

Universität des Saarlandes



Fachrichtung 6.1 – Mathematik

Preprint Nr. 221

**A Variational Approach to the Registration
of Tensor-Valued Images**

Sebastiano Barbieri, Martin Welk and Joachim Weickert

Saarbrücken 2008

A Variational Approach to the Registration of Tensor-Valued Images

Sebastiano Barbieri

School of Engineering Science
Simon Fraser University
8888 University Drive
V5A 1S6, Burnaby, BC, Canada
sba61@sfu.ca

Martin Welk

Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University
Campus E1.1, 66041 Saarbrücken, Germany
welk@mia.uni-saarland.de

Joachim Weickert

Mathematical Image Analysis Group
Faculty of Mathematics and Computer Science, Saarland University
Campus E1.1, 66041 Saarbrücken, Germany
weickert@mia.uni-saarland.de

Edited by
FR 6.1 – Mathematik
Universität des Saarlandes
Postfach 15 11 50
66041 Saarbrücken
Germany

Fax: + 49 681 302 4443
e-Mail: preprint@math.uni-sb.de
WWW: <http://www.math.uni-sb.de/>

Abstract

We present a variational framework for the registration of tensor-valued images. It is based on an energy functional with four terms: a data term based on a diffusion tensor constancy constraint, a compatibility term encoding the physical model linking domain deformations and tensor reorientation, and smoothness terms for deformation and tensor reorientation. Although the tensor deformation model employed here is designed with regard to diffusion tensor MRI data, the separation of data and compatibility term allows to adapt the model easily to different tensor deformation models. We minimise the energy functional with respect to both transformation fields by a multiscale gradient descent. Experiments demonstrate the viability and potential of this approach in the registration of tensor-valued images.

1 Introduction

Image registration denotes the transformation of multiple images of the same object into a common coordinate system [10]. It is of particular importance in medical imaging where data from different imaging modalities or images taken at different times need to be fused to allow diagnosis, surgery or treatment planning.

Registration involves two tasks: first, to find a map that yields for each location in the template image the corresponding location in the reference image; second, to deform the template image accordingly such that the coordinates of corresponding locations in the reference and deformed template image are identical. From the computer vision viewpoint, the first task is a correspondence problem, similar to optic flow or stereo vision, while the second one consists in a geometric transformation commonly known as warping.

Compared to the case of scalar-valued images, like computerised tomography or standard magnetic resonance (MR) images, registration of tensor fields [16] involves a substantial additional difficulty: While deforming scalar-valued images comes down to a simple relocation of values in the image domain (which usually will entail a resampling and thus some interpolation), the values of tensor fields refer to the geometry of the underlying domain, and have therefore to be modified to conform with the transformation of the domain. The local deformation that takes effect here is given by the Jacobian of the location map. Moreover, this induced transformation of tensors does not only affect the warping step but needs to be accounted for already in establishing the correspondence map.

Our contribution. This paper is dedicated to a variational framework for the registration of tensor-valued images that combines the correspondence

and warping problems in a single minimisation task solved by a gradient descent procedure. A preliminary version of this framework has been presented at the CVPR workshop *Tensors in Image Processing and Computer Vision* [3].

In this approach, directional information is used throughout the method by consistently integrating reorientation information. The physical model underlying the relation between domain transformation and tensor reorientation is tailored to basic requirements of DT-MRI processing but not specifically adjusted to empirical (biological, physiological) parameters.

An advantage of the variational approach is that it makes the assumptions of the model explicit and permits to manipulate them in a highly flexible way. In this sense our model is generic: For example, the physical model for tensor deformation could well be replaced with a refined model based on the physiology of neural tissue, or likewise with a completely different transformation behaviour if required so by the physical nature of some tensor imaging application.

We present here the variational model and demonstrate experimentally its potential as a versatile tool for image registration of tensor data.

Related work. Different approaches to DT-MRI registration have been proposed in the literature. One group of approaches treats the correspondence and warping problems independently. In this case, the reorientation is often entirely avoided in the correspondence part by relying only on derived quantities that do not involve directional information [12, 15, 19, 13]. From the so obtained transformation of the image domain one has then to extract the reorientation information for the warping step. Based on a single-fibre model, it is often assumed herein that the shape of the tensors remains unchanged, such that only a rotation needs to be extracted. Strategies for this problem are discussed in [1]: first, the finite-strain method that relies on a polar decomposition of the local deformation matrix of the domain, and second, the *preservation of principal directions* approach in which the effect of the transformation on the principal eigenvector of each tensor determines the local reorientation. A direct application of the local deformation to the tensors, which does not preserve the shapes of the tensors, is discussed in [13]. An overview of these techniques is found in [7].

An algebraic approach to DT-MRI registration is proposed in [8]. Here, tensor transformation is restricted to rotations, and the correspondence problem is solved by a local optimisation of transformation parameters. Furthermore, a multi-scale refinement and the inclusion of feature correspondences are discussed. A feature-based registration approach is also found in [14].

An approach that integrates directional information already into the computation of correspondences can be found in [17, 18]. It uses a diffusion tensor constancy constraint that combines an affine domain transformation with a rotation of the tensor values by polar decomposition, and fits a piecewise affine transformation to the image data. A gradient descent method proposed in [5] is based on a variational model in which tensor reorientation follows the action of the deformation field on the principal eigenvector. In [6] a variational model is proposed that incorporates the deformation by the Jacobian of the displacements in the data term.

In [11], registration based on direction-insensitive quantities is compared with one using the matrix entries. For the registration of DT-MRI data with data from other imaging modalities, one can use mutual information approaches, see [9]. Here again the warping has to be performed as a separate step. Also fluid dynamics PDEs have been used to attack registration problems of this type [4].

Structure of the paper. Our variational model is presented in Section 2. We develop the energy functional in detail and give an outline of the minimisation procedure. Section 3 is devoted to the experimental evaluation. A series of experiments in 2D on synthetic and DT-MRI data demonstrates basic features of the method. By two 3D experiments the viability of the approach for 3D data is verified, and a comparison to a finite-strain based registration procedure is made. A summary and outlook is presented in Section 4.

2 The Variational Model

Assume we are given two tensor fields $R, T : \mathbb{R}^d \rightarrow \text{SPD}(d)$ where $\text{SPD}(d)$ denotes the cone of symmetric positive definite real $d \times d$ -matrices. The tensor field R is the reference image, while T is the template image. We aim at recovering from these images two maps: first, the *displacement field* $u : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that encodes the transformation of the image domain \mathbb{R}^d , and second, the *reorientation field* $P : \mathbb{R}^d \rightarrow G$ that governs the corresponding transformation of the tensor values. Here, $G \subset \text{GL}(d, \mathbb{R})$ is the group of admissible transformation matrices.

This separation serves two purposes. On one side, it allows the variational formulation to be simple, because it is not necessary in this way to solve the equations that relate both fields for one of them. On the other side, it keeps the model generic in the sense that a different relationship between the two fields that may be imposed in a different physical setting can easily be plugged into the same model.

2.1 Energy Functional

Following this idea, our variational registration model requires minimisation of an energy functional

$$E[u, P] = D[u, P] + w_1 C[u, P] + w_2 S_u[u] + w_3 S_P[P] \quad (1)$$

with four components:

- a *data term* D that enforces the match between the transformation fields and the tensor fields,
- the *compatibility term* C that encodes the physically motivated relation between the displacement and reorientation fields,
- the *displacement smoothness term* S_u that implements a regularity assumption on the displacement field, and
- the *reorientation smoothness term* S_P that promotes regularity of the reorientation field.

Each term penalises violations of the respective model requirement by a nonnegative value. The weights w_1 , w_2 , and w_3 balance the influences of the different terms. Let us now discuss these four terms in detail.

Data term. The requirement that under the displacement and reorientation as prescribed by u and P the template T should become as similar as possible to the reference R is expressed by the *diffusion tensor constancy constraint (DTCC)*

$$(P^T)^{-1}T(x + u(x))P^{-1} = R(x) . \quad (2)$$

For the further mathematical derivations it is advantageous to distribute the actions of u and P to the two tensor fields, i.e., $T(x + u) = P^T R P$, see also [8, 17].

Using the Frobenius norm $\|\cdot\|_F$ of the matrix difference, we formulate eventually

$$D[u, P] = \int_{\Omega} \|T(x + u) - P^T(x)R(x)P(x)\|_F^2 dx , \quad (3)$$

where $\Omega \subset \mathbb{R}^d$ is the image domain.

Compatibility term. This term depends on the physical model for the relationship between displacements and tensor reorientation. Pure translations do not require reorientation of tensors. The local deformation of the domain that needs to be accounted for is expressed by the Jacobian $J(x + u)$ of the

location map $x \mapsto x + u$. Geometrically, its influence can be decomposed into three components: rotation, shearing, and scaling.

With focus on DT-MRI data, we assume that scaling of the domain does not affect the tensors, since absolute diffusivities depend rather on the physiology of the tissue on a microscopic scale. In contrast, it is natural to demand that the local rotation should act on the tensors by appropriately rotating their eigensystems.

The most difficult component from the modelling viewpoint is shearing. In the case of DT-MRI, the way how diffusion tensors in tissue change when this tissue is sheared may depend on the underlying fibre composition and the physiology of the fibres, and can therefore not be inferred exactly from the pure second-order tensor model. Based on the assumption that only small shears take place, we propose two versions of the model: one in which the shearing component of $J(x + u)$ acts directly on the tensors, and one in which no shearing at all takes place. The difference between the two lies in the choice of the matrix group G : For the first model it is the group of all matrices with determinant one, $G = \text{SL}(d, \mathbb{R})$, while for the second model we choose the group of (proper) rotation matrices $G = \text{SO}(d)$.

In both cases, the actual compatibility between the displacement field and the reorientation field is measured by

$$C[u, P] = \int_{\Omega} \|(P J(x + u))^p - \det(J(x + u))^{p/d} I\|_{\text{F}}^2 dx, \quad (4)$$

where I is the unit matrix. For the exponent p , the standard value in this paper will be 1.¹ This term expresses that P should equal the inverse of $J(x + u)$ except for its scaling component that is given by the determinant. In the models without shear, the least-squares error from fitting the rotation matrix $P \in \text{SO}(d)$ to the rescaled Jacobian from $\text{SL}(d, \mathbb{R})$ also accounts for the shear.

Smoothness terms. Measured data are inevitably contaminated with noise, an overfitting to which needs to be avoided. This is achieved by regularity conditions imposed on the transformation fields.

To enforce regularity of the vector field u , we use the Euclidean norm $\|\cdot\|_2$ of its spatial gradient, leading to

$$S_u[u] = \int_{\Omega} \sum_{j=1}^d \|\partial_j u\|_2^2 dx. \quad (5)$$

¹In our paper [3], we used $p = 2$ for 2D tensor registration, which seemed to be advantageous in early experiments. However, the case $p = 1$ is algorithmically simpler, and experimental exploration has proven its equal qualitative performance. One comparison is presented in Section 3.

Similarly, we impose regularity on the reorientation by

$$S_P[P] = \int_{\Omega} \sum_{j=1}^d \|\partial_j P\|_{\mathbb{F}}^2 dx . \quad (6)$$

2.2 Gradient Descent

To minimise the functional (1), we use a gradient descent procedure. To this end, we employ the Euler-Lagrange formalism from the calculus of variations to obtain the variational derivatives of E with respect to u and P .

We must, however, take into account that u and P are already multi-channel quantities, thus we need to specify suitable parametrisations. This is easy for the displacement field u as its values belong to a vector space. Here, the displacements u_1, \dots, u_d in the coordinate directions serve as parameters. For P , however, we have to parametrise the Lie group $\text{SO}(d)$ or $\text{SL}(d, \mathbb{R})$. We restrict our discussion here to the 2D case with and without shear, and the 3D case without shear. Assuming small deformations (P close to I), the following parameter sets can be used:

- *Case $\text{SO}(2)$ (2D, pure rotations)*: one rotation angle α , such that

$$P(\alpha) = \begin{pmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{pmatrix} . \quad (7)$$

- *Case $\text{SL}(2, \mathbb{R})$ (2D, rotations and shear)*: The action of P is decomposed into two rotations, with a shear inbetween that leaves the x axis fixed. Representing the rotation angles as $\frac{1}{2}(\alpha \pm \beta)$, we arrive at the parametrisation

$$P(\alpha, \beta, \gamma) = \begin{pmatrix} \cos \frac{\alpha+\beta}{2} & -\sin \frac{\alpha+\beta}{2} \\ \sin \frac{\alpha+\beta}{2} & \cos \frac{\alpha+\beta}{2} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 2\gamma & 1 \end{pmatrix} \begin{pmatrix} \cos \frac{\alpha-\beta}{2} & -\sin \frac{\alpha-\beta}{2} \\ \sin \frac{\alpha-\beta}{2} & \cos \frac{\alpha-\beta}{2} \end{pmatrix} \quad (8)$$

in which α corresponds to the “net rotation”, β gives the shear direction and γ the shear magnitude.

- *Case $\text{SO}(3)$ (3D, pure rotations)*: Making use of Euler angles α, β, θ , we parametrise P by

$$P(\alpha, \beta, \theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \alpha & -\sin \alpha \\ 0 & \sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{pmatrix} \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix} . \quad (9)$$

Denoting by F the integrand of the energy functional, the gradient descent in 2D with the artificial time t reads

$$\begin{aligned}
\frac{\partial}{\partial t}u_1 &= \frac{d}{dx}F_{u_{1x}} + \frac{d}{dy}F_{u_{1y}} - F_{u_1} , \\
\frac{\partial}{\partial t}u_2 &= \frac{d}{dx}F_{u_{2x}} + \frac{d}{dy}F_{u_{2y}} - F_{u_2} , \\
\frac{\partial}{\partial t}\alpha &= \frac{d}{dx}F_{\alpha_x} + \frac{d}{dy}F_{\alpha_y} - F_{\alpha} , \\
\frac{\partial}{\partial t}\beta &= \frac{d}{dx}F_{\beta_x} + \frac{d}{dy}F_{\beta_y} - F_{\beta} , \\
\frac{\partial}{\partial t}\gamma &= \frac{d}{dx}F_{\gamma_x} + \frac{d}{dy}F_{\gamma_y} - F_{\gamma} ,
\end{aligned} \tag{10}$$

where the last two equations are omitted when no shear is used. Analogously, the equations in the 3D case read

$$\begin{aligned}
\frac{\partial}{\partial t}u_1 &= \frac{d}{dx}F_{u_{1x}} + \frac{d}{dy}F_{u_{1y}} + \frac{d}{dz}F_{u_{1z}} - F_{u_1} , \\
\frac{\partial}{\partial t}u_2 &= \frac{d}{dx}F_{u_{2x}} + \frac{d}{dy}F_{u_{2y}} + \frac{d}{dz}F_{u_{2z}} - F_{u_2} , \\
\frac{\partial}{\partial t}u_3 &= \frac{d}{dx}F_{u_{3x}} + \frac{d}{dy}F_{u_{3y}} + \frac{d}{dz}F_{u_{3z}} - F_{u_3} , \\
\frac{\partial}{\partial t}\alpha &= \frac{d}{dx}F_{\alpha_x} + \frac{d}{dy}F_{\alpha_y} + \frac{d}{dz}F_{\alpha_z} - F_{\alpha} , \\
\frac{\partial}{\partial t}\beta &= \frac{d}{dx}F_{\beta_x} + \frac{d}{dy}F_{\beta_y} + \frac{d}{dz}F_{\beta_z} - F_{\beta} , \\
\frac{\partial}{\partial t}\theta &= \frac{d}{dx}F_{\theta_x} + \frac{d}{dy}F_{\theta_y} + \frac{d}{dz}F_{\theta_z} - F_{\theta} .
\end{aligned} \tag{11}$$

To compute these gradient descents, we use explicit (forward Euler) discretisations. Higher order schemes bear no advantage as we aim at a steady state. We discretise spatial derivatives by Sobel operators.

Note that the evaluation of $T(x+u)$ involves a resampling which is performed here by bilinear interpolation. While some interpolation artifacts implied by this proceeding could be reduced by using a shape interpolation approach, see e.g. [2], bilinear interpolation appears to be closer to the physical sampling process which also leads to partial volume effects in the measurements.

Multiscale procedure. A disadvantage of the gradient descent approach is that it can be trapped in local minima, preventing it particularly from




Figure 1: Colour coding schemes for visualisation of tensor field transformations. *Left: (a)* Displacement vectors in 2D. *Right: (b)* Scale for rotation angles and shear factors.

capturing correctly large displacements (i.e., larger than 1 to 2 pixels). Even when it converges to find large displacements, it is slow in this case. To improve convergence in this situation, we use a simple multiscale procedure: We perform the registration for a sequence of Gaussian-smoothed image pairs $(R_k, T_k), (R_{k-1}, T_{k-1}), \dots, (R_0, T_0)$ where $R_0 = R, T_0 = T$ and $R_i = G_{\sigma_i} * R_0, T_i = G_{\sigma_i} * T_0$ with Gaussians G_{σ_i} of decreasing standard deviation $\sigma_k > \sigma_{k-1} > \dots > \sigma_1 > 0$, initialising the transformation fields in step $i < k$ with the results of the preceding step $i + 1$.

3 Experiments

We start with experiments in 2D. To visualise our results, we represent tensor fields by ellipse glyphs where the directions and lengths of the principal axes of each ellipse represent eigenvectors and eigenvalues of the corresponding symmetric positive definite tensor. Colours encode the dominant eigenvector direction. Displacement fields are converted into colour images using the colour coding scheme shown in Figure 1(a). Angles α and shear factors γ of the reorientation field are represented in separate colour images using the scale given in Figure 1(b).

Weight parameters have been chosen manually in all of our experiments. Further investigation will be necessary to address a possible automatic choice of these parameters.

3.1 Synthetic 2D Experiment

In our first experiment (Figure 2), we use a synthetic 2D data set containing a “fibre” as reference image (a). The template image (b) is obtained from it by a consistent 3° rotation applied to the domain and tensors. A caveat




Figure 2: Registration of a synthetic tensor field. (a) *Top left*: Clipping (rows 0 to 19, columns 54 to 73) from a 127×127 synthetic tensor field featuring a simplified “fibre”, used as reference image. (b) *Top middle*: Template image obtained by 3° rotation of (a). (c) *Top right*: Result of variational registration with deformation model involving rotation and shear. (d) *Second row, left*: Displacement field computed in the registration process (all 127×127 pixels). (e) *Second row, middle*: Rotation angle α from the reorientation field. (f) *Second row, right*: Shear factor γ . (g) *Bottom left*: Registration result using the deformation model involving only rotations. (h) *Bottom middle*: Displacement field computed in registration with rotation-only model. (i) *Bottom right*: Rotation angle α from rotation-only model. Inserts in (e), (f) and (i) show correct parameter values (3° , 0.0).

about this test image pair is that the correspondence map is dramatically underdetermined

For registration, we set the weights in (1) as $w_1 = 2 \cdot 10^6$, $w_2 = 5 \cdot 10^7$, $w_3 = 4 \cdot 10^7$. Our complete 2D model with rotation and shear yields the registered image (c). No multiscale procedure has been used here. In (d–f) we show the displacement field and two components of the reorientation field. Due to the mentioned underdetermination of the example, the model attains an optimum with nonzero shear. In particular, the vertically periodic structures visible in (e, f) indicate that the process slightly overfits and matches aliasing-type structures generated by the resampling of the rotated data. Subfigures (g–i) show the results of registration with the rotation-only model.

As a measure of registration quality we use the value of the data term $D[u, P]$ for the computed displacement and reorientation fields. By this means, we can compare different registration models. As this measure could, however, be misled by overfitting, its utilisation is restricted to cases with equal smoothness parameters. Table 1 comprises values of the data term for registration with and without shearing and, for comparison, for a model that does not account for reorientation as well as for the known correct displacement and reorientation field (ground truth). The nonzero error measured in the latter case is caused by the resampling in generating the data and in the registration itself.

As expected, the model with reorientation reduces the error considerably. Since by construction the data fit a pure rotation model, it is naturally that the gain by including shear is smaller. Nevertheless, the underdetermination of the correspondence allows a further improvement by a nonzero shear.²

²In [3], where $p = 2$ was used in the compatibility term, data term measurements for the registration results were significantly lower than those reported here. It is only in this synthetic example that such a large difference between $p = 1$ and $p = 2$ is observed. It can be ascribed to the dramatic underdetermination of the correspondence map by the given data.

Table 1: Data term values measured to assess registration quality for the fibre images from Figure 2(a,b) by different models.

Model with rotation and shear	$2.14 \cdot 10^4$
Model with rotation only	$3.04 \cdot 10^4$
Model without reorientation	$4.60 \cdot 10^4$
Ground truth	$1.29 \cdot 10^4$




Figure 3: 127×127 fields of 2D tensors extracted from two DT-MRI data sets of the same human brain. The data set in (a) was obtained with 10 times averaging, while that in (b) was measured without averaging, which leads to a great difference in noise level.

3.2 Real-World 2D Experiments

For our next three experiments, we use 2D tensor fields extracted from two DT-MRI data sets of the same human brain. The 2×2 tensors are made up by the components of the 3×3 tensors referring to the section plane. Figure 3 shows the images in a grey-value representation, where each of the 2×2 tiles of each image represents one matrix component.

We start by a setup similar to our synthetic experiment, registering the 2D DT-MRI image from Figure 3(b) to a rotated version. The original image serves as reference, and a copy rotated by 3° as template. Adapting to the higher noise level, we set $w_1 = 3.6 \cdot 10^7$, $w_2 = 9 \cdot 10^7$, $w_3 = 7.2 \cdot 10^8$. We use the multiscale procedure with $k = 4$ and standard deviations 4.0, 2.0, 1.0, 0.5. In Figure 4 we show reconstructed displacement and reorientation fields together with clippings of the template and registered image superposed to the reference image. Accuracy measurements are compiled in Table 2. It can be seen in the colour-coded display, Figure 4(e,f) that the angle and shear estimates in most of the image area approximate the correct (ground truth) values 3° and 0, respectively. Histograms shown in Figure 5 confirm this. Significant deviations occur only in those outer regions of the data set which are dominated by noise tensors close to zero.




Figure 4: Registration of a 127×127 field of 2D tensors extracted from a DT-MRI dataset of a human brain as reference image and a copy rotated by 3° as template. (a) *Top left*: Reference image, rows 55 to 66, columns 29 to 40 shown. (b) *Top middle*: Template image, same area. (c) *Top right*: Registered image, same area. (d) *Bottom left*: Displacement field. (e) *Bottom middle*: Rotation angle from the reorientation field. (f) *Bottom right*: Shear factor. Inserts in (e), (f) show ground truth parameter values.




Figure 5: Histograms of estimated deformation parameters for 2D tensors. (a) *Left*: Rotation angles from Figure 4(e). (b) *Right*: Shear factors from Figure 4(f).




Figure 6: Registration of two 127×127 fields of 2D tensors extracted from different DT-MRI data sets of a human brain. (a) *Left*: Colour-coded displacement field that shows an approximately constant shift between the two images. (b) *Middle*: Rotation angle from the reorientation field. (c) *Right*: Shear factor from the reorientation field. – Inserts in (b), (c) show ground truth values.

While the improvement by a model accounting for reorientation is again evident, the remaining mismatch is much larger than in the synthetic model due to the high spatial variation of the data and resulting large resampling error.

We turn now to register the two different 2D DT-MRI images onto each other. As both data sets were acquired during a single measurement, there is only a slight displacement, and virtually no reorientation between them. Using the same parameters as before, we obtain the values shown in Figure 6. In this case, there is no significant influence of the reorientation, which indicates that also no substantial overfitting takes place.

In our last 2D experiment, we use the same two 2D DT-MRI images as before but rotate the reference image by 3° . The weight parameters are the same as before. Results are found in Table 3 and the Figure 7. In the top row (standard setting), a good capture of rotation is achieved in the central region where most anisotropic tensors reside (white matter). One observes, however,

Table 2: Data term measurements for the registration of a 2D DT-MRI data set and its rotated version by different models.

Model with rotation and shear	$1.02 \cdot 10^6$
Model with rotation only	$1.70 \cdot 10^6$
Model without reorientation	$2.39 \cdot 10^6$
Ground truth	$0.89 \cdot 10^6$

that angles are somewhat underestimated. The middle row demonstrates the importance of the compatibility term in our model: By setting its weight to zero, the displacements are still fairly well estimated but deformations are estimated reasonably only in a part of the image domain. Finally, the bottom row shows a result obtained by the model from [3] with exponent $p = 2$ in the compatibility term. With a suitably adapted value of the compatibility weight w_1 , results of both models (top row and bottom row) look almost identical.

3.3 Experiments with 3D Data

In our first 3D experiment (Figure 8) we register a rotated volume taken from one of our DT-MRI data sets onto its unrotated counterpart, using the $SO(3)$ reorientation model. For the visualisation of exemplary data we use ellipsoid glyphs analogous to the ellipse glyphs mentioned above for the 2D case. For the transformation fields we apply the 2D colour coding within the planes shown. The results show that the homogeneous rotation field and corresponding displacements are captured well, again with an underestimation of reorientation angles, see also the histograms in Figure 9.

Our second 3D experiment (Figure 10) consists in registering equally sized volumes from both our DT-MRI data sets. First we use our variational model with $SO(3)$ reorientation. Similarly as in the registration of 2D sections, the displacement field is dominated by an almost constant translation, whereas the rotation angles range around zero, with absolute values not exceeding 0.1° , see also the histogram in Figure 11(a). In order to visualise the small deviations from zero, the angles have been amplified by a factor 10 before applying the colour scale.

For comparison, we register the same volumes by a different model: Gradient descent is performed only for the displacement field, while the reorientation component is computed by the finite-strain model in each iteration step. The resulting Euler angles are shown in the bottom row of Figure 10. It can be seen that the reorientation fields are less smooth than before and vary in

Table 3: Data term measurements for the registration of two 2D DT-MRI data sets, one of them rotated, by different models.

Model with rotation and shear	$1.16 \cdot 10^6$
Model with rotation only	$1.42 \cdot 10^6$
Model without reorientation	$1.51 \cdot 10^6$

fact up to absolute values of approx. 0.3° , compare also the histogram in Figure 11(b).

4 Summary and Outlook

We have introduced a flexible variational framework for the registration of tensor data, based on a functional involving data, compatibility, and smoothness terms. We have presented an experimental proof-of-concept with respect to its applicability on synthetic and real-world data.

A more extensive evaluation of the model in 3D and its comparison to other registration approaches are part of our ongoing work. We also aim at investigating the use of the framework with other physical compatibility models.

Acknowledgement

We thank O. Gruber and I. Henseler from Saarland University Hospital, Homburg, for providing the DT-MRI measurements. The visualisation tools for tensor images were written by S. Didas and L. Pizarro.

References

- [1] D. C. Alexander, C. Pierpaoli, P. J. Basser, and J. C. Gee. Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging*, 20:1131–1139, 2001.
- [2] V. Arsigny, P. Fillard, X. Pennec, and N. Ayache. Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine*, 20(2):411–421, 2006.
- [3] S. Barbieri, M. Welk, and J. Weickert. Variational registration of tensor-valued images. In *Proc. CVPR Workshop on Tensors in Image Processing and Computer Vision*, Anchorage, Alaska, USA, 2008.
- [4] C. Brun, N. Lepore, X. Pennec, Y.-Y. Chou, O. L. Lopez, H. J. Aizenstein, J. T. Becker, A. W. Toga, and P. M. Thompson. Comparison of standard and Riemannian fluid registration for tensor-based morphometry in HIV/AIDS. In *MICCAI Workshop on Statistical Registration: Pair-Wise and Group-Wise Alignment and Atlas Formation*, 2007.

- [5] Y. Cao, M. I. Miller, S. Mori, R. L. Winslow, and L. Younes. Diffeomorphic matching of diffusion tensor images. In *CVPRW '06: Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*, page 67, Washington, DC, USA, 2006. IEEE Computer Society.
- [6] O. Faugeras, C. Lenglet, T. Papadopoulo, and R. Deriche. Non rigid registration of diffusion tensor images. Technical Report 6104, INRIA Sophia-Antipolis, France, 2007.
- [7] J. C. Gee and D. C. Alexander. Diffusion-tensor image registration. In J. Weickert and H. Hagen, editors, *Visualization and Processing of Tensor Fields*, pages 327–342. Springer, Berlin, 2006.
- [8] A. Goh and R. Vidal. Algebraic methods for direct and feature based registration of diffusion tensor images. In H. Bischof, A. Leonardis, and A. Pinz, editors, *Computer Vision – ECCV 2006, Part III*, volume 3953 of *Lecture Notes in Computer Science*, pages 514–525. Springer, Berlin, 2006.
- [9] C. A. Kemper. Incorporation of diffusion tensor MRI in non-rigid registration for image-guided neurosurgery. Master’s thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, Massachusetts, USA, 2003.
- [10] J. Modersitzki. *Numerical Methods for Image Registration*. Oxford University Press, Oxford, 2004.
- [11] H.-J. Park, M. Kubicki, M. E. Shenton, A. Guimond, R. W. McCarley, S. E. Maier, R. Kikinis, F. A. Jolesz, and C.-F. Westin. Spatial normalization of diffusion tensor MRI using multiple channels. *NeuroImage*, 20(4):1995–2009, Dec. 2003.
- [12] G. K. Rohde, S. Pajevic, C. Pierpaoli, and P. J. Basser. A comprehensive approach for multi-channel image registration. In J. C. Gee, J. B. A. Maintz, and M. W. Vannier, editors, *Biomedical Image Registration*, volume 2717 of *Lecture Notes in Computer Science*, pages 214–223. Springer, Berlin, 2003.
- [13] J. Ruiz-Alzola, C.-F. Westin, S. K. Warfield, C. Alberola, S. Maier, and R. Kikinis. Nonrigid registration of 3D tensor medical data. *Medical Image Analysis*, 6(2):143–161, 2002.

- [14] R. Sierra. Nonrigid registration of diffusion tensor images. Master's thesis, Dept. of Electrical Engineering, ETH, Zürich, Switzerland, 2001.
- [15] R. Verma and C. Davatzikos. Matching of diffusion tensor images using Gabor features. In *Proc. IEEE International Symposium on Biomedical Imaging*, pages 396–399, Arlington, VA, USA, 2004. IEEE Computer Society Press.
- [16] J. Weickert and H. Hagen, editors. *Visualization and Processing of Tensor Fields*. Springer, Berlin, 2006.
- [17] H. Zhang, P. Yushkevich, and J. Gee. Deformable registration of diffusion tensor MR images with explicit orientation optimization. In J. Duncan and G. Gerig, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2005*, volume 3749 of *Lecture Notes in Computer Science*, pages 725–732. Springer, Berlin, 2005.
- [18] 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.
- [19] H. Zhang, P. A. Yushkevich, and J. C. Gee. Towards diffusion profile image registration. In *Proc. IEEE International Symposium on Biomedical Imaging*, pages 324–327, Arlington, Virginia, USA, 2004.




Figure 7: Registration of two 127×127 fields of 2D tensors extracted from different DT-MRI data sets of a human brain. The test images were the same as in Figure 6 but with the reference image rotated by 3° . *Top row*: Standard model with exponent $p = 1$ and parameters $w_1 = 3.6 \cdot 10^8$, $w_2 = 9 \cdot 10^7$, $w_3 = 7.2 \cdot 10^8$. (a) *Top left*: Displacement field showing a rotation superposed to the constant shift. (b) *Top middle*: Rotation angle. (c) *Top right*: Shear factor. *Middle row*, (d)–(f): Same but with suppressed compatibility term, $w_1 = 0$. While the displacement field (d) is almost identical to (a), both deformation components are substantially misestimated. *Bottom row*, (g)–(i): Model with $p = 2$ (compare [3]) and parameters $w_1 = 3.6 \cdot 10^7$, $w_2 = 9 \cdot 10^7$, $w_3 = 7.2 \cdot 10^8$. Note that the compatibility weight w_1 is rescaled to accommodate the change in the compatibility energy. – Inserts in angle and shear fields show ground truth values.




Figure 8: Registration of 3D tensor data. A $40 \times 40 \times 40$ volume from a human brain DT-MRI data set serves as reference image. The template is a copy of the same data, rotated first by 5° around the x -axis and then by 5° around the y -axis. The reorientation model allows only rotations. *Top row:* (a) *Left:* Voxels 20 to 35 in x direction, 5 to 20 in y direction from layer 30 in z direction of the reference volume. (b) *Middle:* Same voxels from the template volume. (c) *Right:* Same voxels from the registered volume. *Second row:* (d) *Left:* Displacement field in the x - y central plane. (e) *Middle:* Displacement field in the x - z central plane. (f) *Right:* Displacement field in the y - z central plane. *Bottom row:* Euler angles in the x - y central plane. (g) *Left:* Angle α . (h) *Middle:* Angle β . (i) *Right:* Angle θ . Inserts in (g)–(i): ground truth.




Figure 9: 3D tensor registration results. (a) *Left*: Histogram of Euler angles α from Figure 8(g). (b) *Right*: Histogram of Euler angles β estimated by finite-strain model, see Figure 8(h).




Figure 10: Registration of 3D tensor data. Reference and template were $40 \times 40 \times 40$ regions of the 3D data sets from which also the 2D data in Figure 3 have been extracted. A $40 \times 40 \times 40$ volume from one human brain DT-MRI data set is registered onto an equally sized volume from a different DT-MRI measurement of the same subject. *Top row:* (a) *Left:* Displacement field in the x - y central plane. (b) *Middle:* Displacement field in the x - z central plane. (c) *Right:* Displacement field in the y - z central plane. *Second row:* Euler angles in the x - y central plane, amplified by a factor 10 (i.e., the colour scale from Figure 1(b) covers the interval $[-0.5^\circ, 0.5^\circ]$). (d) *Left:* Angle α . (e) *Middle:* Angle β . (f) *Right:* Angle θ . *Bottom row:* Same Euler angles for a registration based on a finite-strain model (see text). (g) *Left:* Angle α . (h) *Middle:* Angle β . (i) *Right:* Angle θ .




Figure 11: Comparison of 3D tensor registration results. (a) *Left*: Histogram of Euler angles β estimated by our variational method, see Figure 10(e). (b) *Right*: Histogram of Euler angles β estimated by finite-strain model, see Figure 10(h).