Strategies for Data Reorientation during Non-Rigid Warps of Diffusion Tensor Images

D.C. Alexander, J.C. Gee, and R. Bajcsy

GRASP Lab. and Dept. Radiology,
University of Pennsylvania, 3401, Walnut St., Philadelphia. PA 19104, USA.

Abstract. This paper describes work on the registration of diffusion tensor images of the human brain. An existing registration algorithm, the multi-resolution, elastic matching algorithm, [1-3], has been adapted for this purpose. One problem with the application of such a method to this new data type is that transformations of the image affect the DT values at each voxel, as the orientation can change with respect to the surrounding anatomical structures. Three methods for the estimation of an appropriate reorientation of the data from the local displacement field, which describes the image transformation, are presented and tested. Results indicate that the best matches are obtained from a reorientation strategy that takes into account the effects of local shearing on the data as well as the rigid rotational component of the displacement. The methods presented here may be useful for the computation of region based similarity measures of single valued intensity images, which also vary with local image orientation.

1 Introduction

Diffusion tensor (DT) imaging is a recent innovation in MRI (magnetic resonance imaging), [4]. In DT imaging, the measurement acquired at each voxel in an image volume is a symmetric second order tensor, which describes the local water diffusion properties of the material being imaged. The DT may be thought of as a Gaussian density describing the probability of the final position of a molecule, initially at the centre of a voxel, after some fixed time. DT imaging of the human brain has provoked particular interest because of the added insight it provides into the structure of white matter regions. Neuronal axons are fibres that form the connections between different cells of the brain. In these axons water is free to diffuse along the fibre, but diffusion in perpendicular directions is greatly restricted by the cell wall. DT measurements taken from areas in the brain where large numbers of parallel axons are bundled together, for example the white matter tracts, thus tend to exhibit a high degree of anisotropy and the principal direction (PD) of the DT points along the axes of the bundled fibres. By associating neighbouring voxels according to the PD of their DTs, connection pathways within the brain can be traced and mapped, see for example [5], although the resolution at which these images can currently be obtained is such that only major pathways can be extracted reliably. There are an increasing number of clinical applications of DT imaging, for example, the analysis of stroke and multiple sclerosis, [6].

Three images derived from a slice of a DT image of the human brain are shown in Figure 1. The larger image on the left shows a line at each pixel, which indicates the PD of the DT projected into the xy plane. For pixels at which the PD is close to the z-direction, these lines are less meaningful and are drawn in correspondingly lighter colours. The two images on the right show two common indices that are derived from the DT. The lattice anisotropy, [7], image (top) is hyperintense in white matter regions where the diffusion is strongly weighted in one direction. In the PD image, the crescent shaped regions of high anisotropy at the top and bottom of the image have configuration that suggests they are bundles of fibres running from one side of the brain to the other. The DT trace (bottom) is a measure of the total diffusion at a point and thus tends to be largest in areas where the diffusion is unrestricted in all directions such as regions of CSF (cerebro-spinal fluid).
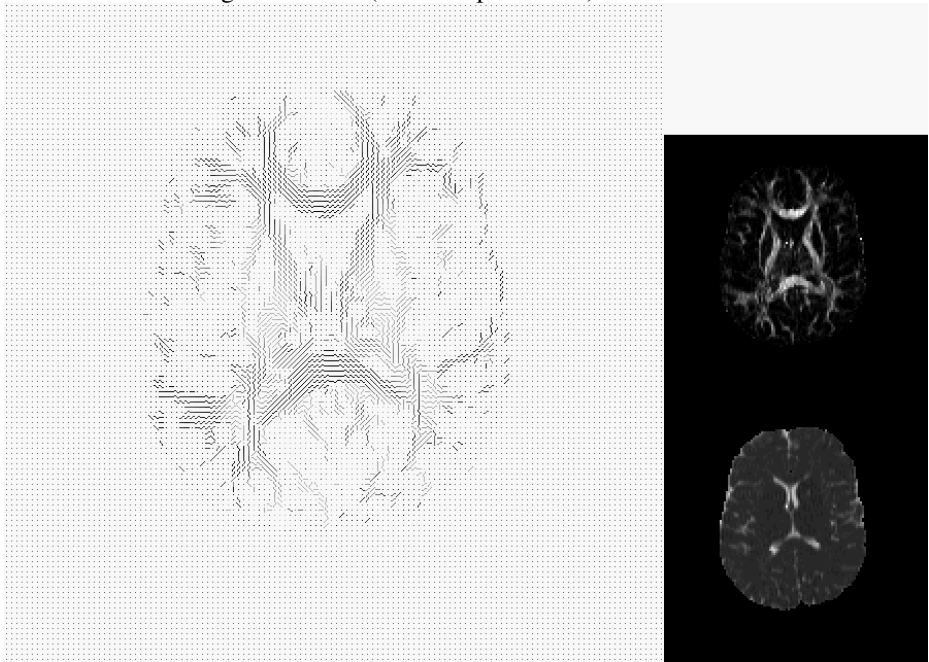

Figure 1 Images derived from a slice of a DT image. Left: principal DT eigenvector projected into the xy-plane. Lighter lines indicate greater z-component in the PD. No line is drawn at points where the anisotropy falls below a certain threshold. Top right: lattice anisotropy image. Bottom right: DT trace image.

Here we consider spatial normalisation of DT images and, in particular, the adaptation of the elastic matching algorithm, [1-3], to work with this new data type. Adaptation of the algorithm is complicated by the fact that each DT has an associated orientation. Image transformations tend to change the orientation of DTs with respect to the surrounding anatomical structure of the image. For a simple image transformation, such as rigid rotation, DT orientation can be preserved by applying a similar transformation to each DT. However, more complex transformations are generally required for accurate spatial normalisation of brain images. The elastic matching algorithm, [1-3], provides spatial normalisation via high dimensional warps

described by a displacement field. The required reorientation of the DTs is dependent on the local properties of the displacement field. In this paper, we discuss strategies for estimation of this reorientation.

Motivation for this work is twofold. Firstly, spatial normalisation of MR images from large groups of patients is of great benefit in clinical studies of anatomical variation over population groups. Spatial normalisation of DT images is required to assist clinical studies of the variation of diffusion properties. It may also be of use in the analysis of the variation of pathways within the brain once the technology to extract these pathways reliably is available. Secondly, the added structure in DT images may allow better anatomical matches to be made between images acquired from different patients, than can be obtained by matching on single valued intensity images. Statistical models of accurate matches based on detailed information, such as combinations of the DT with complementary structural information, could be used to constrain spatial normalisations computed using less expressive data. We also note that the reorientation strategies outlined here may be useful for the computation of region based similarity measures for single valued intensity images undergoing non-rigid transformations. Such measures have already proved effective for image matching, but further improvements might be obtained if the effects of local image reorientation on their values are taken into account.

In the next section, the elastic matching algorithm is described briefly and some issues of its application to DT imagery are discussed. In section 3, we present three strategies for estimating the required reorientation of the DTs during an image transformation. Some experiments to compare these strategies are detailed in section 4 and quantitative results are provided. Finally, conclusions are drawn in section 5.

2   Elastic Matching of DT Images

In this section, we give a brief description of the elastic matching algorithm. Details of the algorithm can be found in, [1-3]. Issues concerning is application to DT imagery, in particular the requirement for reorientation of the data as image transformations are applied, are then discussed.

2.1   Elastic Matching Algorithm

The elastic matching algorithm computes a displacement field describing a warp of one image, which aligns it with another similar image. The warp is computed by iteratively minimising an energy function. The basic energy function contains two terms: a term derived from the similarity of the image data at corresponding points in the fixed target and warped source images, and a term that expresses the amount of deformation caused by the warp.

To provide a starting point for the algorithm that is close to the final solution, a global affine transformation is first applied to the source image. We use Wood's AIR (Automated Image Registration) algorithm, [8,9], to compute the global transformation. Elastic matching is then applied at consecutive levels of a multi-resolution pyramid, from low resolution to high resolution, to find a warp that

optimises the balance between pointwise similarity of the image data and deformation of the source image.

2.2   Application to DT Images

Two issues arise with the application of the elastic matching algorithm to this new type of data. Firstly, how to define the similarity between data at corresponding positions in the two images and secondly, consistency of DT orientations.

For intensity images, for example proton density, T1- or T2-weighted images, it is common to use the squared intensity difference or the cross-correlation of the two images as the measure of local similarity. A detailed discussion of the similarity measures that can be used for comparing DTs is beyond the scope of this paper. We have found that a similarity measure, S_T, based on the magnitude of tensor difference produces good results. Experimental justification of this choice, as opposed to other tensor comparison measures, matching on PD alone or cross correlation of indices derived from the DT, is presented in [10]. For two DTs, $\mathbf{D_1}$ and $\mathbf{D_2}$, S_T is given by

$$S_T = -\frac{1}{\sigma}\sqrt{trace\left[(\mathbf{D_1} - \mathbf{D_2})^2\right]}. \tag{1}$$

The expression in the square root is equivalent to the sum of squared differences of the elements of the DT matrix, and σ is a weighting parameter. S_T is used as the similarity measure in all the experiments presented in this paper.

Here we are concerned with the issue of DT orientation. In the next section we present various strategies for extracting the appropriate reorientation of the DT from the local displacement gradient. In doing so, the image transformation is modelled as locally affine and so we will first consider how we expect basic affine transformations of an image region to affect the DTs at voxels within the region. An affine transformation is comprised of the following components:

- **Rigid translation** of an image region should have no effect on the DTs in the region, since the measured diffusion is not dependent on absolute position.
- **Rigid rotation** of a piece of tissue with respect to a laboratory frame would cause the diffusion characteristics of that tissue to be rotated by the same amount with respect to the laboratory frame of reference. Thus rigid rotation of an image region should be accompanied by a similar transformation of the DTs.
- **Scaling** part (or all) of an image is analogous to changing the size of a particular tissue region. Such a change in size of a region of tissue would generally correspond to an increased number of cells comprising the region. We would not expect the microscopic structure of the tissue to be affected and thus we do not expect the point-wise diffusion properties of the material to change.
- **Shear** of an image region, in general, has the both stretching and rotational effects. Stretching can be viewed in a similar way to scaling and we assume that it does not affect the DT values. The rotational effect of a shear is more complex than the effects of rigid rotation, since its effect on a DT depends on the DTs original orientation with respect to the shear. The orientation of a fibre whose axis is in the same direction as a shearing force would be unaffected by the deformation, but if the axis is in a perpendicular direction the fibre would

undergo some reorientation. We thus expect to observe similar effects in the orientation of DTs in image regions undergoing these transformations. This is illustrated in 2D in Figure 2.
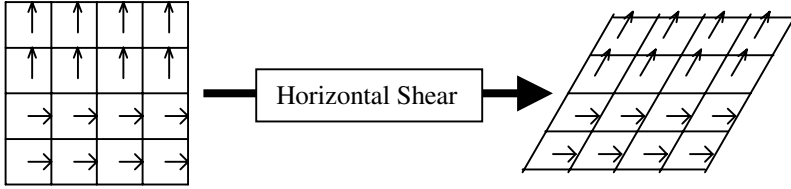

Figure 2 Illustration on the effects of image shear on the orientation of DTs. The arrow in each square of the grid represents the PD of a DT before and after a shear is applied.

We assume throughout that the fundamental diffusion characteristics at any point in the image are unaffected by an image transformation, i.e., the shape of the DT is unaffected so its eigenvalues are preserved. The only change that occurs as a result of the transformation is a reorientation of the axes of the tensor. If $\mathbf{T_1}$ is the DT matrix at a point in an undeformed image, then the matrix, $\mathbf{T_2}$, of the DT in the transformed version of the image is the same matrix after a similarity transform has been applied to rotate the frame of reference of the measurement:

$$\mathbf{T_2} = \mathbf{R}^\mathbf{T}\mathbf{T_1}\mathbf{R} \,, \tag{2}$$

where $\mathbf{R}$ is an appropriate rotation matrix, estimations for which are discussed in the next section.

3   DT Reorientation Strategies

In this section, three strategies are detailed for the extraction of the appropriate reorientation of the DT at each point, i.e., estimation of the rotation matrix, $\mathbf{R}$, in (2). The first two strategies use methods from classical continuum mechanics, which extract the rigid rotation component from the local displacement gradient. The third method accounts for the additional effects of local shearing forces, which is ignored in the first two strategies, by applying the local displacement directly to the principal axes of the DT.

3.1   Small Strain Strategy

Suppose we have two image volumes, I_1 and I_2, and a displacement field, $\underline{u}(\underline{X}) = (u_x(X,Y,Z), u_y(X,Y,Z), u_z(X,Y,Z))$, that describes the warp between the two images, so that if $\underline{x} = \underline{X}+\underline{u}$, then $I_2(\underline{x}) = I_2(\underline{X}+\underline{u}) = I_1(\underline{X})$. The *displacement gradient*,

$$\mathbf{J_u} = \frac{d\underline{u}}{d\underline{X}} = \begin{bmatrix} \dfrac{\partial u_x}{\partial X} & \dfrac{\partial u_x}{\partial Y} & \dfrac{\partial u_x}{\partial Z} \\ \dfrac{\partial u_y}{\partial X} & \dfrac{\partial u_y}{\partial Y} & \dfrac{\partial u_y}{\partial Z} \\ \dfrac{\partial u_z}{\partial X} & \dfrac{\partial u_z}{\partial Y} & \dfrac{\partial u_z}{\partial Z} \end{bmatrix}, \tag{3}$$

at a point, describes the relative displacements of local points and can be decomposed into a rigid rotation component and pure deformation component. If we adopt elastic body models for the two image volumes, as in the elastic matching algorithm, and further assume that the strain on the elastic volume is everywhere small, then a simple additive approximation to this decomposition can be used: $\mathbf{J_u} = \mathbf{E} + \mathbf{\Omega}$. In this decomposition, $\mathbf{E}$, is a symmetric matrix that represents the pure deformation component and $\mathbf{\Omega}$ is the skew-symmetric matrix, given by,

$$\mathbf{\Omega} = \begin{bmatrix} 0 & \dfrac{1}{2}\left(\dfrac{\partial u_x}{\partial Y} - \dfrac{\partial u_y}{\partial X}\right) & \dfrac{1}{2}\left(\dfrac{\partial u_x}{\partial Z} - \dfrac{\partial u_z}{\partial X}\right) \\ -\dfrac{1}{2}\left(\dfrac{\partial u_x}{\partial Y} - \dfrac{\partial u_y}{\partial X}\right) & 0 & \dfrac{1}{2}\left(\dfrac{\partial u_y}{\partial Z} - \dfrac{\partial u_z}{\partial Y}\right) \\ -\dfrac{1}{2}\left(\dfrac{\partial u_x}{\partial Z} - \dfrac{\partial u_z}{\partial X}\right) & -\dfrac{1}{2}\left(\dfrac{\partial u_y}{\partial Z} - \dfrac{\partial u_z}{\partial Y}\right) & 0 \end{bmatrix}. \tag{4}$$

$\mathbf{\Omega}$ represents the relative displacements due to the rigid rotation component at point $\underline{X}=(X,Y,Z)$ in the undeformed image I_1. It can be shown, [11], that $\mathbf{\Omega}$ is the relative displacement matrix equivalent to a rotation about axis $\underline{\omega}$, through angle $|\underline{\omega}|$, where

$$\underline{\omega} = -\Omega_{yz}\underline{i} - \Omega_{zx}\underline{j} - \Omega_{xy}\underline{k} . \tag{5}$$

Ω_{nm} in equation (5) is the entry in the n-th row and m-th column of the matrix $\mathbf{\Omega}$ and $\underline{i}, \underline{j}$, and $\underline{k}$ are unit vectors in the directions of the axes of the original co-ordinate frame, $\underline{X}$.

The transformation matrix, $\mathbf{R}$, corresponding to this rotation can be obtained from the general formula for rotation through angle θ about an axis given by the unit vector (r_x, r_y, r_z), [12]:

$$\mathbf{R} = \begin{bmatrix} r_x^2(1-\cos\theta)+\cos\theta & r_xr_y(1-\cos\theta)-r_z\sin\theta & r_xr_z(1-\cos\theta)+r_y\sin\theta \\ r_xr_y(1-\cos\theta)+r_z\sin\theta & r_y^2(1-\cos\theta)+\cos\theta & r_yr_z(1-\cos\theta)-r_x\sin\theta \\ r_xr_z(1-\cos\theta)-r_y\sin\theta & r_yr_z(1-\cos\theta)+r_x\sin\theta & r_z^2(1-\cos\theta)+\cos\theta \end{bmatrix} \tag{6}$$

.

When the transformation is described locally by a displacement vector field, $\mathbf{J_u}$ is obtained by numerical differentiation of $\underline{u}$. For a global affine transformation, described by a matrix $\mathbf{G}$, $\mathbf{J_u}$, and consequently $\mathbf{R}$, is constant over the entire image volume and $\mathbf{J_u}$ can be obtained directly from $\mathbf{G}$.

3.2   Finite Strain Strategy

Malvern, [11], states that the small strain assumption is justified if the angle of rotation is small compared to one radian. The displacement gradient can always be decomposed into deformation and rigid rotation components, but the additive approximation to this decomposition is inaccurate for larger angles of rotation. In such cases, called "finite strain" cases, [11], a more complex decomposition must be used. Here we just give the expression for the rigid rotation matrix, a derivation of this expression can be found in [11]. For the finite strain case, **R** is given by:

$$\mathbf{R} = \mathbf{F}.\left(\mathbf{F^T}.\mathbf{F}\right)^{-\frac{1}{2}}. \tag{7}$$

Where **F** is the *deformation gradient* given by,

$$\mathbf{F} = \frac{d\underline{x}}{d\underline{X}} = \mathbf{I} + \mathbf{J_U}. \tag{8}$$

Computation of **R** in the finite strain case is more computationally complex, since an eigen-decomposition is required to compute the matrix power in (7), [11].

3.3   Eigenvector Deformation Strategy

The deformation component of the displacement gradient includes transformations such as shearing and non-uniform scaling. As illustrated in Figure 2, these transformations can also affect the orientation, but both the previous strategies discard the deformation component. An alternative approach is to consider the action of the displacement directly on unit vectors in the principal diffusion directions.

There are a number of possible strategies that could be adopted to estimate the appropriate reorientation of the DT in this way. In general, the action of the displacement on a particular elliptical contour of the DT will yield a contour that is no longer elliptical. We wish to preserve the fundamental shape of the DT and, furthermore, the property we are most concerned with preserving is the PD. Thus, in the method we have chosen to use here, we ensure that the PD is mapped perfectly to its image under the local displacement. Orientation of the DT in the orthogonal directions is computed in such a way as to ensure that the other eigenvectors of the DT are as close as possible to their images. The method proceeds as follows:

- Compute unit eigenvectors, $\underline{e}_1$, $\underline{e}_2$, $\underline{e}_3$, of the DT.
- Apply the local deformation gradient, **F**, to the principal eigenvector, $\underline{e}_1$, to find its image in the deformed configuration, $\mathbf{F}(\underline{e}_1)$.
- Compute the rotation matrix, $\mathbf{R_1}$, that maps $\underline{e}_1$ onto a unit vector in the direction of $\mathbf{F}(\underline{e}_1)$. The axis and angle of this rotation are obtained from the vector and scalar products of $\underline{e}_1$ and $\mathbf{F}(\underline{e}_1)$, and $\mathbf{R_1}$ can then be calculated using equation (6).

A secondary rotation is required to map the second principal eigenvector, $\underline{e}_2$, from its position *after* the first rotation to a direction as close as possible to its image under **F**.

- Find the images of $\underline{e}_2$ under transformations **F** and $\mathbf{R_1}$.
- Find the projection, $P(\underline{e}_2)$, of $\mathbf{F}(\underline{e}_2)$ onto a plane perpendicular to $\mathbf{R_1}(\underline{e}_1)$. Note that the 2nd principal eigenvector of the rotated DT, $\mathbf{R_1}(\underline{e}_2)$, already lies in this plane.

- Compute a second rotation, $\mathbf{R_2}$, that rotates $\mathbf{R_1}(\underline{e}_2)$ onto a unit vector in the direction of $P(\underline{e}_2)$. The axis of this rotation is $\mathbf{R_1}(\underline{e}_1)$ and the angle is found from the dot product of $\mathbf{R_1}(\underline{e}_2)$ with $P(\underline{e}_2)$.

Note that in this scheme, the DT reorientation is not constant for global affine transformations, unlike the previous two methods. Although the deformation gradient, $\mathbf{F}$, is constant, its effect on the eigenvectors of the DT depends on the orientation of $\underline{e}_1$ and $\underline{e}_2$, so a separate reorientation must be computed at every voxel.

4   Experiments and Results

In this section, a set of experiments is described, which test the performance of the different DT reorientation strategies proposed in the previous section. Experiments are performed over a single pair of DT images of the human brain. Details of data acquisition are given first, followed by a description of the experiments and, finally, results from these experiments and some discussion.

4.1   Data

Both images are taken from young female subjects. The general image acquisition methodology is identical to that reported in [4]. Images were acquired using a 1.5T GE Signa Horizon EchoSpeed spectrometer. Each DT image consists of 33 contiguous axial slices, with slice thickness 3.5mm, 220mm field of view and 128x128 in-plane resolution. Six gradient directions were sampled and 4 images were acquired for each direction. Four images with no diffusion weighting were also acquired and so a total of 28 T2-weighted acquisitions were made per slice of the DT image volume.

Brain regions were extracted from the background by hand in both images and background voxels are set to a value outside the measurement range.

4.2   Experiments

In all the experiments, the elastic matching algorithm is run at four resolutions from one sixteenth to one half of the full, isotropic, 1mm^3, volume, which is interpolated from the original image. The weighting of the deformation term in the elastic matching energy function is varied from 1.0 to 0.125, through successive negative powers of two for successive levels of the pyramid.

In order to select an optimal value for σ, which controls the weighting of the similarity term of (1), a number of landmark correspondences were defined in each image. For each reorientation strategy, the value of σ was chosen to be that for which the summed distance between landmarks after matching is minimal. For the match obtained with σ set at this optimum value, a second measure is computed, which indicates the similarity of PDs at corresponding points. Care must be taken with measures of this type, as the PD is poorly defined in regions where the diffusion is isotropic, and the difference cannot be relied upon as a measure of match quality. For this reason, the final measure of match quality, E_{PD}, is an average of the dot product of the two PDs weighted by the geometric mean of their lattice anisotropies.

$$E_{PD} = \frac{\sum \sqrt{v(\mathbf{D}_1).v(\mathbf{D}_2)}.\big(pd(\mathbf{D}_1).pd(\mathbf{D}_2)\big)}{\sum \sqrt{v(\mathbf{D}_1).v(\mathbf{D}_2)}}. \tag{9}$$

The sums are taken over the region within both hand extracted brain regions. $\mathbf{D}_1$ and $\mathbf{D}_2$ are the DTs in corresponding positions of the two images, $v(\mathbf{D})$ represents the lattice anisotropy of $\mathbf{D}$ and $pd(\mathbf{D})$ is a unit vector in the PD of $\mathbf{D}$.

4.3   Results

Separate matches are computed using each of the DT reorientation strategies and once more with no DT reorientation. The quality of each match is assessed using E_{PD} and results are given in Table 1. We also give the magnitude of the tensor difference ($-\sigma S_T$ of (1)) averaged over the overlap region to give a measure, E_T, in each case. Note that *higher* values of E_{PD} indicate better match quality, whereas *lower* values of E_T are more desirable.

Table 1 Quality measures for the matches between the pair of brain volumes computed using each of the DT reorientation strategies.

Reorientation strategy	Small Strain	Finite Strain	Eigenvector deformation.	No reorientation.
E_T	627.8	625.3	622.2	623.4
E_{PD}	0.768	0.768	0.772	0.771

The optimal value of the similarity weighting parameter, σ, was found to be around 200 for each strategy. The value of the landmark distance measure is noisy in the vicinity of this minimum point and no significant differences in this measure can be observed between the difference reorientation strategies. The values of E_{PD} and E_T, however are more stable and the differences in these values observed in Table 1 are consistent in the vicinity of $\sigma=200$. With σ set this value, we find that both the small strain and finite strain reorientation strategies produce poor performance and the matches obtained with them are worse than with no reorientation at all. The eigenvector deformation strategy is the only strategy that improves the quality of match over the control case in which no reorientation is applied.

5   Conclusions

We have presented three methods for estimation of the appropriate reorientation of DTs in image volumes undergoing non-rigid transformations, which we model as locally affine. Of the three methods the best results were obtained from the eigenvector deformation strategy, which explicitly computes the effects of the local displacement field on vectors along the principal axes of the diffusion ellipsoid. This is the only method of the three presented that takes into account the reorientation effects of the deformation component of the image transformation.

The numerical differences in the results above are small. This fact may reflect that the structure of these particular brains is such that little reorientation is required

during matching, or it may be that the reorientation is not correct throughout the matched image volume. Experiments on an extended data set are planned. Simple experiments with synthetic data can verify that for a known displacement field, the expected reorientation of the data is extracted correctly by the methods described above. However, we cannot verify that a good voxel-voxel anatomical match is made for the real data.

Although the orientation of the DTs is updated throughout the matching process, there is currently no explicit term in the energy function for their orientation. The inclusion of such a term should allow the elastic matching algorithm to exploit the correspondence of orientational information in the DTs while computing the match and so should further improve the quality of the matches obtained. As mentioned in the introduction, this issue also arises when matching intensity images using region based similarity measures and the improvement of these matches is added motivation for continued investigation in this area.

References

1. R. Bajcsy and S. Kovacic, "Multi-Resolution Elastic Matching", Computer Vision, Graphics and Image Processing, Vol., 46, pp. 1-21, 1989.

2. J.C. Gee and D.R. Haynor, "Numerical Methods for High Dimensional Warps". Chapter in "Brain Warping", ed. A.W. Toga, Academic Press, 1998

3. J.C. Gee and R.K. Bajcsy, "Elastic Matching: Continuum Mechanical and Probabilistic Analysis". Chapter in "Brain Warping", ed. A.W. Toga, Academic Press, 1998

4. C. Pierpaoli, P. Jezzard, P.J. Basser, A. Barnett, G. Di Chiro, "Diffusion Tensor MR Imaging of the Human Brain", Radiology, Vol. 201, No. 3, pp. 637-648, 1996.

5. C. Poupon, J.-F. Mangin, V. Frouin, J. Regis, F. Poupon, M. Pachot-Clouard, D. LeBihan, and I. Bloch, "Regularisation of MR Diffusion Tensor Maps for Tracking White Matter Bundles", Proc. MICCAI'98, MIT, Boston, 1998.

6. M.A. Horsfield, H.B.W. Larsson, D.K. Jones, A. Gass, "Diffusion magnetic resonance imaging in multiple sclerosis", Journal of Neurology, Neurosurgery and Psychiatry, Vol. 64 (Supplement), pp S80-S84, 1998.

7. C. Pierpaoli and P.J. Basser, "Toward a Quantitative Assessment of Diffusion Anisotropy", Magnetic Resonance Medicine, Vol. 36, pp. 893-906, 1996.

8. R.P. Woods, S.T. Grafton, C.J. Holmes, S.R. Cherry and J.C. Mazziotta, "Automated Image Registration: I General Methods and intra-subject intra-modality validation", Journal of Computer Assisted Tomography, Vol. 22 pp. 141-154, 1998.

9. R.P. Woods, S.T. Grafton, J.D.G. Watson, N.L. Sicotte and J.C. Mazziotta, "Automated Image Registration: II Inter-subject validation of linear and non-linear models", Journal of Computer Assisted Tomography, Vol. 22 pp. 155-165, 1998.

10. D.C. Alexander, R. Bajcsy and J.C. Gee, "Elastic Matching of Diffusion Tensor MRIs", accepted for publication in Computer Vision and Image Understanding, 1999.

11. L.E. Malvern, "Introduction to the Mechanics of a Continuous Medium", Prentice-Hall, Inc. Englewood Cliffs, N.J., 1969.

12. K.S. Fu, R.C. Gonzalez and C.S.G. Lee, "Robotics: Control, Sensing, Vision and Intelligence", New York: McGraw-Hill, 1987.