gradient $\mathbf{H}^{\text{Tensor}}(\mathbf{x})$ is computed, and from which, after non-maximum suppression, we can obtain a final edge magnitude $H^{\text{Tensor}}(\mathbf{x})$.

Edge detection on FA map

Fractional anisotropy is an important measure of the degree of anisotropy. By performing edge detection on the FA map, one can detect edges formed by highly anisotropic constituents of the brain against those which are less anisotropic. In TIMER, the edge information furnished by the FA map delineates the white matter. More details on how FA can be computed are shown in the [Fractional anisotropy \(FA\)](#) section. Edges from tensors and edges from FA map

are complementary to each other and, by using both, potentially all major kinds of tissue boundaries, that is, those formed between white matter (WM), gray matter (GM) and cerebro-spinal fluid (CSF), can be detected and aligned in the registration. We denote the edge magnitude returned by the FA map at point \mathbf{x} as $H^{\text{FA}}(\mathbf{x})$.

Tensor geometric features

Tensor \mathbf{D} can be decomposed into the form:

$$\mathbf{D} = \lambda_1 \hat{\mathbf{e}}_1 \hat{\mathbf{e}}_1^T + \lambda_2 \hat{\mathbf{e}}_2 \hat{\mathbf{e}}_2^T + \lambda_3 \hat{\mathbf{e}}_3 \hat{\mathbf{e}}_3^T. \quad (9)$$

where $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and $\hat{\mathbf{e}}_i$ is the normalized eigenvector corresponding to λ_i . Geometrically, diffusion tensor \mathbf{D} can be represented by an ellipsoid with three axes oriented in the directions of its eigenvectors and the semi-axis length proportional to the square root of its eigenvalues. Three different ellipsoidal shapes give rise to three geometrical structures for the diffusion tensors: prolate (linear) structure, where diffusion is mainly in the direction corresponding to $\hat{\mathbf{e}}_1$; oblate (planar) structure, where diffusion is restricted to a plane spanned by $\hat{\mathbf{e}}_1$ and $\hat{\mathbf{e}}_2$; and spherical structure with isotropic diffusion. In TIMER, we have incorporated fractional anisotropy (FA) and principal diffusivities (PD) (Basser and Pierpaoli, 1996) as part of the attribute vector to capture the tensor geometric properties, giving us microstructural information of the brain.

Fractional anisotropy (FA)

Fractional anisotropy (FA) (Basser and Pierpaoli, 1996) is defined as:

$$\text{FA}(\mathbf{x}) = \sqrt{\frac{3}{2} \left\{ \frac{\sum_{k=1}^3 [\lambda_k^{(\mathbf{D})}(\mathbf{x}) - \bar{\lambda}^{(\mathbf{D})}(\mathbf{x})]^2}{\sum_{k=1}^3 [\lambda_k^{(\mathbf{D})}(\mathbf{x})]^2} \right\}} \quad (10)$$

where $\bar{\lambda}^{(\mathbf{D})}(\mathbf{x})$ is the mean diffusivity of $\mathbf{D}(\mathbf{x})$, defined as: $\text{trace}(\mathbf{D}(\mathbf{x}))/3$. FA measures the fraction of the anisotropic portion of tensor $\mathbf{D}(\mathbf{x})$. For an isotropic medium, $\text{FA} = 0$; for a fully anisotropic medium, FA is 1. Its value is characteristically high for a voxel containing white matter.

Principal diffusivities (PDs)

Principal Diffusivities (PDs) are simply the eigenvalues of $\mathbf{D}(\mathbf{x})$:

$$\lambda_1^{(\mathbf{D})}(\mathbf{x}) \geq \lambda_2^{(\mathbf{D})}(\mathbf{x}) \geq \lambda_3^{(\mathbf{D})}(\mathbf{x}). \quad (11)$$

They are measures of diffusivities in the principal directions. Some regions of the brain, such as the ventricles, contain mostly CSF and microstructural barriers to water mobility are sparse in these regions, although a few membranes may be present, and as a result PDs typically have high values. Owing to the inherently anisotropic nature of WM, one is able to roughly separate WM and non-WM tissues utilizing the FA map. By the same token, PD values of CSF are much larger than those of other tissues, so CSF and non-CSF tissues can be roughly separated.

Driving voxels

Brain medical images are inherently high dimensional and computations involved in image registration can be prohibitive. In order to overcome this, a small subset of voxels with distinctive attributive vectors, called *driving voxels* (Shen and Davatzikos, 2002), are initially selected to drive the registration, obtaining a coarse estimation of the registration. More driving voxels are included as the registration progresses to participate in refining the registration. TIMER selects, as driving voxels, a combination of voxels with the highest tensor edge magnitudes, FA map edge magnitudes and FA




Fig. 4. Driving voxels superimposed on the FA image. Shown on the left are the initial driving voxels and on the right the additional driving voxels, in red, when the registration progresses.

values, since these voxels represent important anatomical structures and are typically relatively easy to find in images with sufficient contrast. The driving voxels for different stages of registration are shown in Fig. 4.

Summary of TIMER

A flowchart summarizing the working mechanism of TIMER is shown in Fig. 5. A more detailed description is provided in the following.

- 1. Histogram Matching and Affine Registration:** Transform the subject DT image to the template space using global affine transformation parameters which are computed based on the FA maps. Reorient the tensors (Xu et al., 2003) of the transformed subject image.
- 2. Downsampling:** Downsample the DT images to an appropriate scale. Deformation field determined in this scale will be used to drive the deformation in a subsequent higher scale.
- 3. Compute Attribute Vectors:** Compute an attribute vector for each voxel in the template and subject DT images, consisting of tensor edge magnitude, FA map edge magnitude, FA value, principal diffusivities, tensor mean principal diffusivities, and also principal variabilities.
- 4. Select Driving Voxels:** Determine a set of driving voxels based on tensor edge magnitudes, FA map edge magnitudes, and also FA values.
- 5. Locate Correspondence:** For each template driving voxel, search in its neighborhood for all subject voxels with similar attribute vectors, which is determined by a threshold which decreases with time. Tentatively deform the subvolume centered at the current template driving voxel, weighted by a Gaussian kernel, to every similar voxel found in the subject and obtain an overall degree of similarity for the deformed subvolume. This template subvolume will be deformed to the subject voxel with the largest degree of similarity, provided that the degree of similarity is over a certain predefined threshold. Note that the principle of consistent transformation (Christensen, 1999; Shen and Davatzikos, 2002) is observed in this process. For each subject driving voxel, search its neighborhood for a template driving voxel with the most similar attribute vector. If this similarity is above a threshold, which decreases with time, then the subvolume will also be deformed in the direction of the subject. If deformation field from the previous

scale exists, load the deformation field for the guidance of the current cycle of correspondence matching.

- 6. Regularize Deformation Field:** Refine deformation field utilizing smoothness constraints (Shen and Davatzikos, 2002).
- 7. Upsample Deformation Field:** Upsample the obtained deformation field and utilize it to guide the correspondence matching of the subsequent higher scale.
- 8. Morph Subject DT Image and Reorient Tensors:** If all the scales have been considered, morph the subject DT image to the template space according to the determined deformation field.

The HAMMER (Shen and Davatzikos, 2002) algorithm should be referred to if more information on step 4 to step 8 is needed.

Evaluation of TIMER

A number of experiments were performed to evaluate the accuracy of TIMER. We started off by computing a group-averaged image, which could essentially give us a visual indication of the accuracy of TIMER. This was done by first registering the diffusion tensor images of different subjects onto a randomly chosen template and then computing the average image of all the registered images. A sharp average image indicates good registration accuracy. To further evaluate the accuracy of TIMER, we simulated a set of deformation fields, which we used as ground truths. Based on these deformation fields, the respective simulated brain images could be constructed, and after registering them back onto the template, we could compare the deformation fields estimated by TIMER with the ground truths and quantify the registration accuracy. To gauge the effectiveness of TIMER in aligning the brain microstructures, especially those in the white matter, a fiber tracking experiment was performed, where the consistency of the fiber bundles in the simulated brains were compared with those in the template. We also performed an experiment using simulated brain atrophy to gauge whether an atrophic region remains detectable after registration. Lastly, we performed a noise sensitivity study to evaluate how well TIMER could perform under different signal-to-noise ratios (SNRs). Whenever appropriate, results obtained using DTI registration algorithms proposed by Yang et al. (2008a) and Zhang et al. (2006) will be included for comparison, and they will be referred to as YANG and ZHANG¹ respectively in rest of the paper.

¹ The version used in our experiments is an extension of (Zhang et al., 2006), which can handle relatively larger deformations, as mentioned in (Zhang et al., 2007).

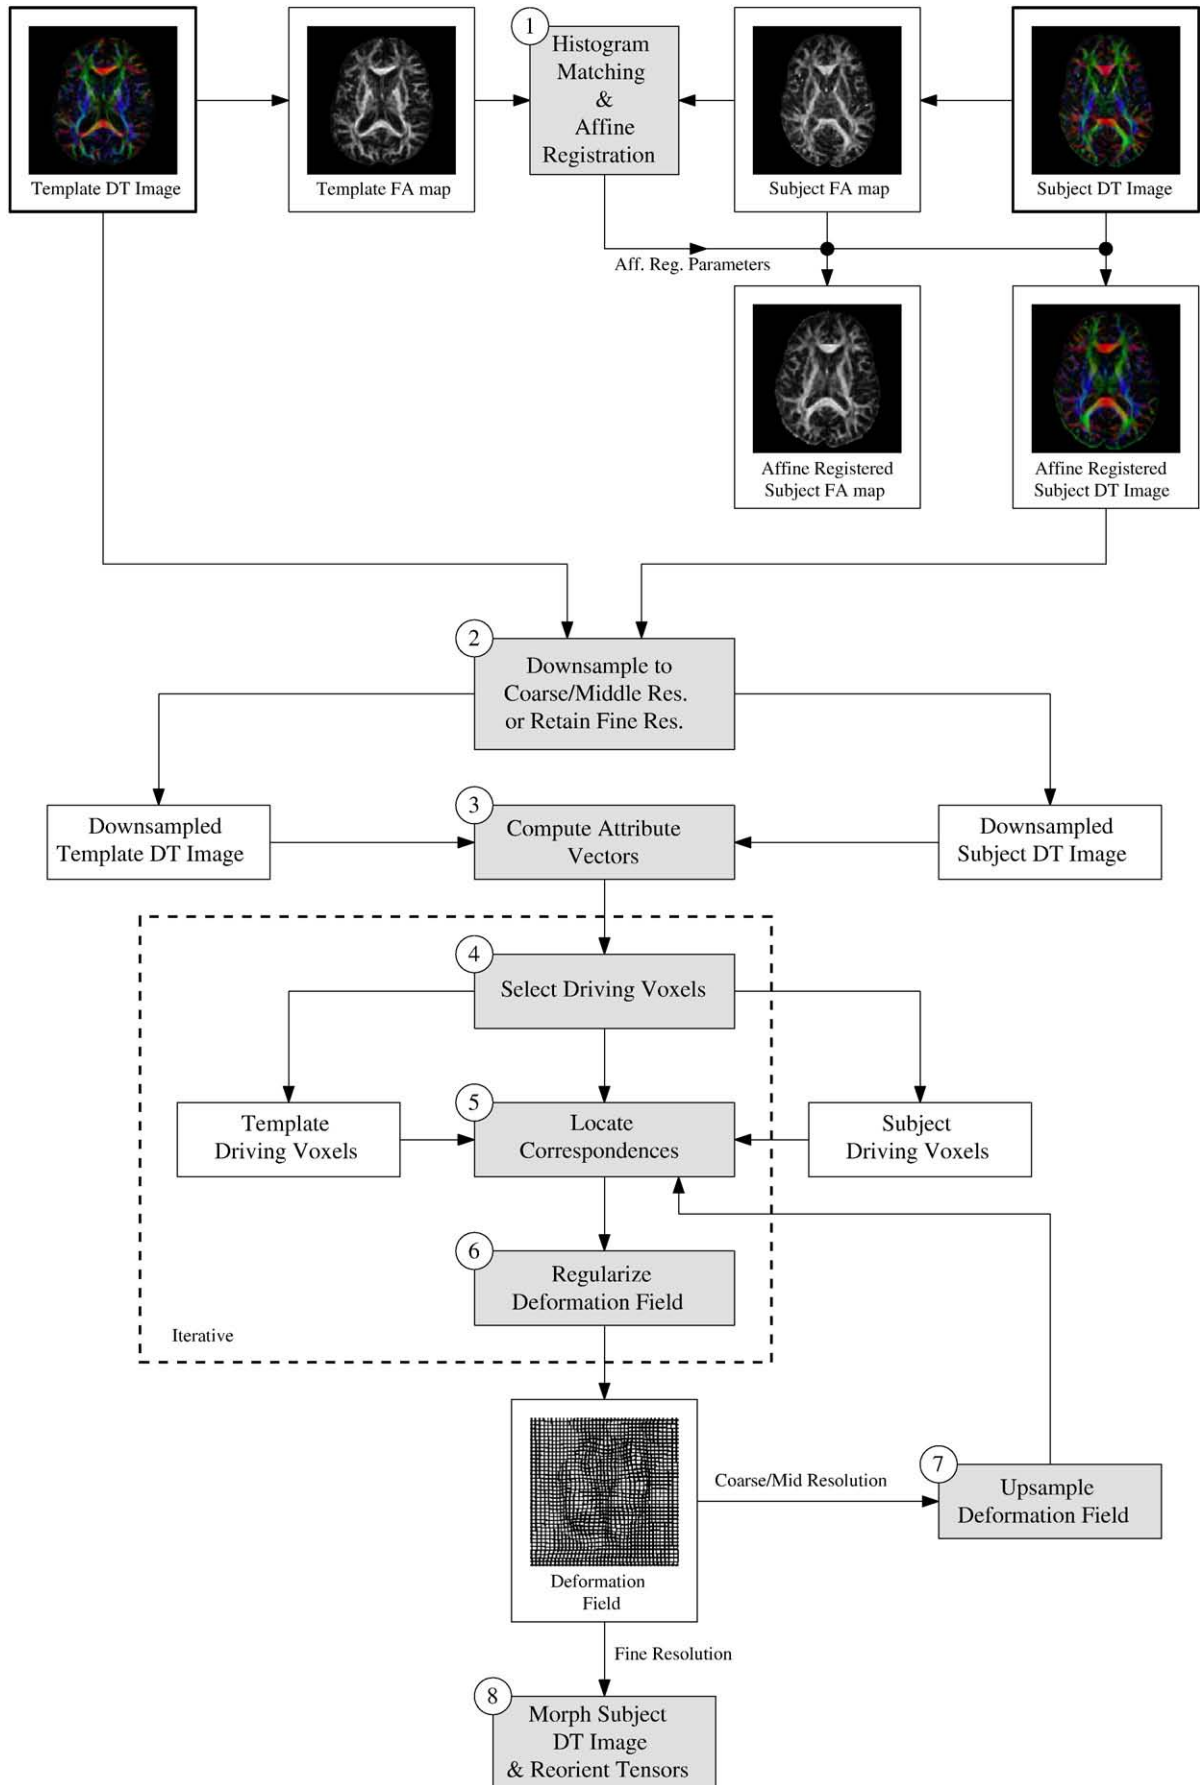


Fig. 5. A flowchart summarizing the main components of TIMER. The numbers in the circles refer to the corresponding text in the Summary of TIMER section.

Real subjects

A direct and intuitive way of gauging the accuracy of a registration algorithm is by inspecting the sharpness of an average image obtained from a set of aligned images. Inaccurate registration often results in the wrong alignment of anatomical structures, and when averaged, will result in fuzziness. To evaluate TIMER, we took a similar approach where one subject was selected as the template from the dataset, and the other 21 subjects were then registered onto this template to form the average image. For easier inspection, a standard deviation map was generated. The standard deviation value at each voxel location gives us information on the extent of tensor deviation of the aligned images from the average image. For this purpose, we employed the tensor normalized standard deviation, defined by Jones et al. (2002) as:

$$\frac{1}{\|\bar{\mathbf{D}}\|_F} \sqrt{\frac{1}{N-1} \sum_{k=1}^N \|\mathbf{D}_k - \bar{\mathbf{D}}\|_F^2} \quad (12)$$

where $\|\cdot\|_F$ is the Frobenius norm, and $\bar{\mathbf{D}}$ the tensor mean, computed by element-wise averaging, of the same voxel location from a set of N diffusion tensor images, and \mathbf{D}_k is the tensor from the k th image. Another useful metric for alignment evaluation is the normalized scalar product applied to the FA map of each individual registered image with respect to the template. Given the FA maps of the template, $I(\mathbf{x})$, and the subject, $J(\mathbf{x})$, the normalized scalar is defined as:

$$S = \frac{\sum_{\mathbf{x} \in V} I(\mathbf{x}) \cdot J(\mathbf{x} + \mathbf{u}(\mathbf{x}))}{[\sum_{\mathbf{x} \in V} I^2(\mathbf{x}) \cdot \sum_{\mathbf{x} \in V} J^2(\mathbf{x} + \mathbf{u}(\mathbf{x}))]^{\frac{1}{2}}} \quad (13)$$

where $\mathbf{u}(\mathbf{x})$ is the subject-to-template deformation field and V denotes the brain image volume.

Simulated subjects

To further evaluate the accuracy of TIMER, we generated a set of simulated deformation fields using the statistical model of deformation (SMD) proposed by Xue et al. (2006). In this approach, Wavelet-Packet Transform (WPT) of the training deformations and their Jacobians, in conjunction with a Markov random field (MRF) spatial regularization, are used to capture both coarse and fine characteristics of the training deformations in a statistical fashion. Simulated deformations can then be generated by randomly sampling the resultant statistical distribution (Xue et al., 2006). One human brain was utilized as the template and the 20 simulated deformation fields, which also served as the ground truths, were applied to the template, resulting in 20 simulated human brains. These 20 simulated brains were then registered back onto the template using TIMER and the deformation fields estimated by the registration were compared with the ground truths. The accuracy of the estimated deformation fields were evaluated by the average difference \bar{m}_d and the standard deviation (SD) σ_d of the two deformation fields, i.e. (Yang et al., 2008b):

$$\bar{m}_d = \frac{1}{|V|} \sum_{\mathbf{x} \in V} \|\mathbf{u}_g(\mathbf{x}) - \mathbf{u}_e(\mathbf{x})\| \quad (14)$$

and

$$\sigma_d = \sqrt{\frac{1}{|V|} \sum_{\mathbf{x} \in V} (\|\mathbf{u}_g(\mathbf{x}) - \mathbf{u}_e(\mathbf{x})\| - \bar{m}_d)^2} \quad (15)$$

where $\|\cdot\|$ is the Euclidean distance, $|V|$ the number of voxels in the brain image volume V , and $\mathbf{u}_g(\mathbf{x})$, $\mathbf{u}_e(\mathbf{x})$ the template-to-subject deformation vectors at location \mathbf{x} of the ground truth deformation field and the estimated deformation field, respectively.




Fig. 6. Group-averaged images resulting from the registration of the 21 subjects. The FA weighted first principal directions are shown in their color coded representations: green for the anterior–posterior direction, blue for the superior–inferior direction, and red for the left–right direction. The tensors in the yellow boxes are shown in their FA weighted ellipsoidal representations in the bottom panels.




Fig. 7. White matter edges generated from the FA map of the average tensor image is overlaid on the template. The proper alignment of the edges is indicative of good registration accuracy.

Fiber tracking

Local diffusion patterns characterized by the restricted motion of cellular fluid within brain white matter gives an estimation of the

orientation of the underlying fibers. The potential of DTI, capitalizing on this fact, to reveal white matter integrity, disruption and pathology makes it the preferred modality for studying white matter diseases. Accurate DTI registration is essential to establish inter-subject homology for further statistical inference. Using a tractography method known as *Fiber Assignment by Continuous Tracking* (FACT) (Mori et al., 1999; Xue et al., 1999), fiber bundles passing through some regions of interest (ROIs) were tracked, extracted, and compared for quantifying how well TIMER performs in these specific ROIs. Fiber tracking was performed in DtiStudio (Jiang et al., 2006), with the FA start tracking threshold set to 0.25 and the FA stopping threshold set to 0.2. A turning angle of more than 70° also stops the tracking. The same fiber tracking procedure was performed on the TIMER registered simulated brains, their group-averaged mean and also the template itself. By comparing the similarity of the fiber tract bundles, we could gauge the effectiveness of TIMER in registering the images. The distance of two fiber bundles was measured in a way similar to that used in (Zhang et al., 2006). The distance measure is defined as:

$$\frac{1}{|\mathcal{F}| + |\mathcal{G}|} \left[\sum_{F_i \in \mathcal{F}} \min_{G_j \in \mathcal{G}} d(F_i, G_j) + \sum_{G_i \in \mathcal{G}} \min_{F_j \in \mathcal{F}} d(F_j, G_i) \right] \quad (16)$$

where $d(F_i, G_j)$ is a pairwise distance between fibers \mathcal{F} and \mathcal{G} , which in our case, is defined as the mean of the closest distance for every point




Fig. 8. Normalized standard deviation. The images in the top row are slices of the normalized standard deviation map generated from the TIMER registered images; the images in the bottom row are those from affine registration.

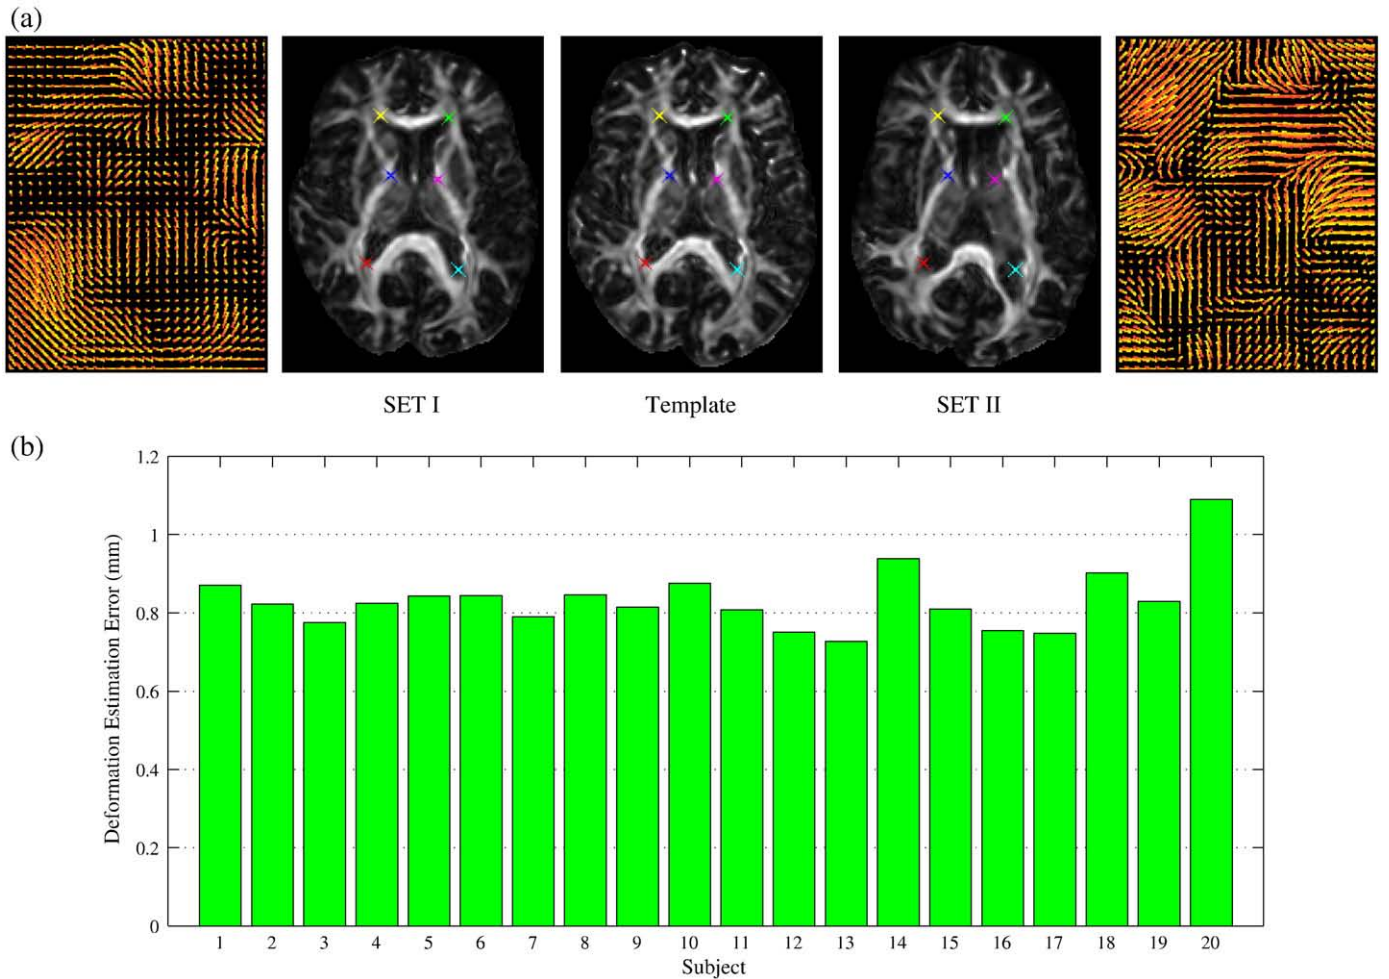


Fig. 9. (a) Difference between SET I and SET II. Middle: Template. Middle left and left: an image in Set I with its template-to-subject deformation field; Middle right and right: an image in Set II with its deformation field. The deformation fields of images in Set II are characteristically more complex and larger in magnitudes compared with the more conservative deformation fields of Set I. In the different panels, markers with the same color are placed at the same locations to facilitate comparison. (b) Deformation field errors of TIMER for SET I. The average deformation field error of the 20 simulated brains is 0.83 mm with a standard deviation of 0.48 mm.

of two fibers (Gerig et al., 2004). When two fiber bundles are perfectly aligned, the value returned is zero. This measure approximately establishes anatomic correspondences between points along fibers of different subjects.

Atrophy detection

We simulated hypothetical abnormality, affecting the tensors in a specific brain region, and gauge the detectability of the abnormal region after registration. We simulated atrophic brains based on the healthy brains that we had created in the *Simulated subjects* section and, after registering them onto the template, we wanted to determine whether the atrophic region was still detectable, gauged using *t*-test. We first defined a region of interest on a main fiber pathway in the template image, and we could then map this region of interest onto the simulated brains based on their respective simulated deformation fields. These regions on the simulated brains were then perturbed with a hypothetical atrophy, which effectively made the tensors more isotropic (rounder and less elongated ellipsoids) and hence had smaller FA values. Let \mathbf{D} be the tensor with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and eigenvectors \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 . For a degradation percentage η , we simulated the atrophy as follows:

- Step 1: Let $\lambda = \max(\lambda_1, \lambda_2, \lambda_3)$
- Step 2: Set $\lambda_i = (1 - \eta)\lambda_i + \eta\lambda$, $\forall i$

- Step 3: Set

$$\mathbf{D}' = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3) \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)^T$$

- Step 4: Randomly perturb the tensor angle with an orthogonal rotation matrix² \mathbf{R} by a maximum value of θ_{\max} :

$$\mathbf{D}'' = \mathbf{R}\mathbf{D}'\mathbf{R}^T. \quad (17)$$

Directions of anisotropic tensors are perturbed more than the isotropic ones.

Obviously, at $\eta = 1.0$ the resultant tensor \mathbf{D}' will be perfectly isotropic, the FA value of which is 0; at $\eta = 0.0$, the original eigenvalues are retained. By performing paired *t*-tests on the FA maps of the healthy brains and those with simulated atrophy, we could gauge how well the detectability of the atrophic regions was retained after registration with TIMER. Detected regions with a false discovery rate (FDR) of 0.05 was retained and deemed significant. The Dice ratio (Dice, 1945) was computed for the detected region with respect to the ground truth atrophic region, a higher value of which indicates a more

² See Appendix for explanation on how a random rotation matrix can be generated.

Table 1

Averages and standard deviations of deformation estimation errors (mm).

		Whole brain			Cortical region		
		TIMER	YANG	ZHANG	TIMER	YANG	ZHANG
SET I	Mean	0.83	1.36	0.84	0.86	1.44	0.92
	Std. deviation	0.48	0.82	0.76	0.47	0.82	0.81
SET II	Mean	1.11	1.87	1.44	1.18	2.03	1.68
	Std. deviation	0.82	1.48	1.75	0.88	1.59	1.96

accurate match of the detected region and hence less false positives and/or false negatives.

Noise sensitivity

Due to the signal-attenuation nature of diffusion encoding and the inadvertent T_2 relaxation resulting from the long TE necessary to accommodate the sensitization gradient pulses, the signal-to-noise ratio (SNR) of DTI is inherently low (Chen and Hsu, 2005). The

simplest means of improving the effective SNR is by acquiring more signals for averaging, using more diffusion sensitization directions, encoding levels, or a combination of these. Regardless of the strategy, the scan time is necessarily and proportionally lengthened. For these reasons, it is important that algorithms developed for diffusion tensor images are robust to noise and are able to perform satisfactorily in low SNR conditions. A simple SNR sensitivity evaluation of a registration algorithm can be carried out by generating some noisy images, register them onto the template, and compare their registration accuracy with respect to their noise free counterparts. In order to ensure that, after adding noise, the diffusion tensors remain symmetric positive definite (SPD), we first decompose the diffusion tensors using Cholesky decomposition (Bergmann et al., 2007):

$$\mathbf{D} = \mathbf{L}\mathbf{L}^T \tag{18}$$

where \mathbf{L} is a lower triangular matrix with strictly positive diagonal entries, and \mathbf{L}^T denotes the transpose of \mathbf{L} . Adding noise on the

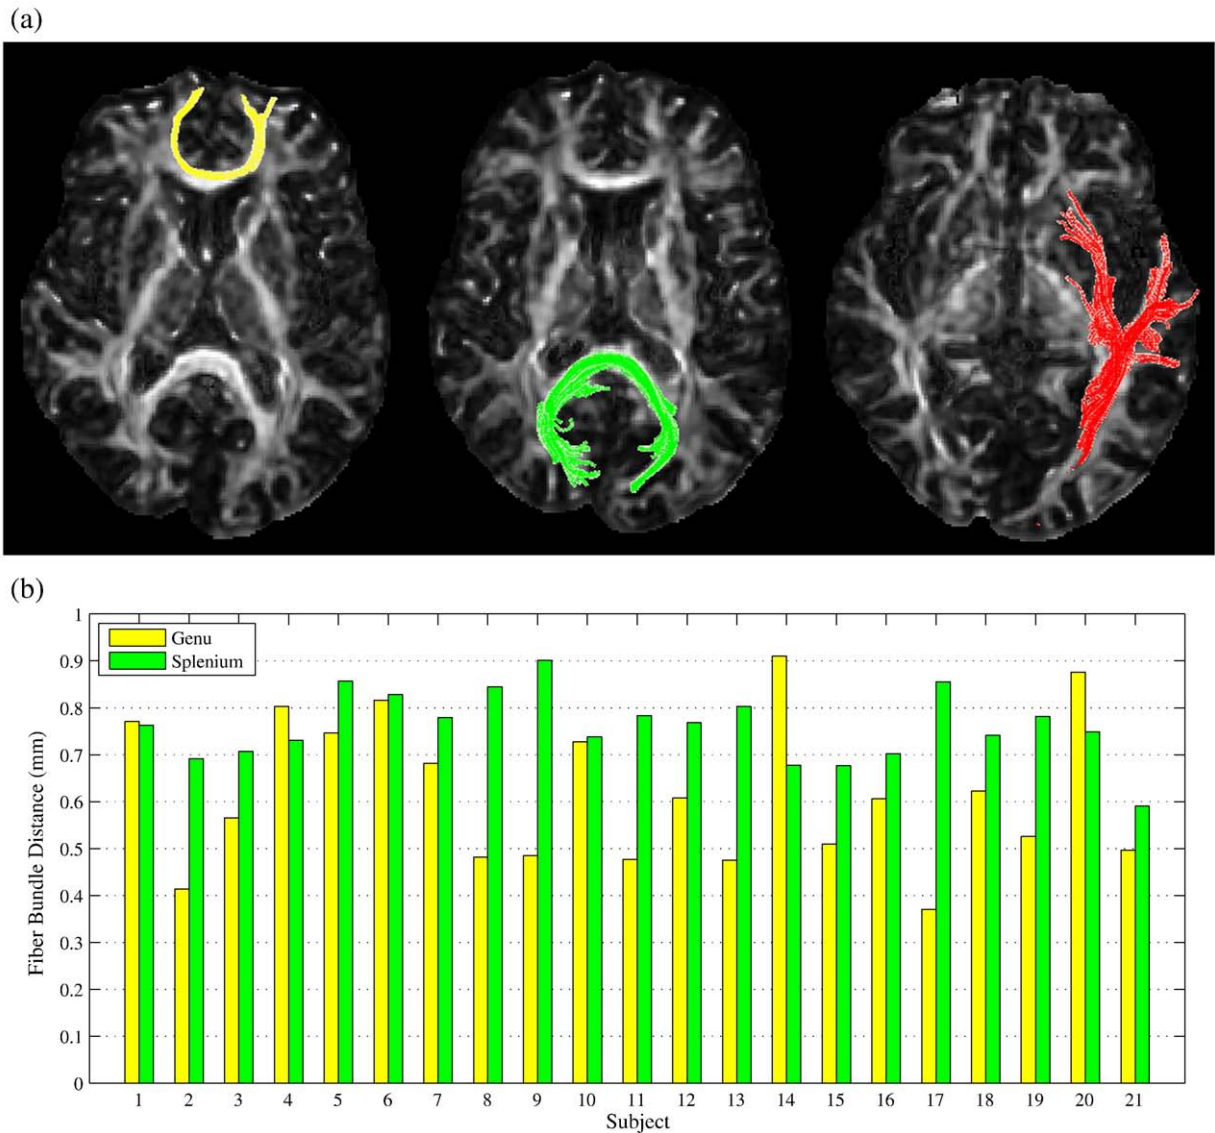


Fig. 10. (a) Template fiber bundles used as the ground truths for evaluation. The fiber bundles shown in yellow and green reside in the genu and the splenium of the corpus callosum, respectively. In red is a fiber bundle near the cortical surface. (b) Distances of fiber bundles extracted from the simulated brains (SET I) with respect to the template. The yellow and green bars show the results for the genu and the splenium fiber bundles, respectively. Subject 21 is the group-averaged mean. The average errors for the genu and the splenium fiber bundles are 0.62 mm and 0.77 mm, respectively. Results from SET II (red fibers in (a)) are not shown in this graph as their higher fiber distance values would overwhelm those yielded by SET I. Moreover, SET I and SET II consist of different simulated subjects.

Table 2
Average fiber distances (mm).

		TIMER	YANG	ZHANG
SET I	Genu	0.62	0.78	0.67
	Splenium	0.77	0.82	1.00
SET II	Cortical	1.25	1.49	1.91

elements of \mathbf{L} , forming $\tilde{\mathbf{L}}$, we can obtain a SPD noisy diffusion tensor $\mathbf{D} = \tilde{\mathbf{L}}\tilde{\mathbf{L}}^T$.

Results

Real subjects

The average images generated by registering 21 subjects onto the template are shown in Fig. 6. It can be observed that for FA map based affine registration, the average image, shown in Fig. 6(b), is fuzzy especially in areas near the cortical surface. In comparison, TIMER yields an average image, shown in Fig. 6(c), which shows much improved sharpness. By superimposing the edges of the FA map of the average image on the FA map of the template, as shown in Fig. 7, we can see that the white matter edges of the images are properly aligned, indicative of good registration accuracy. For convenient visual inspection, the normalized standard deviation maps are shown in Fig. 8. The images yielded by TIMER are generally darker in both cortical and subcortical regions, implying less registration variability and closer resemblance of tensors at corresponding locations of different subjects. The scalar product values, averaged over all 21 registered images, yielded by TIMER, YANG, ZHANG and affine registration are 0.9126, 0.9094, 0.8894, and 0.8293, respectively. The relative high value given by TIMER once again validate its performance superiority compared to other methods.

Simulated subjects

We generated two sets (SET I & SET II) of simulated deformation fields. SET I consisted of 20 deformation fields, which were more regular and generally had smaller displacement magnitudes, whereas SET II contained the same number of deformation fields, but were more complex and had larger magnitudes, generated using a larger variance in SMD (Xue et al., 2006). Fig. 9(a) shows some typical deformation fields and images from SET I and SET II. The deformation estimation errors for SET I are shown in Fig. 9(b). The average error for the 20 simulated brains is 0.83 mm, which indicates that TIMER is accurate to a subvoxel level. A summary of the results

Table 3
Average Dice ratios of detected regions relative to the ground truth atrophic region.

η	TIMER	YANG	ZHANG	AFFINE
0.05	0.81	0.61	0.41	0.00
0.10	0.89	0.75	0.35	0.00
0.40	0.83	0.59	0.25	0.00

for the whole brain and also the cortical region in comparison with YANG and ZHANG is provided in Table 1. For SET I, the accuracy yielded by TIMER is closely followed by that of ZHANG and is better than YANG by a larger margin. But for SET II, the improvement brought by TIMER is clearly quite significant. For all cases, TIMER is also more consistent as indicated by the smaller standard deviation values. Similar conclusion can be drawn from the results of the cortical region.

Fiber tracking

We present here two sets of results. For SET I, two ROIs were selected so that two fiber bundles, one residing in the genu and the other in the splenium of the corpus callosum (CC), could be extracted for comparison. For SET II, we evaluated TIMER in a more difficult situation where a fiber bundle near the cortical surface was considered. Registration of near cortical surface regions is often the Achilles' heel of many registration algorithms. All the fiber bundles are shown in Fig. 10(a). Detailed results for TIMER for the genu and splenium fiber bundles are shown in Fig. 10(b). A summary of the results of all fiber bundles, with those of YANG and ZHANG included, is shown in Table 2. Due to the inherent nature of fiber tracking, accurate registration of the fiber bundles are not only determined by the correct estimation of the deformation field, as was examined in the previous experiment, but also by the correct reorientation of tensors. TIMER incorporates a tensor reorientation scheme which attempts to estimate the underlying fiber orientation by considering an appropriate small neighborhood around each voxel (Xu et al., 2003), and is hence less vulnerable to noise. This is attested by the results yielded by TIMER in comparison with YANG and ZHANG, where fiber bundles extracted from images registered with TIMER are consistently closer to the ground truths.

Atrophy detection

We used $\eta=0.05,0.10,0.40$ and $\theta_{\max}=10^\circ$ to generate 3×20 simulated atrophic brains using SET I. Fig. 11 shows the region with simulated atrophy and their ellipsoidal representation images before and after the introduction of atrophy. The atrophy detection results




Fig. 11. Simulated atrophy. The content in the green box (a) is enlarged with its ellipsoidal representation images before and after the introduction of atrophy shown in (b) and (c), respectively. The tensors are fatter when atrophic. The changes in directions can be noticed from the change of colors.

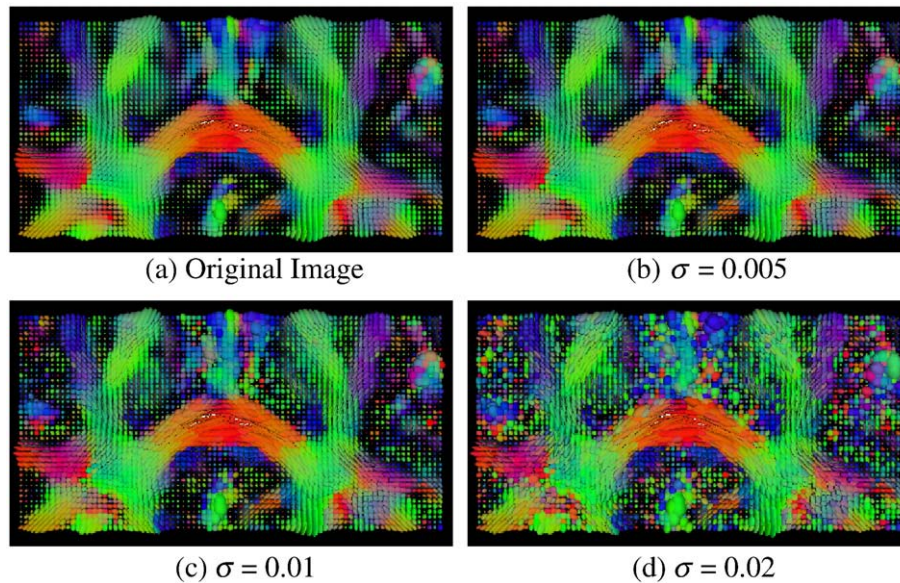


Fig. 12. Noise degraded images. The original image, on the top left, followed by images with increasing degrees of noise ($\sigma=0.005, 0.01, \text{ and } 0.02$).

are shown in Table 3. It is worth noting that in order to maintain the detection specificity of the atrophic region, i.e. high Dice ratio, the registration algorithm used, while maintaining the detectability of the atrophic region, should not be overtly sensitive to the changes brought about by the abnormalities. If not so, the registration of the surrounding non-atrophic regions will be affected and will hence result in low specificity due to false positives. This problem is mitigated by the multiscale mechanism in TIMER, the coarse scale features of which effectively help avoid the aforementioned pitfall by considering a regional neighborhood which is often larger than the atrophic region. TIMER therefore has a less tendency of being influenced by the atrophic region itself and propagating the abnormalities to other non-atrophic regions in the course of registration. And it is thus not surprising here that the Dice ratios yielded by TIMER are consistently higher. On average, TIMER yields a 0.2 unit improvement over YANG and 0.5 unit over ZHANG. The affine registered images fail to capture the subtlety of the atrophy, rendering the atrophic region undetectable.

Noise sensitivity

We tested the SNR sensitivity of TIMER by measuring the decrease of registration accuracy (via measuring the deformation estimation error), using 10 randomly selected subjects from SET I and SET II. Additive Gaussian noise of standard deviations, $\sigma=0.005, 0.01, 0.02$, is used. Some sample images are shown in Fig. 12. The results indicate that TIMER is robust to noise, as validated by the low percentages of accuracy decreases — 0.0076%, 0.60%, and 1.33%, respectively.

Discussion

In medical imaging, it is important to build deformable anatomical models that take into account the underlying anatomy, and not simply the similarity of image intensities. To this end, TIMER attaches a multiscale attribute vector to each voxel, reflecting its underlying anatomical structure in a local scale, and also its relationship to more distant voxels in a more global scale. A rich enough attribute vector can potentially differentiate different parts of the anatomy that would otherwise look similar. Relative high variability of brain structures makes certain features, especially

those capturing detailed features, vary dramatically across individual brains, thus confounding the image matching procedure. It is hence desirable that features used for image matching be robust to structural variations across individuals. Since orientation variation can be part of the confounding factor in the matching process, features used should also preferably be invariant to rotation. A number of such features can be collected to form attribute vectors for correspondence matching in registration. Attribute vectors are a critical part of TIMER. An attribute vector is defined for each voxel in a volumetric image, and it reflects the underlying structure at multiple scales. Each attribute vector in TIMER characterizes the voxel in terms of its regional tensor mean, regional tensor variance, diffusivity, tensor edge magnitude, FA map edge magnitude and FA value. As can be attested by the experiments, attribute vector as such is sufficiently rich for registration guidance. The attribute vectors are used to morph the subject onto the template using a deformable scheme similar to that of HAMMER (Shen and Davatzikos, 2002). After registration, tensor at each voxel is reoriented using the Procrustean estimation approach proposed in (Xu et al., 2003). Registration of images with dimension $256 \times 256 \times 70$ takes approximately 90 min on a 2.66 GHz Linux machine. Images with smaller dimension, e.g., $128 \times 128 \times 46$, take around 30 min.

Tract-based spatial statistics (TBSS)

In their TBSS paper (Smith et al., 2006), Smith et al. argue that, to sufficiently align images for subsequent statistical analysis, a non-linear alignment with *intermediate* degrees of freedom (DoF) has to be used. The main concern is that while low-DoF barely guarantees the proper alignment of structural details, high-DoF jeopardizes structural topology, causing biologically incorrect results. The formulation of TIMER ascertain that the pitfalls at both DoF extremes can be avoided. While the correspondence matching mechanism of TIMER is inherently high-DoF, it is regularized by a smoothness constraint which ensures proper topological preservation. In fact, it is perceivable that TIMER can easily be incorporated in the framework of TBSS for effective analysis of tract statistics. TIMER can be employed to non-linearly register the images onto a selected template, form the mean

Table 4
Improvement of TIMER over YANG and ZHANG in terms of paired *t*-test *p*-values.

	YANG	ZHANG
Real subjects* (Real subjects section)	$p = 4.92 \times 10^{-4}$	$p = 3.54 \times 10^{-14}$
Simulated subjects** (Simulated subjects section)	$p = 8.69 \times 10^{-25}$	$p = 2.98 \times 10^{-5}$
Fiber tracking*** (Fiber tracking section)	$p = 4.57 \times 10^{-6}$	$p = 2.68 \times 10^{-8}$

*All scalar product values, **All estimated deformation field errors, ***All fiber bundle distances.

diffusion tensor image and its FA map, and generate the tract skeleton. After projecting the individual subjects' FA values onto the skeleton, the spatially aligned FA values can be employed for voxelwise cross-subject statistical analysis.

Brain tumors

Existing brain image registration algorithms are generally unable to account for topological differences between a brain with tumor and a normal brain. The effects of edema, tissue death and resorption, tissue infiltration, tumor mass-effect, cause severe deformation atypical of natural intersubject variability to which these algorithms are tuned. As a future work, TIMER can be tuned to account for abnormality as such. One possible approach is by adopting the method proposed by Mohamed et al. (2006), where a biomechanical three-dimensional model is employed to simulate the deformation caused by the tumor bulk and the peri-tumor edema. The estimation of the deformation field is decomposed into two components, one representing inter-individual brain shape differences, and the other representing tumor-induced deformation. For facilitating registration, a simulation of tumor mass-effect is performed on the template image to generate an image that is similar to the image of the brain with tumor.

Concluding remarks

In conclusion, *Tensor Image Morphing for Elastic Registration* or TIMER is proposed in this paper as an accurate diffusion tensor registration algorithm. The main novelty of TIMER lies in its direct extraction of regional and edge information from the tensors to guide the registration in a hierarchical fashion. Using these features to drive the registration, it is found that subvoxel accuracies can be achieved with an average deformation field estimation accuracy of 0.83 mm. The directional structure of the tensors are also retained accurately, as can be attested by the fiber tracking experiment, in which TIMER yields an average accuracy of 0.62 mm, 0.77 mm and 1.32 mm for the genu, splenium and near cortical fiber bundles, respectively. Pathological cues such as atrophy detectability is also retained after the registration, which is indicated by the typically high Dice ratios. Table 4 provides a summary of improvements brought about by TIMER in terms of *p*-values obtained from a series of paired *t*-tests. Although TIMER performs reasonably well in all experiments performed, we feel that there is still room for improvement. One possible direction is to employ an example based learning method to learn separate sets of appropriate features for registration guidance of different brain regions. Research in this direction can be conducted with the hope of a more refined DTI registration algorithm.

Acknowledgments

The authors would like to express their gratitude to Gary Zhang of Penn Image Computing and Science Lab (PISCL) for helping them understand his algorithm better and for his graciousness in optimizing the code for the comparisons conducted in this paper. This work was supported in part by grants 1R01EB006733, 1R03EB008760 and R01EB008374, as well as the UNC RRC grant.

Appendix: Generation of 3D random rotation matrix

A rotation matrix to rotate about any arbitrary axis can be constructed using:

$$\mathbf{R} = \begin{bmatrix} tx^2 + c & txy - sz & txz + sy \\ txy + sz & ty^2 + c & tyx - sx \\ txz - sy & tyz + sx & tz^2 + c \end{bmatrix} \quad (19)$$

where

$$c = \cos \theta, \quad s = \sin \theta, \quad t = 1 - \cos \theta \quad (20)$$

and (x,y,z) is a unit vector on the axis of rotation. For a tensor \mathbf{D} with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3$ and eigenvalues \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 we can find a random axis of rotation in the plane defined by \mathbf{v}_2 and \mathbf{v}_3 , i.e.:

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \frac{a\mathbf{v}_2 + b\mathbf{v}_3}{\sqrt{a^2 + b^2}} \quad (21)$$

where a and b are random arbitrary numbers. We can then rotate the principal direction as defined by \mathbf{v}_1 using (19) by a random angle, θ , between 0 and θ_{\max} .

References

- Alexander, D.C., 2006. Visualization and processing of tensor fields. Springer Berlin Heidelberg, Ch. An Introduction to Computational Diffusion MRI: the Diffusion Tensor and Beyond. Springer Berlin, Heidelberg, pp. 83–106.
- Arsigny, V., Fillard, P., Pennec, X., Ayache, N., 2006. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magn. Reson. Med.* 56 (2), 411–421.
- Basser, P.J., Pierpaoli, C., 1996. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *J. Magn. Reson. Ser. B* 111 (3), 209–219.
- Bergmann, Ø., Christiansen, O., Lie, J., Lundervold, A., June 2009. Shape-adaptive DCT for denoising of 3D scalar and tensor valued images. *Journal of Digital Imaging* 22 (3), 297–308.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 8, 679–714.
- Chen, B., Hsu, E.W., 2005. Noise removal in magnetic resonance diffusion tensor imaging. *Magn. Reson. Med.* 54, 393–407.
- Christensen, G.E., 1999. Consistent linear-elastic transformations for image matching. *Information Processing in Medical Imaging*. Springer Berlin, Heidelberg, pp. 224–237.
- Dice, L.R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26, 297–302.
- Gerig, G., Gouttard, S., Corouge, I., 2004. Analysis of brain white matter via fiber tract modeling. *Proceedings of the IEEE EMBS. Institute of Electrical and Electronics Engineering (IEEE)*, pp. 4421–4424.
- Huang, H., Zhang, J., Wakana, S., Zhang, W., Ren, T., Richards, L., Yarowsky, P., Donohue, P., Graham, E., van Zijl, P.C.M., Mori, S., 2006. White and gray matter development in human fetal, newborn and pediatric brains. *NeuroImage* 33 (1), 27–38.
- Jiang, H., van Zijl, P.C.M., Kim, J., Pearlson, G.D., Mori, S., February 2006. DtiStudio: resource program for diffusion tensor computation and fiber bundle tracking. *Comput. Methods Programs Biomed.* 81 (2), 106–116.
- Jones, D.K., Griffin, L.D., Alexander, D.C., Catani, M., Horsfield, M.A., Howard, R., Williams, S.C.R., 2002. Spatial normalization and averaging of diffusion tensor MRI data sets. *NeuroImage* 17, 592–617.
- Lazebnik, S., Schmid, C., Ponce, J., 2005. A sparse texture representation using local affine regions. *IEEE Trans. Pattern Anal. Mach. Intell.* 27, 1265–1278.
- Lee, A., Lepore, N., Barysheva, M., Chou, Y., Brun, C., Madsen, S., McMahon, K.L., de Zubicaray, G., Meredith, M., Wright, M., Toga, A., Thompson, P., May 2008. Comparison of fractional and geodesic anisotropy in diffusion tensor images of 90 monozygotic and dizygotic twins. *ISBI'2008. Institute of Electrical and Electronics Engineering (IEEE)*, pp. 943–946.
- Liu, T., Li, H., Wong, K., Tarokh, A., Guo, L., Wong, S.T., 2007. Brain tissue segmentation based on DTI data. *NeuroImage* 38 (1), 114–123.
- Lowe, D., 2004. Distinctive image features form scale-invariant key-points. *Int. J. Comput. Vis.* 2, 91–110.
- Mohamed, A., Zacharakis, E.I., Shen, D., Davatzikos, C., 2006. Deformable registration of brain tumor images via a statistical model of tumor-induced deformation. *Med. Image Anal.* 10, 752–763.
- Mori, S., Crain, B.J., Chacko, V.P., van Zijl, P.C.M., 1999. Three-dimensional tracking of axonal projections in the brain by magnetic resonance imaging. *Ann. Neurol.* 47 (2), 265–269.
- Mori, S., Itoh, R., Zhang, J., Kaufmann, W.E., Zijl, P.C.M.V., Solaiyappan, M., Yarowsky, P., 2001. Diffusion tensor imaging of the developing mouse brain. *Magn. Reson. Med.* 46 (1), 18–23.
- Mukundan, R., Ramakrishnan, K., 1998. *Moment Functions in Image Analysis – Theory and Applications*. World Scientific Publishing, Singapore.
- Pierpaoli, C., Jezzard, P., Basser, J.P., Barnett, A., Chiro, G.D., 1996. Diffusion tensor MR imaging of human brain. *Radiology* 201 (3), 637–648.

- Shen, D., 2007. Image registration by local histogram matching. *Pattern Recogn.* 40, 1161–1171.
- Shen, D., Davatzikos, C., 2002. HAMMER: hierarchical attribute matching mechanism for elastic registration. *IEEE Trans. Med. Imag.* 21 (11), 1421–1439.
- Smith, S., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T., Mackay, C., Watkins, K., Ciccarelli, O., Cader, M., Matthews, P., Behrens, T., 2006. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. *NeuroImage* 31 (4), 1487–1505.
- Westin, C.F., Peled, S., Gudbjartsson, H., Kikinis, R., Jolesz, F.A., April 1997. Geometrical diffusion measures for MRI from tensor basis analysis. *ISMRM '97. International Society for Magnetic Resonance in Medicine (ISMRM)*, p. 1742.
- Westin, C.F., Maier, S.E., Mamata, H., Nabavi, A., Jolesz, F.A., Kikinis, R., 2002. Processing and visualization of diffusion tensor MRI. *Med. Image Anal.* 6 (2), 93–108.
- Xu, D., Mori, S., Shen, D., van Zijl, P.C.M., Davatzikos, C., July 2003. Spatial normalization of diffusion tensor fields. *Magn. Reson. Med.* 50 (1), 175–182.
- Xue, R., van Zijl, P.C.M., Crain, B.J., Solaiyappan, M., Mori, S., 1999. In vivo three-dimensional reconstruction of rat brain axonal projections by diffusion tensor imaging. *Magn. Reson. Med.* 42 (6), 1123–1127.
- Xue, Z., Shen, D., Karacali, B., Stern, J., Rottenberg, D., Davatzikos, C., 2006. Simulating deformations of MR brain images for validation of atlas-based segmentation and registration algorithms. *NeuroImage* 33 (3), 855–866.
- Yang, J., Shen, D., Davatzikos, C., Verma, R., 2008a. Diffusion tensor image registration using tensor geometry and orientation features. *MICCAI'08*, 11 (Pt. 2). Springer Berlin, Heidelberg, pp. 905–913.
- Yang, J., Shen, D., Misra, C., Wu, X., Resnick, S., Davatzikos, C., Verma, R., 2008b. Spatial normalization of diffusion tensor images based on anisotropic segmentation. *Medical Imaging 2008: Image Processing*, 6914. SPIE, p. 69140L.
- Zhang, H., Yushkevich, P.A., Alexander, D.C., Gee, J.C., 2006. Deformable registration of diffusion tensor MR images with explicit orientation optimization. *Med. Image Anal.* 10 (5), 764–785.
- Zhang, H., Avants, B.B., Yushkevich, P.A., Woo, J.H., Wang, S., McCluskey, L.F., Elman, L.B., Melhem, E.R., Gee, J.C., November 2007. High-dimensional spatial normalization of diffusion tensor images improves the detection of white matter differences: an example study using amyotrophic lateral sclerosis. *IEEE Trans. Med. Imag.* 26 (11), 1585–1597.