TIMER: Tensor Image Morphing for Elastic Registration

Pew-Thian YAP[†] Guorong WU[†] Hongtu ZHU[‡] Weili LIN[†] Dinggang SHEN[†]
[†]Department of Radiology, [‡]Department of Biostatistics, and BRIC
University of North Carolina at Chapel Hill, NC 27599
{ptyap,grwu,weili_lin,dgshen}@med.unc.edu, hzhu@bios.unc.edu

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 voxel neighborhoods 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, and fiber tracking indicate that TIMER performs better than the other methods in comparison [12, 14].*

1. Introduction

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

To quantify diffusion tensor abnormalities based on 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. [12] 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 [12], 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 [9] 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 structural image registration to establish correspondence matching as required in the registration process, 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 previous methods. In view of this, new

978-1-4244-3993-5/09/$25.00 ©2009 IEEE

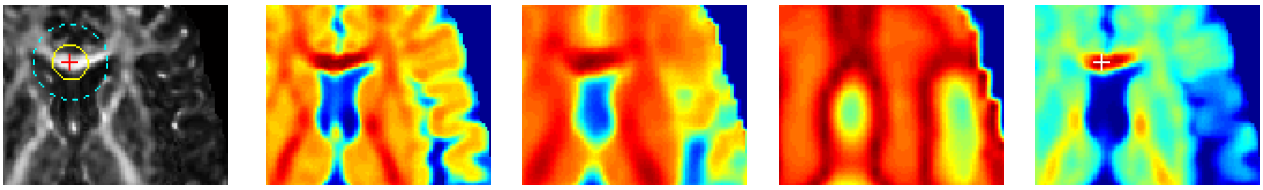

(a) (b) (c) (d) (e)

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

approaches are proposed to work directly on the tensors. However, methods involving single voxel correspondence matching can be susceptible to the effect of noise and also spurious variation in 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 stuctures. 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 its 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 resolutions. 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 irrespective of the neighborhood size 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 the tissue boundaries, local features such as SIFT [7] and RIFT [6] can be incorporated to aid the registration. To this end, we extend Canny edge detector to work directly on the diffusion tensors. The 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 abovementioned, regional information is employed to form an attribute vector for each voxel, which serves as morphological signatures to reflect the anatomical context of each voxel at multiple scales. These attribute vectors are utilized to assist an automated correspondence detection which is formulated in a hierarchical and spatially deformable fashion [9]. 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 tensors are finally reoriented using the method proposed in [10].

2. Tensor Image Morphing for Elastic Registration (TIMER)

2.1. Attribute Vectors

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, an attribute vector is attached to each voxel, reflecting its underlying anatomical structure in a local scale and also its relationship to more distant voxels in a global scale. A rich enough attribute vector can potentially differentiate different parts of the anatomy that would otherwise look similar. In this paper, building upon the ideas in [9], we formulate a deformable registration scheme which aims to minimize the effect of local minima, which is a critical problem since brain anatomy is inherently high dimensional, complex and ambiguous. The energy function being optimized is successively approximated using a sequence of lower dimensional smooth energy functions. This is accomplished in a hierarchical fashion in which initially only a small set of the voxels with highly distinctive attribute vectors are used to drive the volumetric deformation, which is equivalent to an approximate optimization of the energy function. In the later stages, the number of voxels is gradually increased to optimize the energy function more accurately in a higher dimensionality. The driving voxels and hence their attribute vectors are selected hierarchically according to how uniquely they stand out among other attribute vectors in the image, hence reducing ambiguity and local minima. To make the feature vector more discriminative, the above attribute vector is computed at multiple scale so that different anatomical structures are accounted for.

Figure 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 attibute vector in TIMER is designed to comprise of 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 [2], higher order moments [8], and also other inter-tensor measures [1]. For local boundary features, SIFT [7] and RIFT [6] are possible choices. And for tensor geometrical features, prolateness, oblateness, and sphericity measures [12] and orientation features [12] can be employed. However, we find that the current choices of features used in TIMER are sufficiently distinctive to give good registration accuracy. These multiscale features are grouped into an attribute vector for each voxel, i.e.:

$$\begin{aligned}\mathbf{a} = \big[&\big(\mathbf{a}_{\mathrm{F}}^{\text{Regional}}, \mathbf{a}_{\mathrm{F}}^{\text{Edge}}, \mathbf{a}_{\mathrm{F}}^{\text{Geometrical}}\big),\\ &\big(\mathbf{a}_{\mathrm{M}}^{\text{Regional}}, \mathbf{a}_{\mathrm{M}}^{\text{Edge}}, \mathbf{a}_{\mathrm{M}}^{\text{Geometrical}}\big),\\ &\big(\mathbf{a}_{\mathrm{C}}^{\text{Regional}}, \mathbf{a}_{\mathrm{C}}^{\text{Edge}}, \mathbf{a}_{\mathrm{C}}^{\text{Geometrical}}\big)\big] \quad (1)\end{aligned}$$

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:

$$\begin{aligned}&m(\mathbf{a}(\mathbf{x}_1), \mathbf{a}(\mathbf{x}_2)) =\\ &\Pi_{s\in\{\mathrm{F,M,C}\}}\big[\Pi_i(1 - |a_{s,i}^{\text{Regional}}(\mathbf{x}_1) - a_{s,i}^{\text{Regional}}(\mathbf{x}_2)|)\times\\ &\qquad\Pi_j(1 - |a_{s,j}^{\text{Edge}}(\mathbf{x}_1) - a_{s,j}^{\text{Edge}}(\mathbf{x}_2)|)\times\\ &\qquad\Pi_k(1 - |a_{s,k}^{\text{Geometrical}}(\mathbf{x}_1) - a_{s,k}^{\text{Geometrical}}(\mathbf{x}_2)|)\big] \quad (2)\end{aligned}$$

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-the 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 $s \in \mathrm{F, M, C}$, 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 of registration. More detailed descriptions of the features used in TIMER is furnished in the following sections.

2.2. Regional Statistical Features

For each point in the diffusion tensor image, multi-scale regional tensor distribution information can be extracted from its multi-scale neighborhoods. Specifically, for a given tensor $\mathbf{D}(\mathbf{x})$, we can extract the regional tensor distribution information from its neighborhood $\{\mathbf{D}(\mathbf{t})|\mathbf{t} \in \mathcal{N}(\mathbf{x})\}$, where $\mathcal{N}(\mathbf{x})$ is the set of coordinates of voxels in the neighborhood of $\mathbf{x}$. By varying the size of neighborhood $\mathcal{N}(\mathbf{x})$ or scaling the image, we can obtain a rich set of multi-scale regional tensor distribution information, which can be used to extract tensor information at different scales and hence drive the registration hierarchically. 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 commonly used 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 tensors at neighborhoods of multiple scales in the log-space. A matrix exponential operation can then be performed to transform the mean $\mathbf{M}(\mathbf{x})$ back to the original

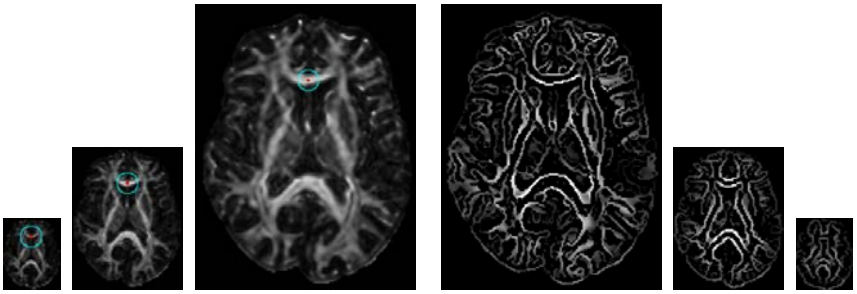

Figure 2. Computation of multiscale features. Instead of changing neighborhood size, the image is progressively downsampled and multiscale features are computed using a fixed neighborhood radius as indicated by the circle. This approach cuts down on the computation time. The first three images on the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

tensor space, obtaining $\widehat{\mathbf{M}}(\mathbf{x})$, 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 [12]. The variance $\mathbf{V}(\mathbf{x})$ can be dealt with in a similar way.

2.2.1 Regional Mean:

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

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

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

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

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

2.2.2 Regional Variance:

Similarly, we can define the regional variance as:

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

and the principal varibilities as:

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

These eigenvalues are further scaled according to the following equation to yield their mean normalized values:

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

We have used $\mathbf{M}^2$ instead of $\mathbf{M}$ to match the dynamic range of $\mathbf{V}$.

2.3. 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.

2.3.1 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 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, which is then 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. For each voxel in the volume, a gradient $\mathbf{H}^{\text{Tensor}}(\mathbf{x})$ can be computed, and from which, after non-maximum suppression, a final edge magnitude $H^{\text{Tensor}}(\mathbf{x})$ can be obtained.

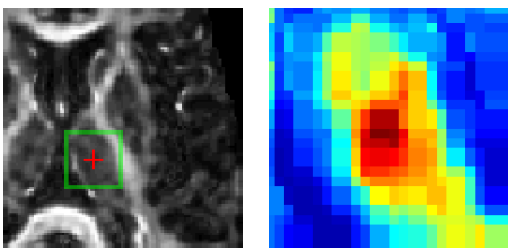

Figure 3. Distinctiveness of attribute vector. The similarity map in the green box is magnified for a closer inspection. Dark red indicates high similarity and dark blue otherwise.

2.3.2 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. 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 detection at point $\mathbf{x}$ as $H^{\text{FA}}(\mathbf{x})$.

3. Experimental Results

A number of experiments were designed to evaluate the accuracy of TIMER. Whenever appropriate, results obtained using DTI registration algorithms proposed by Yang et al. [12] and Zhang et al. [14] will be included for comparison, and they are referred to as YANG and ZHANG[1] respectively in rest of the paper.

3.1. Dataset

The dataset consists of diffusion tensor images of 22 subjects, acquired using a 1.5T scanner. Each of the dataset consisted of 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 240×240 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.

3.2. Real Subjects

One subject was selected from the dataset and taken as the template. 21 subjects were then registered onto this template. By averaging all the registered images, we could visually inspect the accuracy of the registration. The result is shown in Fig. 4. It can be observed that for the FA map based affine registered images, their average image, shown in Fig. 4(b), is fuzzy especially in areas near the cortical surface. In comparison, after registration with TIMER, the average image, Fig. 4(c), shows much improved sharpness. Jones et al. [5] defines the normalized standard deviation of tensors as:

$$\frac{1}{||\bar{\mathbf{D}}||_{\text{F}}}\sqrt{\frac{1}{N-1}\sum_{k=1}^{N}||\mathbf{D}_k - \bar{\mathbf{D}}||_{\text{F}}^2} \tag{8}$$

[1]The version used in our experiments is an extension of [14], which can handle relatively larger deformations as mentioned in [13].

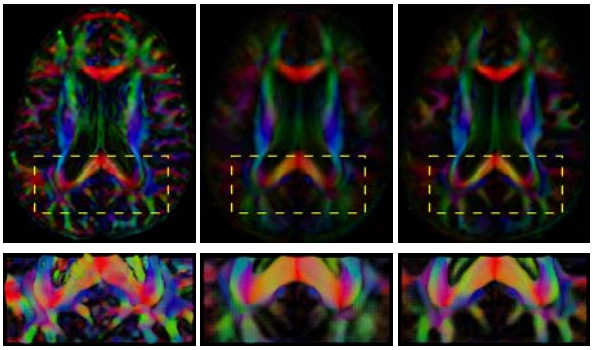

(a) Template (b) Affine (c) TIMER

Figure 4. 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.

where $||\cdot||_{\text{F}}$ 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 kth image. The results for a few slices of the images registered with TIMER and also affine registration are shown in Fig. 5. The images yielded by TIMER are generally darker and is hence indicative of less registration variability. To further quantify the accuracy of the registration, a normalized scalar product of the FA map of each individual registered image with that of the template can be taken. Given the FA maps of the template, $I(\mathbf{x})$, and the subject, $J(\mathbf{x})$, the normalized scalar is defined as:

$$S = \frac{\sum_{\mathbf{x}\in\mathcal{V}} I(\mathbf{x}) \cdot J(\mathbf{x}+\mathbf{u}(\mathbf{x}))}{\left[\sum_{\mathbf{x}\in\mathcal{V}} I^2(\mathbf{x}) \cdot \sum_{\mathbf{x}\in\mathcal{V}} J^2(\mathbf{x}+\mathbf{u}(\mathbf{x}))\right]^{\frac{1}{2}}} \tag{9}$$

where $\mathbf{u}(\mathbf{x})$ is the subject-to-template deformation field and $\mathcal{V}$ denotes the image volume. The average values over all registered images yielded by TIMER, YANG, ZHANG and affine registration are 0.9126, 0.9094 ($p = 4.92 \times 10^{-4}$)[2], 0.8894 ($p = 3.54 \times 10^{-14}$), and 0.8293 ($p = 1.59 \times 10^{-21}$), respectively. For real images, the main problem lies in the fact that one-to-one correspondences between anatomical structures of two images do not always exist. The relatively high value yielded by TIMER is indicative of its ability to cope with situations as such.

3.3. Simulated Deformation Fields

To further evaluate the accuracy of TIMER, we generated 20 simulated deformation fields using the statistical model of deformation (SMD) proposed by Xue et al.

[2]p-values are obtained from paired t-tests performed using the individual values returned by all images. The same goes for the rest of the paper.

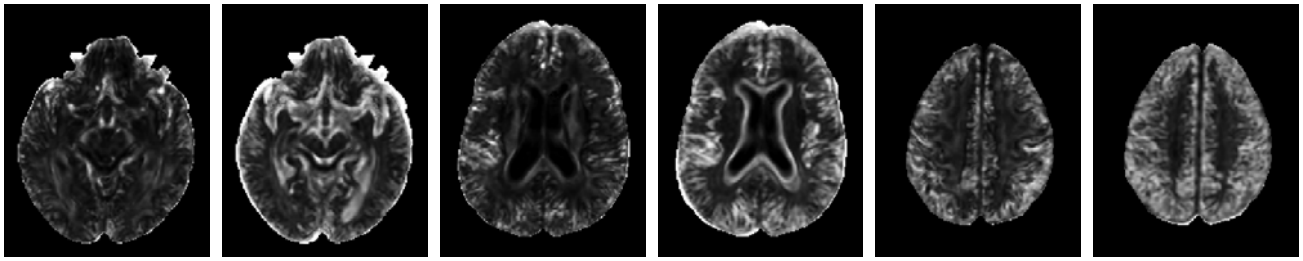

Figure 5. Normalized standard deviation. Starting from the left, the odd and even images are those generated using TIMER and affine registration respectively.

[11]. 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 $\bar{m}_{\mathrm{d}}$ and the standard deviation (SD) σ_{d} of the two deformation fields, i.e. :

$$\begin{aligned} \bar{m}_{\mathrm{d}} &= \frac{1}{|\mathcal{V}|}\sum_{\mathbf{x}\in\mathcal{V}} ||\mathbf{u}_{\mathrm{g}}(\mathbf{x}) - \mathbf{u}_{\mathrm{e}}(\mathbf{x})||, \\ \sigma_{\mathrm{d}} &= \sqrt{\frac{1}{|\mathcal{V}|}\sum_{\mathbf{x}\in\mathcal{V}} \left(||\mathbf{u}_{\mathrm{g}}(\mathbf{x}) - \mathbf{u}_{\mathrm{e}}(\mathbf{x})|| - \bar{m}_{\mathrm{d}}\right)^2} \end{aligned} \tag{10}$$

where $||\cdot||$ is the Euclidean distance, $|\mathcal{V}|$ the number of voxels in the image volume $\mathcal{V}$, and $\mathbf{u}_{\mathrm{g}}(\mathbf{x})$, $\mathbf{u}_{\mathrm{e}}(\mathbf{x})$ the template-to-subject deformation vectors at location $\mathbf{x}$ of the ground truth deformation field and the estimated deformation field, respectively. The average error for the 20 simulated brains is 0.83mm, which indicates that TIMER is accurate to a sub-voxel level. A summary of the results for the whole brain in comparison with YANG and ZHANG is shown in Table 1. In fact, as can be seen from the table, two sets of results of are included, one (SET I) for the previously mentioned 20 simulated images and the other (SET II) for 20 more images simulated in the same manner but with more complex local deformations. The deformation fields for SET I are more regular and generally have smaller displacement magnitudes, whereas those of SET II are more complex and have larger magnitudes, generated using a larger variance in SMD [11]. 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. Also shown in Table 1 are the results for the cortical regions and similar conclusions can be drawn. Overall, TIMER shows improvement over YANG ($p = 8.69 \times 10^{-25}$) and ZHANG ($p = 2.98 \times 10^{-5}$).

3.4. 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 intersubject homology for further statistical inference. Using a tractography method know as *Fiber Assignment by Continuous Tracking* (FACT) [4], fiber bundles passing through some regions of interest (ROIs) were tracked, extracted, and compared for quantifying how well TIMER performs in these specific ROIs. 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 6. The similarity of two fiber bundles was then measured in a similar way to that used in [14]. The similarity measure is defined as:

$$\frac{1}{|\mathcal{F}| + |\mathcal{G}|}\left[\sum_{F_i\in\mathcal{F}} \min_{G_j\in\mathcal{G}} d(F_i, G_j) + \sum_{G_j\in\mathcal{G}} \min_{F_i\in\mathcal{F}} d(F_i, G_j)\right] \tag{11}$$

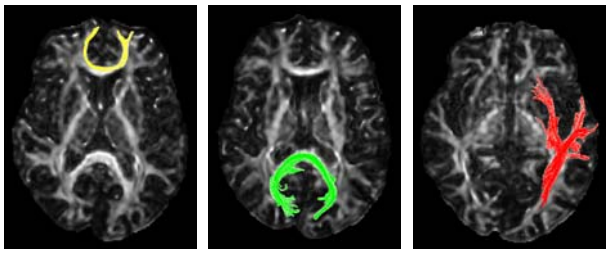

Figure 6. Fiber bundles extracted for comparison. The fiber bundles shown in yellow and green reside in the genu and the splenium of the corpus callosum, respectively. In red is a fiber bundle near the cortical surface.

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

Table 2. Average fiber distances (*mm*). Values in the brackets are the standard deviations.

		TIMER	YANG	ZHANG
SET I	Genu	0.62 (0.025)	0.78 (0.043)	0.67 (0.019)
	Splenium	0.77 (0.004)	0.82 (0.004)	1.00 (0.043)
SET II	Cortical	1.25 (0.125)	1.49 (0.115)	1.91 (0.208)

where $d(F_i, G_j)$ is a pairwise distance between fibers $F_i \in \mathcal{F}$ and $G_i \in \mathcal{G}$, which in our case, is defined as the mean of the closest distance for every point of two fibers [3]. 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. 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 neighorbood around each voxel [10], and is hence less vulnerable to noise. This is attested by the results yielded by TIMER in comparison with YANG ($p = 4.57 \times 10^{-6}$) and ZHANG ($p = 2.68 \times 10^{-8}$), where fiber bundles extracted with TIMER are consistently closer to the ground truths.

4. Conclusion

In conclusion, *Tensor Image Morphing for Elastic Registration* or TIMER is proposed as a relatively accurate diffusion tensor registration algorithm in this paper. 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. 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

This work was supported in part by grants 1R01EB006733, 1R03EB008760 and R01EB008374, and also UNC RRC grant.

References

[1] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 56(2):411–421, 2006.

[2] P. J. Basser and C. Pierpaoli. Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor MRI. *Journal of Magnetic Resonance*, Series B 111(3):209–219, 1996.

[3] G. Gerig, S. Gouttard, and I. Corouge. Analysis of brain white matter via fiber tract modeling. In *Proceedings of the IEEE EMBS*, pages 4421Ű–4424, 2004.

[4] H. Jiang, P. C. van Zijl, J. Kim, G. D. Pearlson, and S. Mori. DtiStudio: resource program for diffusion tensor computation and fiber bundle tracking. *Computer Methods and Programs in Biomedicine*, 81(2):106–116, February 2006.

[5] D. K. Jones, L. D. Griffin, D. C. Alexander, M. Catani, M. A. Horsfield, R. Howard, and S. C. R. Williams. Spatial normalization and averaging of diffusion tensor MRI data sets. *Neuroimage*, 17:592–617, 2002.

[6] S. Lazebnik, C. Schmid, , and J. Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:1265–1278, 2005.

[7] D. Lowe. Distinctive image features form scale-invariant key-points. *International Journal of Computer Vision*, 2:91–110, 2004.

[8] R. Mukundan and K. Ramakrishnan. *Moment functions in image analysis – Theory and Applications*. World Scientific Publishing, Singapore, 1998.

[9] D. Shen and C. Davatzikos. HAMMER: Heirarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging*, 21(11):1421–1439, 2002.

[10] D. Xu, S. Mori, D. Shen, P. C. M. van Zijl, and C. Davatzikos. Spatial normalization of diffusion tensor fields. *Magnetic Resonance in Medicine*, 50(1):175–182, July 2003.

[11] Z. Xue, D. Shen, B. Karacali, J. Stern, D. Rottenberg, and C. Davatzikos. Simulating deformations of MR brain images for validation of atlas-based segmentation and registration algorithms. *NeuroImage*, 33(3):855–866, 2006.

[12] J. Yang, D. Shen, C. Davatzikos, and R. Verma. Diffusion tensor image registration using tensor geometry and orientation features. In *MICCAI'08*, volume 11 (Pt. 2), pages 905–913, 2008.

[13] H. Zhang, B. B. Avants, P. A. Yushkevich, J. H. Woo, S. Wang, L. F. McCluskey, L. B. Elman, E. R. Melhem, and J. C. Gee. High-dimensional spatial normalization of diffusion tensor images improves the detection of white matter differences: An example study using amyotrophic lateral sclerosis. *IEEE Transactions on Medical Imaging*, 26(11):1585–1597, November 2007.

[14] H. Zhang, P. A. Yushkevich, D. C. Alexander, and J. C. Gee. Deformable registration of diffusion tensor mr images with explicit orientation optimization. *Medical Image Analysis*, 10(5):764–785, 2006.