

A KERNEL-BASED GRAPHICAL MODEL FOR DIFFUSION TENSOR REGISTRATION

A. Sotiras^{1,2}, R. Neji^{1,2,3}, J.-F. Deux⁴, N. Komodakis⁵, G. Fleury³ and N. Paragios^{1,2}

¹Laboratoire MAS, Ecole Centrale Paris, Châtenay-Malabry, France

²Equipe GALEN, INRIA Saclay - Île-de-France, Orsay, France

³Département Signaux et Systèmes Électroniques, Supélec, Gif-sur-Yvette, France

⁴Centre Hospitalier Universitaire Henri Mondor, Créteil, France

⁵Computer Science Department, University of Crete, Crete, Greece

ABSTRACT

In this paper, we propose a novel method for the spatial normalization of diffusion tensor images. The proposed method takes advantage of both the diffusion information and the spatial location of tensor in order to define an appropriate metric in a probabilistic framework. A registration energy is defined in a Reproducing Kernel Hilbert Space (RKHS), encoding the image dissimilarity and the regularity of the deformation field in both the translation and the rotation space. The problem is reformulated as a graphical model where the latent variables are the rotation and the translation that should be applied to every tensor and the observed variables are the tensors themselves. Efficient linear programming is used to minimize the resulting energy. Quantitative and qualitative results on a manually annotated dataset of diffusion tensor images demonstrate the potential of the proposed method.

Index Terms— Diffusion tensor imaging, spatial normalization, kernels, markov random fields, discrete optimization

1. INTRODUCTION

Diffusion Tensor Imaging (DTI) is a fairly new modality that is able to provide clinicians with useful information about the structure and the geometry of the observed tissues. It has been mainly used to study the connectivity between the different structures of the human brain. DTI has been lately used to study the human skeletal muscles as diffusion can provide information about the structure and the organization of the muscle fibers.

For clinicians to be able to draw valid conclusions from their clinical studies, medical image registration is an important tool. In general, it consists of estimating a spatial transformation such that two images are aligned. In particular, Diffusion Tensor (DT) registration not only tries to recover the spatial correspondences but also reorient the tensors accordingly to account for the rotational component of the spatial deformation [1]. The directional information of the diffusion tensors as well as the high-dimensionality of the data further complicates the registration process. The existing diffusion tensor registration algorithms can be subdivided in two classes.

The first class of methods extracts features from the diffusion tensor data and uses them in vector-data registration algorithms. In [2], features describing the distribution of tensor geometry over an isotropic neighborhood as well as ones conveying information about the orientation are combined in a vector registration framework. In a similar way, the Geodesic Loxodrome distance and a modified Multidimensional Scaling are used in combination with tensor shape and

orientation information in [3]. More recently, in [4], the proposed algorithm includes regional tensor statistics and edge maps to capture relevant information in a multiscale fashion for hierarchical registration.

The second class of approaches takes advantage of standard scalar registration techniques by defining appropriate metrics for tensor matching. In particular, in [5], a fluid registration framework followed by preservation of principal directions (PPD) tensor reorientation scheme is proposed where the symmetrized Kullback-Leibler divergence between probabilities is used to compare the tensors. Methods that allow for an explicit reorientation scheme and consider the exact gradient of the objective function are proposed in [6, 7]. The first builds upon the Demons algorithm for the scalar images by introducing an exact Finite-Strain (FS) differential. The second is based on a piecewise affine deformation model where an optimization of the rotational component of the transformation is possible.

In this paper, both spatial and diffusion information are used in order to define an appropriate metric between tensors. Both types of information are combined in a probabilistic framework where a probabilistic kernel maps diffusion tensor images in an implicit way to a Reproducing Kernel Hilbert Space (RKHS). In this space, local smoothness constraints impose the regularity for both the spatial transformation and the reorientation of the tensors. What is more, we extend the framework proposed in [8] to diffusion tensor images. The fast Primal-Dual (fast-PD) optimization method [9] is used to minimize the proposed energy in a discrete setting where the six-dimensional deformation space is quantized. The quaternion representation of rotations is considered allowing for proper interpolation and sampling of the rotations space.

2. MAPPING DIFFUSION PROBABILITIES TO A HILBERT SPACE

Diffusion tensors refer to the covariance of a Gaussian probability over the displacements \mathbf{r} of the water protons given a diffusion (mixing) time t which is provided by the following expression:

$$p(\mathbf{r}|t, \mathbf{D}) = \frac{1}{\sqrt{\det(\mathbf{D})(4\pi t)^3}} \exp\left(-\frac{\mathbf{r}^t \mathbf{D}^{-1} \mathbf{r}}{4t}\right).$$

The spatial arrangement of the tensors conveys valuable information. In order to take the spatial context into account, we proposed in [10] to incorporate the spatial information in a probabilistic setting. Given a diffusion tensor \mathbf{D} localized at voxel \mathbf{x} , the probability of the position \mathbf{y} of the water molecule previously localized at \mathbf{x} can

The two first authors contributed equally to this work.

be obtained in a straightforward way:

$$p(\mathbf{y}|\mathbf{x}, t, \mathbf{D}) = \frac{1}{\sqrt{\det(\mathbf{D})(4\pi t)^3}} \exp\left(-\frac{(\mathbf{y} - \mathbf{x})^t \mathbf{D}^{-1}(\mathbf{y} - \mathbf{x})}{4t}\right)$$

Let us now consider the normalized L^2 inner-product between two Gaussian diffusion probability densities p_1 and p_2 with parameters (x_1, \mathbf{D}_1) and (x_2, \mathbf{D}_2) respectively. It is given by the following expression:

$$k_t(p_1, p_2) = \frac{\int p_1(\mathbf{y}|\mathbf{x}_1, t, \mathbf{D}_1) p_2(\mathbf{y}|\mathbf{x}_2, t, \mathbf{D}_2) d\mathbf{y}}{\sqrt{\int (p_1(\mathbf{y}|\mathbf{x}_1, t, \mathbf{D}_1))^2 d\mathbf{y}} \sqrt{\int (p_2(\mathbf{y}|\mathbf{x}_2, t, \mathbf{D}_2))^2 d\mathbf{y}}}$$

It can be shown, based on [11], that the previous equation has an interesting closed form that is a product of two terms:

$$k_t(p_1, p_2) = 2\sqrt{2} \underbrace{\frac{\det(\mathbf{D}_1)^{\frac{1}{4}} \det(\mathbf{D}_2)^{\frac{1}{4}}}{\sqrt{\det(\mathbf{D}_1 + \mathbf{D}_2)}}}_{\text{tensor similarity term}} \times \underbrace{\exp\left(-\frac{1}{4t}(\mathbf{x}_1 - \mathbf{x}_2)^t (\mathbf{D}_1 + \mathbf{D}_2)^{-1} (\mathbf{x}_1 - \mathbf{x}_2)\right)}_{\text{spatial connectivity term}} \quad (1)$$

Therefore k_t takes into account the tensor affinity as well the spatial position. This is crucial since combination of spatial and diffusion information allows for a better modeling of the interactions between tensors. It is easy to verify that the kernel k_t verifies the Mercer property over the space of multivariate normal distributions. The Mercer property accounts for the existence of a mapping ϕ associated with the kernel k_t that provides an embedding from the space of Gaussian probabilities in the RKHS \mathcal{H} such that $k_t(p_1, p_2) = \langle \phi(p_1), \phi(p_2) \rangle_{\mathcal{H}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product of \mathcal{H} . This allows to have the following Hilbertian metric δ_t :

$$\delta_t(p_1, p_2) = \sqrt{2 - 2k_t(p_1, p_2)}$$

In order to ease the notation, in the remainder of the paper we will identify a Gaussian probability distribution with its parameters and denote:

$$\delta_t(p_1(\mathbf{y}|\mathbf{x}_1, t, \mathbf{D}_1), p_2(\mathbf{y}|\mathbf{x}_2, t, \mathbf{D}_2)) = \delta_t((\mathbf{x}_1, \mathbf{D}_1), (\mathbf{x}_2, \mathbf{D}_2)).$$

3. DEFORMABLE REGISTRATION

Let us consider a source DT image $U : \Omega \mapsto S_+(3)$ and a target image V , where Ω is the source image domain and $S_+(3)$ is the space of symmetric positive definite matrices. We aim at computing a deformation field $\mathcal{T} : \Omega \mapsto \mathbb{R}^3 \times \text{SO}(3)$ where $\text{SO}(3)$ is the special orthogonal group. At each point $\mathbf{x} \in \Omega$, $\mathcal{T}(\mathbf{x}) = (\mathbf{t}(\mathbf{x}), \mathbf{R}(\mathbf{x}))$ is a pair composed of a translation vector $\mathbf{t}(\mathbf{x})$ and a rotation matrix $\mathbf{R}(\mathbf{x})$ that deforms U in an image W such that $W(\mathbf{x} + \mathbf{t}(\mathbf{x})) = \mathbf{R}(\mathbf{x})U(\mathbf{x})\mathbf{R}(\mathbf{x})^t$.

We consider a grid-based deformation model that can provide for one-to-one and invertible transformations. The basic idea of the deformation model is that by superimposing a grid $G : [1, K] \times [1, L] \times [1, M]$ (where K, L and M are smaller than the dimensions of the domain) onto the moving image, it is possible to deform the embedded image by manipulating the control points belonging to

the grid. Consequently, the goal is to recover the deformation vector $\mathcal{T}_{\mathbf{p}} = (\mathbf{t}_{\mathbf{p}}, \mathbf{R}_{\mathbf{p}})$ that should be applied to the node \mathbf{p} of the grid, in order for the images to be aligned. In such a framework, the deformation $\mathcal{T}(\mathbf{x}) = (\mathbf{t}(\mathbf{x}), \mathbf{R}(\mathbf{x}))$ that should be applied to an image position \mathbf{x} can be obtained through interpolation of the deformations obtained at the control points:

$$\mathbf{t}(\mathbf{x}) = \sum_{\mathbf{p} \in G} \eta_s(|\mathbf{x} - \mathbf{p}|) \mathbf{t}_{\mathbf{p}}, \quad \mathbf{R}(\mathbf{x}) = \sum_{\mathbf{p} \in G} \eta_r(|\mathbf{x} - \mathbf{p}|) \mathbf{R}_{\mathbf{p}}.$$

η_s and η_r are functions that weight the influence of each control point of the grid to each point of the domain in relation to their spatial distance from it. The group of rotations matrices is not a linear space, this motivates the use of the quaternion representation in order to interpolate the rotations.

Given the above-defined deformation model, the DT images will be deformed in such a way that an appropriately defined dissimilarity criterion with respect to the distance δ_t is minimized:

$$E_d = \frac{1}{|G|} \sum_{\mathbf{p} \in G} \int_{\Omega} \eta_p^{-1}(|\mathbf{x} - \mathbf{p}|) \delta_t((\mathbf{x}, W(\mathbf{x})), (\mathbf{x}, V(\mathbf{x}))) d\mathbf{x}. \quad (2)$$

E_d is simply a data term that will drive the deformation towards a minimal mismatch between the deformed image W and the target image V , written by using the control points of the superimposed grid G . Note that in E_d , only the tensor similarity term in (Eq.1) is relevant, since we will compare tensors that share the same location. The back-projection function η^{-1} computes the influence of the position \mathbf{x} to the control point \mathbf{p} . If the nearest neighbor weighting scheme is considered, then each position \mathbf{x} contributes to only one control point \mathbf{p} with a weight equal to one. In the general case, it takes the following form $\eta_p^{-1}(|\mathbf{x} - \mathbf{p}|) = \frac{\eta_p(|\mathbf{x} - \mathbf{p}|)}{\int_{\Omega} \eta_p(|\mathbf{y} - \mathbf{p}|) d\mathbf{y}}$. It should be noted that different weighting schemes can be used for the interpolation of the displacement field (η_s), the interpolation of the rotations (η_r) as well as for the back-projection to the nodes of the grid (η_p).

In order to account for the ill-posedness of the problem as well as the local structural information of the source image U a regularization term is needed. We suppose that the deformation field is approximately, up to a suitable change in the diffusion time t to account for local rescaling, *locally isometric* in the RKHS \mathcal{H} , i.e. that it preserves the distance δ_t between spatially neighboring Gaussian probabilities when deforming U and accounts for a possible change of scale. Thus, the smoothness term is defined as:

$$E_s = \int_G \int_{\mathbf{z} \in \mathcal{N}_{\mathbf{x}}} |\delta_t((\mathbf{x}, U(\mathbf{x})), (\mathbf{z}, U(\mathbf{z}))) - \delta_{t_{\mathbf{z}\mathbf{x}}}((\bar{\mathbf{x}}, W(\bar{\mathbf{x}})), (\bar{\mathbf{z}}, W(\bar{\mathbf{z}})))| d\mathbf{z} d\mathbf{x} \quad (3)$$

where $\bar{\mathbf{x}} = \mathbf{x} + \mathbf{t}(\mathbf{x})$, $\bar{\mathbf{z}} = \mathbf{z} + \mathbf{t}(\mathbf{z})$, $t_{\mathbf{z}\mathbf{x}} = t \frac{\|\bar{\mathbf{x}} - \bar{\mathbf{z}}\|^2}{\|\mathbf{x} - \mathbf{z}\|^2}$ and $\mathcal{N}_{\mathbf{x}}$ is a local neighborhood of \mathbf{x} on the grid G . We expect the minimization of E_s to favor tensor reorientation so that the local source image structure can be preserved in the deformed image. In other words, it is through the regularization term that the rotations and the translations are coupled in such a way that the local structure in the deformed image, expressed by the distance between spatially neighboring Gaussian probability distributions, remain consistent with the local structure in the source image. An important underlying property is that the coupling provided by the smoothness term constrains the rotations with respect to the displacements, so that tensors do not rotate independently from the translation they undergo with respect to their neighbors.

Following recent ideas in scalar image registration [8] and recent advances in discrete optimization [9], we opt for the use of a discrete optimization technique called Fast-PD [9]. The reason behind our choice is twofold: Fast-PD can provide an optimal solution (up to a user-defined bound) in an efficient way and allows for a gradient-free optimization thus permitting the use of different deformation models.

To be able to apply the Fast-PD optimization, it is obligatory to provide a quantized version of the deformation space. Let $\Theta = (\mathbf{d}^1, \dots, \mathbf{d}^n)$ be a quantized version of the deformation space $\mathbb{R}^3 \times \text{SO}(3)$, then to each quantized deformation \mathbf{d}^i , a label l^i can be assigned to it, thus defining a discrete set of labels $L = \{l^1, \dots, l^n\}$. Then, assigning a label $l_{\mathbf{p}}$ to the node \mathbf{p} , where $l_{\mathbf{p}} \in L$, corresponds to applying the deformation $\mathbf{d}^{l_{\mathbf{p}}}$ to the node, that is translating it by $\mathbf{t}^{l_{\mathbf{p}}}$ and rotating the corresponding tensor by $\mathbf{R}^{l_{\mathbf{p}}}$.

The quantization of spatial displacements is intuitive, the case of rotations is however less straightforward. In order to quantize the group of rotation matrices, we use their quaternion representation. The problem is equivalent to sampling points over the unit sphere \mathbb{S}^3 of \mathbb{R}^4 . We use layered Sukharev Grid sequences [12] that offer a multi-resolution, deterministic and uniform sampling of \mathbb{S}^3 by back-projecting points sampled over a hypercube inscribed in \mathbb{S}^3 outward onto the spherical surface. The set Θ is therefore formed by the pairs of sampled translations and rotations.

Following [8], we cast the registration problem as a discrete multi-labeling problem. In such a context, the goal is to recover the optimal individual label $l_{\mathbf{p}}$ that should be assigned to each node \mathbf{p} of the grid. This can be done using the theory of MRFs, the general form of which is the following:

$$E_{\text{MRF}} = \sum_{\mathbf{p} \in G} V_{\mathbf{p}}(l_{\mathbf{p}}) + \sum_{\mathbf{p} \in G} \sum_{\mathbf{q} \in \mathcal{N}(\mathbf{p})} V_{\mathbf{p}\mathbf{q}}(l_{\mathbf{p}}, l_{\mathbf{q}}) \quad (4)$$

where $V_{\mathbf{p}}(\cdot)$ are the unary potentials that encode the data term and $V_{\mathbf{p}\mathbf{q}}(\cdot, \cdot)$ are the pairwise potentials that encode smoothness constraints. $\mathcal{N}(\mathbf{p})$ represents the neighborhood system of the node \mathbf{p} . The unary potentials will be defined according to the data term in (Eq.2):

$$V_{\mathbf{p}}(l_{\mathbf{p}}) \approx \int_{\Omega} \eta_{\mathbf{p}}^{-1}(|\mathbf{x} - \mathbf{p}|) \delta_{\epsilon}((\mathbf{x}, W(\mathbf{x})), (\mathbf{x}, V(\mathbf{x}))) d\mathbf{x}, \quad (5)$$

where W is deformed by applying the label $l_{\mathbf{p}}$ to the node \mathbf{p} , that is translating it and rotating the corresponding tensor accordingly. Similarly the pairwise potentials are derived following (Eq.3):

$$V_{\mathbf{p}\mathbf{q}}(l_{\mathbf{p}}, l_{\mathbf{q}}) = |\delta_{\epsilon}((\mathbf{p}, U(\mathbf{p})), (\mathbf{q}, U(\mathbf{q}))) - \delta_{\epsilon}(\mathbf{p}, W(\mathbf{p})), (\mathbf{q}, W(\mathbf{q}))| \quad (6)$$

where $(\mathbf{p}, W(\mathbf{p}))$ (resp. $(\mathbf{q}, W(\mathbf{q}))$) are obtained by applying the deformation parameters of the label $l_{\mathbf{p}}$ (resp. $l_{\mathbf{q}}$). More schematically, it is as if any pair of neighboring grid tensors (\mathbf{x}, \mathbf{z}) were linked by a spring of a known rest length $l_0(\mathbf{x}, \mathbf{z}) = \delta_{\epsilon}((\mathbf{x}, U(\mathbf{x})), (\mathbf{z}, U(\mathbf{z})))$. Deforming the tensors will affect the length of the spring which is now given by the equation $l(\mathbf{x}, \mathbf{z}) = \delta_{\epsilon}(\mathbf{x}, W(\mathbf{x})), (\mathbf{z}, W(\mathbf{z}))$. Thus, the minimization of the smoothness energy will amount to keeping the length of the spring l close enough to its rest length l_0 . The number of labels acts as a bottleneck from computational perspective both in terms of time and needed memory. As a result, we use a compositional multiscale approach, that will gradually refine the solution by applying finer grids (and consequently shorter diffusion times t) and label sets.

4. EXPERIMENTAL VALIDATION

For validation purposes, we considered DT images of the calf muscle of 10 healthy subjects. The images were acquired with a 1.5 T MRI scanner using the following parameters : $TR = 3600 \text{ ms}$, $TE = 70 \text{ ms}$ and b value of 700 s.mm^{-2} with 12 gradient directions and 13 repetitions. The size of the obtained volumes is $64 \times 64 \times 20$ voxels with a voxel resolution of $3.125 \text{ mm} \times 3.125 \text{ mm} \times 7 \text{ mm}$. High-resolution T1-weighted images were simultaneously acquired and segmented in 7 muscle groups by an expert.

Based on the given segmentations, the dice overlap, the sensitivity and the specificity of the deformed source segmentation with respect to the target segmentation, are computed. Moreover, four angular similarity criteria are also evaluated on the target mask: the mean difference in the azimuthal angle θ and the polar angle ϕ in spherical coordinates of the principal directions of diffusion, their average angular separation (AAS) as well as the average overlap of eigenvalue-eigenvector pairs (AOE). We also compute the mean difference in fractional anisotropy (FA).

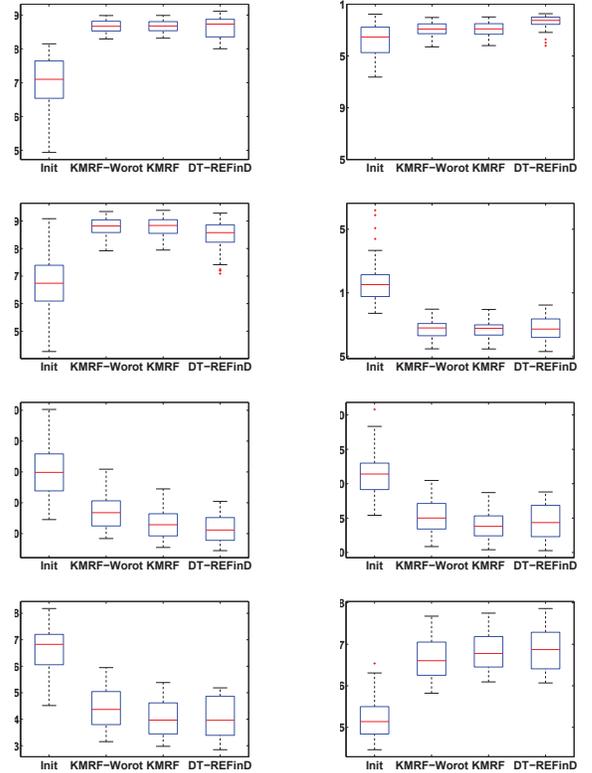


Fig. 1. Boxplots of the evaluation criteria over the 50 registrations before registration (Init), with our method with a single identity rotation label (KMRF-Worot) and several rotation labels (KMRF), as well as the method in [6] (DT-REFinD).

Among the possible 90 registrations, we chose randomly a subset of 50 pairs of DT images. In all our experiments, we used a three-level multiresolution scheme. The grids used at the three levels were of size $6 \times 6 \times 5$, $12 \times 12 \times 10$ and $18 \times 18 \times 15$. The following diffusion times were used: $t = \{2 \cdot 10^5, 5 \cdot 10^4, 2 \cdot 10^4\}$. A number of $73 = 18 \times 4 + 1$ translation labels were used per res-

olution level, sampled along the horizontal and vertical directions as well as the diagonals. For rotation sampling, we generated 10^3 , 10^4 and 10^5 quaternions using Sukharev layered grids. Of these, we selected 100, 50 and 25 for the three levels respectively. These samples were chosen as the closest with respect to the geodesic distance $\arccos(\cdot, \cdot)$ on \mathbb{S}^3 to the identity matrix (or equivalently with the smallest angle). Towards imposing the diffeomorphic property on the deformation field, we use a cubic B-spline interpolation of the displacement field, with the maximum displacement being restricted to 0.4 times the grid spacing. We used a simple trilinear scheme for tensor interpolation and a nearest-neighbor backprojection (η_p).

For the sake of comparison, we provide the values of the computed evaluation criteria before and after registration. We also compare our method to a reference algorithm proposed in [6] (the software is publicly available at <http://www-sop.inria.fr/asclepios/software/MedINRIA/>) and to the result of our method without a rotational component, i.e. with a single rotation label equal to the identity matrix. For the reference algorithm, we considered a three-level multiresolution pyramid with a smoothing kernel of size 1 and a maximum displacement of 4.

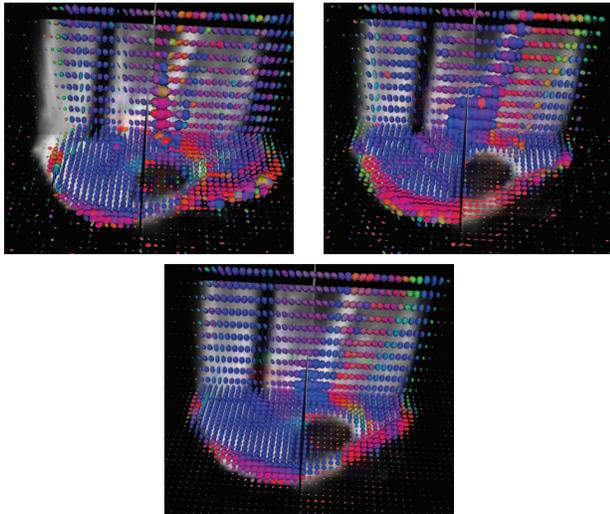


Fig. 2. From left to right and top to bottom: moving, fixed and deformed tensors. All are overlaid on the B0-image of the target subject. RGB colors encode principal directions of diffusion.

We report in (Fig.1) the boxplots of the evaluation criteria over the 50 registrations for our method and for the approaches described above. Our approach improves significantly all the evaluation criteria with respect to the initial state (no registration) and achieves close results to [6]. We run a paired statistical Student t-test with a significance level of 0.05 for comparison and we found that the two approaches performed equivalently for the dice and FA, that our method achieved better results for θ , AAS and sensitivity while [6] performed better in ϕ , AOE and specificity. The inclusion of rotation labels improved (significantly according to the t-test) all the angular criteria with respect to the no-rotation experiments.

For qualitative evaluation, we report in (Fig.2) a view of moving tensors, target tensors and deformed tensors overlaid over the baseline target image. We can see that the spatial mismatch is minimized while the tensor field obtained is smooth and the directions of the deformed tensors are similar to the fixed ones.

5. DISCUSSION

We introduced a novel method for diffusion tensor registration. The main contribution is two-fold. First, we proposed to use a diffusion probability kernel that models both spatial and data dependencies in the tensor field in order to drive the registration process. The proposed formulation allows for the matching of the deformed and the target images while reorienting the tensors and taking into account the local structural information of the source image. Moreover, we showed that the MRF-based formulation for scalar images proposed in [8] can be extended to the case of tensor images. A possible improvement could be to consider automatic and location-dependent adaptive quantization of the search space.

ACKNOWLEDGMENTS

The authors would like to thank M. Maatouk, A. Rahmouni and G. Bassez from Centre Hospitalier Universitaire Henri Mondor for acquiring and segmenting manually the used data.

6. REFERENCES

- [1] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee, "Spatial transformations of diffusion tensor magnetic resonance images," *TMI*, 2001.
- [2] J. Yang, D. Shen, C. Davatzikos, and R. Verma, "Diffusion tensor image registration using tensor geometry and orientation features," in *MICCAI*, 2008.
- [3] M. O. Irfanoglu, R. Machiraju, S. Sammet, C. Pierpaoli, and M. V. Knopp, "Automatic deformable diffusion tensor registration for fiber population analysis," in *MICCAI*, 2008.
- [4] P-T. Yap, G. Wu, H. Zhu, W. Lin, and D. Shen, "Timer: Tensor image morphing for elastic registration," *Neuroimage*, 2009.
- [5] M-C. Chiang, A. D. Leow, A. D. Klunder, R. A. Dutton, M. Barysheva, S. E. Rose, K. McMahon, G. I. de Zubicaray, A. W. Toga, and P. M. Thompson, "Fluid registration of diffusion tensor images using information theory," *TMI*, 2008.
- [6] B. Yeo, T. Vercauteren, P. Fillard, X. Pennec, P. Golland, N. Ayache, and O. Clatz, "DTI registration with exact finite-strain differential," in *ISBI*, 2008.
- [7] H. Zhang, P. A. Yushkevich, D. C. Alexander, and J. C. Gee, "Deformable registration of diffusion tensor MR images with explicit orientation optimization," *MedIA*, 2006.
- [8] B. Glocker, N. Komodakis, G. Tziritas, N. Navab, and N. Paragios, "Dense image registration through MRFs and efficient linear programming," *MedIA*, 2008.
- [9] N. Komodakis, G. Tziritas, and N. Paragios, "Performance vs computational efficiency for optimizing single and dynamic MRFs: Setting the state of the art with primal-dual strategies," *CVIU*, 2008.
- [10] R. Neji, A. Besbes, N. Komodakis, J-F. Deux, M. M., A. Rahmouni, G. Bassez, G. Fleury, and N. Paragios, "Clustering of the human skeletal muscle fibers using linear programming and angular hilbertian metrics," in *IPMI*, 2009.
- [11] T. Jebara, R. Kondor, and A. Howard, "Probability product kernels," *JMLR*, 2004.
- [12] A. Yershova and S. M. LaValle, "Deterministic sampling methods for spheres and $SO(3)$," in *ICRA*, 2004.