Log-Domain Diffeomorphic Registration of Diffusion Tensor Images

Andrew Sweet, Xavier Pennec

► To cite this version:

Andrew Sweet, Xavier Pennec. Log-Domain Diffeomorphic Registration of Diffusion Tensor Images. WBIR 2010 - International Workshop on Biomedical Image Registration , Jul 2010, Lübeck, Germany. pp.198-209, 10.1007/978-3-642-14366-3_18 . inria-00616163

HAL Id: inria-00616163

<https://hal.inria.fr/inria-00616163>

Submitted on 20 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Log-Domain Diffeomorphic Registration of Diffusion Tensor Images

Andrew Sweet¹ and Xavier Pennec¹

Asclepios, INRIA Sophia-Antipolis, France
`andrew.sweet@sophia.inria.fr`

Abstract. Diffusion tensor imaging provides information about deep white matter anatomy that structural magnetic resonance images typically fail to resolve. Non-linear registration of diffusion tensor images, for which a few methods already exist, allows us to capture the deformations of these structures that would otherwise go unobserved. Here, we build on an existing method for diffeomorphic registration of diffusion tensor images, so that it fully incorporates the useful log-domain parameterization of diffeomorphisms. Initially, this allows us to easily include a registration symmetry constraint that is highly desirable for pair-wise registration. More importantly, the parameterization allows simple and proper calculation of statistics on the transformations obtained. We show that the symmetric log-domain method exhibits the most preferable trade-off between image correspondence and deformation smoothness on real data and also achieves the best recovery of synthetic warps.

Keywords. Diffusion tensor imaging, computational anatomy, diffeomorphic registration, exponential map.

1 Introduction and Motivation

Although linear registration allows proper visual comparison of images and can also account for subject movement during image acquisition, its major limitation is that it only accounts for features of global transformations between images, such as position, orientation and scale. Capturing local anatomical differences, such as the size or shape of particular brain regions, requires non-linear registration, which overcomes these limitations by allowing images to transform differently at each point.

The discipline of computational anatomy aims to use these non-linear transformations to compute deformation statistics of anatomical structures that can potentially account for biological variability within a population [1]. Therefore, any method of registration used in this discipline must be able to provide deformations that can easily be used for statistical analysis. One such method makes use of the diffeomorphic demons registration framework [2], described in Sect. 2, and can be adapted to directly estimate a vector space parameterization of the transformation [3], allowing simple statistical calculation.

In this work, we specifically consider non-linear registration of diffusion tensor images (DTIs), which represent the diffusion of water in the brain using a second-order symmetric tensor at each voxel [4]. DTI registration is of particular interest because it provides unique information about deep white matter structures, the deformations of which we propose may be more significant than changes observed from scalar image registration. In a purely mathematical sense, it is also the case that DTI registration can be considered better determined to estimate a transformation from one image to another simply because each voxel of a DTI contains a tensor defined by six unique values, as opposed to the single value of a scalar image. It is for both these reasons that we propose DTI to be a suitable choice for the study of computational anatomy.

Although the diffeomorphic demons framework has already been suitably developed for DTI registration [5], we contribute further developments in Sect. 3 by describing how the previously referenced parameterization, which has already been used for scalar image registration, is also feasible for DTI registration. This section continues by explaining how our method can easily incorporate a registration symmetry constraint, which ensures that the registration process is independent of the order of input images. This constraint has previously been shown as a desirable feature in non-linear pair-wise registration [6] and we also show, in Sect. 5, that it specifically improves the performance of our method as well. The experiments and results in this section also show that the parameterization alone does not significantly affect the registration process.

2 The Diffeomorphic Demons Framework

Throughout this work, it is assumed that there is a non-parametric spatial transformation s from a moving image, M , to a fixed image, F . The ‘demons’ framework, initially described in [7], provides one approach to find this transformation and has seen many developments. Its most recent general form [2] works by attempting to iteratively minimize an energy

$$E(F, M, s, c) = \sigma_i^{-2} \text{Sim}(F, M \circ c) + \sigma_x^{-2} \text{Dist}(s, c) + \sigma_T^{-2} \text{Reg}(s) \quad (1)$$

where c is a non-parametric transformation that should achieve point correspondences between the images, σ_i weights the uncertainty of the images, σ_x weights the spatial uncertainty between c and s and σ_T weights the spatial uncertainty of s alone. This means that solving (1) is equivalent to finding a small update transformation u to compose with the current one such that $c = s \circ u$. The introduction of the hidden variable, c , allows the energy to be split into two forms, each of which can be optimized alternately in the following scheme [8].

1. **Correspondence:** given the current s , find the c that minimizes

$$E_c(F, M, s, c) = \sigma_i^{-2} \text{Sim}(F, M \circ c) + \sigma_x^{-2} \text{Dist}(s, c). \quad (2)$$

2. **Regularization:** given the c found from step 1, find the s that minimizes

$$E_s(s, c) = \sigma_x^{-2} \text{Dist}(s, c) + \sigma_T^{-2} \text{Reg}(s). \quad (3)$$

The regularization step in the demons framework constrains the smoothness of s , but it does not guarantee the smoothness, or even existence, of its inverse transformation. However, for s to be a meaningful transformation in biomedical image registration it is desirable that finding an s from M to F implies the existence of a smooth s^{-1} from F to M . Formally, we require s to be a *diffeomorphism*. Such a constraint can be imposed by exploiting the log-Euclidean framework for diffeomorphisms [9]. This uses the log map to parameterize the update transformation as a stationary velocity field $\mathbf{u} = \log(u)$, which when put through the exponential map gives a diffeomorphic update $u = \exp(\mathbf{u})$.

Although we are free to choose the similarity, distance and regularization criteria in (1), it is typically the case that $\text{Sim}(F, M \circ c) = \|F - M \circ c\|^2$, $\text{Dist}(s, c) = \|s - c\|^2 = \|u\|^2$ and $\text{Reg}(s) = \|\nabla s\|_K^2$, where it should be noted that the regularization criterion makes use of a norm defined in a space K . These choices have the advantage that (2) has an approximate closed form solution that can be found independently at each point of a scalar image [2] and that (3) can be approximately solved by convolution with a Gaussian related to K [8]. Additionally, the number of free parameters in the model is typically reduced by defining $\sigma_i = \|F - M \circ c\|$.

With all these assumptions, a single iteration of the diffeomorphic demons algorithm consists of the two following steps

1. $\mathbf{u}^* = \underset{\mathbf{u}}{\text{argmin}} [E_c^{\text{diffeo}}(F, M, s, \mathbf{u})]$
 $= \underset{\mathbf{u}}{\text{argmin}} [\sigma_i^{-2} \|F - M \circ (s \circ \exp(\mathbf{u}))\|^2 + \sigma_x^{-2} \|\exp(\mathbf{u})\|^2]$
2. $s \leftarrow K_{\text{diff}} \star \exp(K_{\text{fluid}} \star \mathbf{u}^*)$

where $\star$ represents convolution, $K_{\text{fluid}} = \mathcal{G}[\mathbf{0}, \sigma_{\text{fluid}}^2 \mathbf{I}]$ is used to performed fluid-like regularization of the update to the transformation and $K_{\text{diff}} = \mathcal{G}[\mathbf{0}, \sigma_{\text{diff}}^2 \mathbf{I}]$ is used to perform diffusion-like regularization of the updated transformation. We use $\mathcal{G}[\mu, \Sigma]$ to denote a Gaussian distribution with mean μ and covariance Σ .

3 Diffeomorphic Demons Registration of DTIs

3.1 General DTI Registration

While the regularization step in the diffeomorphic demons framework only operates on vector fields, and is therefore independent of the image type, the correspondence step includes a few operations that must be explicitly defined for tensor images.

First of all, there should be a way to perform arithmetic on tensor images, so that the sum of squares similarity criterion can continue to be used. This is achieved by exploiting the log-Euclidean framework for tensors [10], which allows a tensor $T(n)$ at voxel n in image T to be parameterized by a log-tensor $\log(T(n))$. As log-tensors belong to a vector space, tensor arithmetic is accomplished by performing vector arithmetic on the log-tensors and exponentiating the result. Secondly, we must also define how to correctly warp the DTIs. As

we work with discrete images, there must be a way to interpolate tensors so that images can always be compared at the same points. Continuing to use the log-tensor representation means that interpolation is simply achieved by linear component-wise vector interpolation, an approach which has exhibits reasonable performance [10].

However, non-rigid transformations of tensor images also causes the local orientation of the tensors to be lost. The orientation of diffusion tensors is vital, as it is groups of locally aligned tensors that represent anatomical white matter structures in the brain. There are two possible reorientation schemes that attempt to correct for this [11]. The preservation of principal direction (PPD) approach finds the rotation matrix, $R(n)$, which ensures that the principal axis of $T(n)$ is the same as it was before the warp was applied. By contrast, the finite strain (FS) approach finds the $R(n)$ that minimizes the Frobenius distance to the local transformation Jacobian matrix $J(n)$. In practice, this minimization is computed from the polar decomposition of the Jacobian, so that $R(n) = (J(n)J(n)^T)^{-\frac{1}{2}}J(n)$. In both cases, the corrected tensor is given by $T'(n) = R(n)^T T(n) R(n)$ and the log-tensor is corrected in the same way to give $\log(T'(n)) = R(n)^T \log(T(n)) R(n)$.

Given the advantages of using log-tensors in all these aspects of DTI registration, we compute these beforehand, use them in the registration, then take their exponential to produce the final warped image. Therefore, all references to the fixed and moving images, F and M , hereby refer to the log-tensor images.

3.2 DT-REFinD

The DT-REFinD method [5], hereby denoted as DTR, provides a way of performing diffeomorphic demons registration of tensor images and importantly incorporates FS reorientation directly into the energy to be minimized

$$E_c^{\text{DTR}}(F, M, s, \mathbf{u}) = \sigma_i^{-2} \|F - R^T(M \circ (s \circ \exp(\mathbf{u})))R\|^2 + \sigma_x^{-2} \|\exp(\mathbf{u})\|^2 \quad (4)$$

where R can be thought of as an image of rotation matrices, which specifies the tensor reorientation at each voxel. The direct incorporation of the reorientation into the optimization has been shown to improve performance compared to a scheme where orientation is simply corrected after each standard update. Following the notation of [5], we can express the correspondence energy as

$$E_c^{\text{DTR}}(F, M, s, \mathbf{u}) = \left\| \begin{matrix} \varphi_1(F, M, s \circ \exp(\mathbf{u})) \\ \varphi_2(\mathbf{u}) \end{matrix} \right\|^2 = \|\varphi_c(F, M, s, \mathbf{u})\|^2 \quad (5)$$

where $\varphi_1(F, M, s \circ \exp(\mathbf{u})) = \sigma_i^{-1} [F - R^T(M \circ (s \circ \exp(\mathbf{u})))R]$, $\varphi_2(\mathbf{u}) = \sigma_x^{-1} \exp(\mathbf{u})$ and $\varphi_c(F, M, s, \mathbf{u}) = [\varphi_1(F, M, s \circ \exp(\mathbf{u})), \varphi_2(\mathbf{u})]^T$.

In general, the demons algorithm approximates this energy by the 0th and 1st order terms of its Taylor expansion with respect to $\mathbf{u}$

$$E_c^{\text{DTR}}(F, M, s, \mathbf{u}) \approx \left\| \begin{bmatrix} \varphi_1(F, M, s \circ \exp(\mathbf{0})) \\ \varphi_2(\mathbf{0}) \end{bmatrix} + \begin{bmatrix} D\varphi_1(F, M, s \circ \exp(\mathbf{0})) \\ D\varphi_2(\mathbf{0}) \end{bmatrix} \mathbf{u} \right\|^2 \quad (6)$$

which means that finding $\mathbf{u}^*$ by minimizing (6) is equivalent to solving the least squares problem $\|\mathbf{b} - \mathbf{A}\mathbf{u}\|^2$ where $\mathbf{b} = [\varphi_1; \varphi_2]$ and $\mathbf{A} = -[D\varphi_1; D\varphi_2]$. For scalar image registration, the matrix $\mathbf{A}$ has a block diagonal structure which means that each vector component of $\mathbf{u}^*$ can be individually solved. However, DTR’s incorporation of FS reorientation into the energy means that $\mathbf{A}$ is sparse, but no longer has a simple block diagonal structure because the optimal displacement of a tensor at one point affects the reorientation, and therefore optimal displacement, of a tensor at a neighboring point. Accordingly, the current implementation of DTR finds a solution using a least squares conjugate gradients solver in Gmm++ [12], but is memory and processor intensive due to the large size of the sparse matrix $\mathbf{A}$.

4 Log-Domain Diffeomorphic Registration of DTIs

4.1 Log-Domain DT-REFinD

On first appearances, the basic diffeomorphic demons framework satisfies the need for statistical computation because the log map of the final transformation should be its stationary velocity field representation. However, in practice the log map acts as a high-pass filter, and exhibits a lack of stability between the transformation and its stationary velocity field representation [13]. To overcome this problem an alternative approach [3] extends the demons framework by not only parameterizing the update field in the ‘log-domain’, but the current transformation as well, so that the update is directly applied to a stationary velocity field $\mathbf{v}$ whose exponential is s . This direct update on the velocity field is achieved using the Baker-Campbell-Hausdorff function, denoted as $Z(\cdot)$, to approximate the composition for small updates such that $\exp(Z(\mathbf{v}, \epsilon\mathbf{u})) \approx \exp(\mathbf{v}) \circ \exp(\epsilon\mathbf{u})$ [14]. As a result, the two iterative steps of log-domain DTR (LDDTR) are

1. $\mathbf{u}^* = \underset{\mathbf{u}}{\operatorname{argmin}} [E_c^{\text{DTR}}(F, M, \exp(\mathbf{v}), \mathbf{u})]$
2. $\mathbf{v} \leftarrow K_{\text{diff}} \star Z(\mathbf{v}, K_{\text{fluid}} \star \mathbf{u}^*)$

where it should be noted that E_c^{DTR} is the same as in the original algorithm. In other words, exactly the same method can be used to find $\mathbf{u}^*$ because the log-domain approach only affects the representation of the transformation. However, the consequence of using this representation directly is that the regularization is defined for the stationary velocity field $\mathbf{v}$ which parameterizes the transformation, rather than the transformation itself. More specifically, the regularization energy is modified to become $\text{Reg}(s) = \|\nabla \log(s)\|_K^2 = \|\nabla \mathbf{v}\|_K^2$.

4.2 Symmetric Log-Domain DT-REFinD

The diffeomorphic parameterization in the general framework ensures s^{-1} exists and is smooth, but the classic demons algorithm only ever finds the forwards transformation $s : M \rightarrow F$. Running the algorithm with the images exchanged will certainly produce a diffeomorphism $t : F \rightarrow M$, but there is no guarantee

that $t = s^{-1}$. This registration symmetry, or inverse consistency, can be easily achieved in the log-domain by performing two unconstrained optimisations of the current transformation and projecting these onto a new symmetric transformation [3]. The first is the same as usual and finds an update $\mathbf{u}_{\text{forw}}^*$ for s , while the second finds an update $\mathbf{u}_{\text{back}}^*$ for s^{-1} with the fixed and moving images exchanged. As DTR can be performed in the log-domain, the same approach can be exploited to define the two steps of symmetric LDDTR (SLDDTR)

1. (a) $\mathbf{u}_{\text{forw}}^* = \underset{\mathbf{u}}{\operatorname{argmin}} [E_c^{\text{DTR}}(F, M, s \circ \exp(\mathbf{v}), \mathbf{u})]$
(b) $\mathbf{u}_{\text{back}}^* = \underset{\mathbf{u}}{\operatorname{argmin}} [E_c^{\text{DTR}}(M, F, s \circ \exp(-\mathbf{v}), \mathbf{u})]$
2. $\mathbf{v} \leftarrow \frac{1}{2} K_{\text{diff}} \star (Z(\mathbf{v}, K_{\text{fluid}} \star \mathbf{u}_{\text{forw}}^*) - Z(-\mathbf{v}, K_{\text{fluid}} \star \mathbf{u}_{\text{back}}^*))$

where it should be noted that R will be different for the steps 1a and 1b and the second term in step 2 is negated in order to invert the updated inverse transformation, so that the average of both updated transformations is found in the forwards direction. Note that projection of the two updated transformations onto the space of symmetric transformations is simply performed by averaging because the representative stationary velocity fields belong to a vector space.

5 A Comparison of DTI Registration in and out of the Log-Domain

In order to implement the LDDTR and SLDDTR methods described in Sect. 4.1 and Sect. 4.2 respectively, we have adapted the implementation of symmetric and non-symmetric log-domain diffeomorphic demons registration of scalar images [15], so that it uses the implementation of the energy function from [5].

5.1 Validation on Real Data

Diffusion weighted imaging data are provided by the Neuradapt study group and the authors would like to acknowledge M. Vassallo, C. Lebrun and S. Chanalet for making these available. Here we consider a sub-group of 7 subjects from this study. For each subject, a single unweighted ($b = 0$) was acquired along with 23 gradient weighted ($b = 700\text{s/mm}^2$) images with data dimensions of $256 \times 256 \times 26$ and spatial dimensions of $0.9375\text{mm} \times 0.9375\text{mm} \times 5.5\text{mm}$. While the anisotropy of the spatial dimensions is particularly high, we still believe that the data can highlight any differences between the registration methods considered.

DTI reconstructions are performed assuming the usual log-Gaussian noise model and any non-positive tensors, which are physically meaningless, are replaced with a local tensor mean [16]. Each subject's $b = 0$ image is linearly registered to that of the 2mm ICBM-DTI-81 template [17] using the affine version of the robust method described in [18], which is available in [19]. The resultant affine transformations are applied to their corresponding DTIs, using FS reorientation. Finally, the brain extraction toolkit [20] is used to generate a brain

foreground mask from each $b = 0$ image, which is applied to the affinely registered DTI to remove any tensors outside of the brain. Although this tool is primarily designed for use on T1 images, we find that the masks generated using the T2 weighted $b = 0$ images are reasonable after a small erosion.

Every possible unique pair-wise registration is performed between the subjects' DTIs, giving 42 registrations for the non-symmetric methods and 21 for SLDDTR. All algorithms are allowed to iterate ten times using $\sigma_x = 1$, which is enough to ensure reasonable convergence of the solutions. A multi-resolution scheme is not used because it does not significantly improve the convergence or performance, possibly due to the initial affine registration. As LDDTR and SLDDTR optimize a different regularization energy compared to DTR, we register over a range of regularization parameter values $\sigma_T = \{0.6, 0.8, \dots, 2.0\}$ so that the final transformations produced by each of the methods can be compared at a range of harmonic energies (HEs). The HE is defined as the mean square Frobenius norm of the transformation Jacobian and therefore corresponds to the irregularity of the transformation.

Figure 1 demonstrates that at low HEs, the mean square error between the log-tensor images, referred to from here as the log mean square error (LMSE), is relatively similar for all of the methods. One difference is that LDDTR produces a higher LMSE than DTR at the same HE. This difference is accentuated at higher HEs. By contrast, the SLDDTR method achieves a lower LMSE than DTR at the same HE, although the difference is less clear at higher HEs. These observations suggest that using the log-domain parameterization directly has a detrimental effect on performance. Yet the same parameterization also allows easy incorporation of the symmetric constraint, which seems to be beneficial.

Fig. 1: The mean log-tensor image square error plotted against the mean transformation harmonic energy for registration of 42 subject pairs using the DTR, LDDTR and SLDDTR methods.

Figure 2 shows an example of a single registration performed using all three methods with a single reasonable regularization parameter $\sigma_T = 1.4$. The visual

correspondence of the warped DTIs these produce illustrate the similarity in performance.

Fig. 2: An example registration from the DTI of one subject (a) to another (b) using the DTR (c) (LMSE=0.280, HE=0.0623), LDDTR (d) (LMSE=0.199, HE=0.0881) and SLDDTR (e) (LMSE=0.200, HE=0.0742) methods using $\sigma_T = 1.4$. For (a-e) the image intensity represents fractional anisotropy and the color represents the principal axis of the tensor where red=left-right, green=posterior-anterior and blue=inferior-superior. In all cases, the same mid-axial slice is displayed using MedINRIA [19].

5.2 Performance on Synthetic Warps

In order to quantitatively compare the performance of the methods with a known ground truth, we create three random diffeomorphisms for each subject, apply them to the DTI of their respective subject and add noise to the warped DTIs. A single noisy warped DTI is generated according to the following scheme.

1. Create a random velocity field $\mathbf{v}_r$ by sampling a vector for each foreground voxel in the original DTI from $\mathcal{G}[\mathbf{0}, \sigma_r^2 \mathbf{I}]$.
2. Convolve $\mathbf{v}_r$ with $\mathcal{G}[\mathbf{0}, \sigma_s^2 \mathbf{I}]$ to give a smooth random velocity field $\mathbf{v}_s$.
3. Exponentiate $\mathbf{v}_s$ to give a random diffeomorphism $s_s = \exp(\mathbf{v}_s)$.

4. Warp the original DTI with s_s using FS reorientation.
5. Add noise drawn from $\mathcal{G}[\mathbf{0}, \sigma_n^2 \mathbf{I}]$ to the log-tensors in the warped DTI.

For our experiments, we find that using $\sigma_r^2 = 10^4$, $\sigma_s^2 = 7.25^2$ and $\sigma_n^2 = 0.005$ produces warps with similar properties to those found from pair-wise registration of the real data. Specifically, the mean displacement of the random warps is 3.403mm and their mean harmonic energy is 0.0915. Although the noisy warped DTIs are not necessarily anatomically believable, as demonstrated by a single example in Fig. 3, they do provide an opportunity to validate and explain previous observations from experiments on real data.

Fig. 3: An example of a synthetic warp, represented here by its application to a regular grid image (a) (HE=0.0936), applied to the DTI of one subject (b) to produce a warped DTI to which noise is added (c) (LMSE=0.271). For the DTIs, the intensity and color maps are described in Fig. 2. In all cases, the same mid-axial slice is displayed using MedINRIA [19].

All registration methods are applied from the original image to the noisy warped images for each subject using the same range of regularization parameter values as before $\sigma_T = \{0.6, 0.8, \dots, 2.0\}$. The DTR and LDDTR methods are also applied from the noisy images to the originals. As a ground truth is present in this experiment, we additionally consider the distance from the recovered deformation field to the true one $\text{dist}(s, s_{\text{true}}) = \|s - s_{\text{true}}\|$, as well as the distance between their Jacobians $\text{dist}(J(s), J(s_{\text{true}})) = \|J(s) - J(s_{\text{true}})\|$.

Figure 4 shows that at low HEs, all three methods exhibit very similar LMSEs. In accordance with this result, the distances between the transformations and their Jacobians are also relatively similar. At higher HEs, DTR produces a slightly lower LMSE than the LDDTR and SLDDTR methods, but this actually represents an increase in the distance from the true transformation.

The distance from the true transformations appears to be optimal for all methods at an HE of around 0.07. Here, the LDDTR method exhibits slightly better recovery of the true transformation than DTR, but this may occur simply because the synthetic warps really are parameterized by the velocity fields

assumed by LDDTR. Yet the SLDDTR method, which makes the same assumption, recovers the true transformation even better than LDDTR at all HEs, which illustrates that the symmetric constraint is beneficial for DTI registration in the same way that has been previously demonstrated for scalar image registration [3]. Despite this, the Jacobian distances are relatively similar for all methods at all HEs. This suggests that the advantage of the symmetric constraint comes from capturing information in the transformation that is not locally linear.

Fig. 4: The mean log-tensor image square error (a), mean distance between the recovered and true transformations (b) and mean distance between their Jacobians (c) plotted against the mean transformation harmonic energy for registration of 42 subject pairs using the DTR, LDDTR and SLDDTR methods.

6 Conclusions and Further Work

In this work, we have demonstrated that the log-domain parameterization of diffeomorphisms can be fully incorporated into the demons framework for non-linear registration of DTIs. While directly applying this parameterization with

little regularization may reduce registration performance, this can be counteracted by incorporating an inverse consistency constraint into the method. Furthermore, this extra constraint seems to improve performance a little beyond that of the original method.

Although the immediate contributions of this work are not revolutionary, they provide the essential foundations for proper statistical analysis of the structural deformations observable from DTI registration. In fact, the data used in this study are only a small subset from a larger collection of 180 subjects, pairwise registration of which is currently being undertaken. The statistics from these registrations have the potential to demonstrate major modes of deformation and could also be reincorporated back into the method for statistical regularization through non-stationary convolution.

Other important variations on the DTR method are also possible. For example, [5] explains that as the DTR method must solve a large system anyway, we might as well directly solve the DTR equivalent of (1), thereby reducing the dependence on the regularization approximations made in the original demons framework. Additionally, we might also consider other ways of incorporating the symmetric constraint into the DTR method, such as [21] who find a single update from a symmetric energy for log-domain registration of scalar images. Although this means we could avoid solving two systems as in the projection approach, it means the single system to solve is larger and may be more easily subject to local minima. Time permitting, all these variations will be investigated prior to the future work described above, so that all the implications of DTI registration in the log-domain can be fully understood.

References

1. Pennec, X.: Statistical computing on manifolds: from riemannian geometry to computational anatomy. In Nielsen, F., ed.: *Emerging Trends in Visual Computing*. Volume 5416 of LNCS. (2008) 347–386
2. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* **45**(1, Supp.1) (2009) S61–S72
3. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Symmetric log-domain diffeomorphic registration: A demons-based approach. In Metaxas, D., Axel, L., Fichtinger, G., Székely, G., eds.: *Proc. MICCAI 2008*. Volume 5241 of LNCS. (2008) 754–761
4. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. *Biophysical journal* **66**(1) (1994) 259–267
5. Yeo, B.T.T., Vercauteren, T., Fillard, P., Peyrat, J.M., Pennec, X., Golland, P., Ayache, N., Clatz, O.: DT-REFinD: Diffusion tensor registration with exact finite-strain differential. *IEEE Transactions on Medical Imaging* **28**(12) (2009) 1914–1928
6. Klein, A., Andersson, J., Ardekani, B.A., Ashburner, J., Avants, B., Chiang, M.C., Christensen, G.E., Collins, L.D., Gee, J., Hellier, P.: Evaluation of 14 nonlinear deformation algorithms applied to human brain MRI registration. *NeuroImage* **46**(3) (2009) 786–802

7. Thirion, J.P.: Image matching as a diffusion process: An analogy with Maxwell's demons. *Medical Image Analysis* **2**(3) (1998) 243–260
8. Cachier, P., Bardinet, E., Dormont, D., Pennec, X., Ayache, N.: Iconic feature based nonrigid registration: The PASHA algorithm. *Computer Vision and Image Understanding* **89**(2-3) (2003) 272–298
9. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-Euclidean framework for statistics on diffeomorphisms. In Larsen, R., Nielsen, M., Sporring, J., eds.: *Proc. MICCAI 2006*. Volume 4190 of LNCS. (2006) 924–931
10. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Log-Euclidean metrics for fast and simple calculus on diffusion tensors. *Magnetic Resonance in Medicine* **56**(2) (2006) 411–421
11. Alexander, D.C., Pierpaoli, C., Basser, P.J., Gee, J.C.: Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Transactions on Medical Imaging* **20**(11) (2001) 1131–1139
12. Renard, Y., Pommier, J., Fournie, M., Schleimer, B.: Gmm++: http://home.gna.org/getfem/gmm_intro.html
13. Hernandez, M., Olmos, S., Pennec, X.: Comparing algorithms for diffeomorphic registration: Stationary LDDMM and diffeomorphic demons. In Pennec, X., Joshi, S., eds.: *Proc. MFCA 2008*. (2008) 24–35
14. Bossa, M., Hernandez, M., Olmos, S.: Contributions to 3D diffeomorphic atlas estimation: Application to brain images. In Ayache, N., Ourselin, S., Maeder, A., eds.: *Proc. MICCAI 2007*. Volume 4791 of LNCS. (2007) 667–674
15. Dru, F., Vercauteren, T.: An ITK implementation of the symmetric log-domain diffeomorphic demons algorithm. *The Insight Journal* January–June 2009 (2009)
16. Fillard, P., Pennec, X., Arsigny, V., Ayache, N.: Clinical DT-MRI estimation, smoothing and fiber tracking with log-Euclidean metrics. *IEEE Transactions on Medical Imaging* **26**(11) (2007) 1472–1482
17. Mori, S., Oishi, K., Jiang, H., Jiang, L., Li, X., Akhter, K., Hua, K., Faria, A.V., Mahmood, A., Woods, R., Toga, A.W., Pike, G.B., Neto, P.R., Evans, A., Zhang, J., Huang, H., Miller, M.I., van Zijl, P., Mazziotta, J.: Stereotaxic white matter atlas based on diffusion tensor imaging in an ICBM template. *NeuroImage* **40**(2) (2008) 570–582
18. Ourselin, S., Roche, A., Prima, S., Ayache, N.: Block matching: A general framework to improve robustness of rigid registration of medical images. In DiGioia, A.M., Delp, S., eds.: *Proc. MICCAI 2000*. Volume 1935 of LNCS. (2000) 557–566
19. Toussaint, N., Souplet, J.C., Fillard, P.: MedINRIA: Medical image navigation and research tool by INRIA. In: *Proc. MICCAI 2007 Workshop on Interaction in Medical Image Analysis and Visualization*. (2007)
20. Smith, S.M.: Fast robust automated brain extraction. *Human Brain Mapping* **17**(3) (2002) 143–155
21. Sabuncu, M.R., Yeo, B.T.T., van Leemput, K., Vercauteren, T., Golland, P.: Asymmetric image-template registration. In Yang, G.Z., Hawkes, D.J., Rueckert, D., Noble, A., Taylor, C., eds.: *Proc. MICCAI 2009*. Volume 5761 of LNCS. (2009) 565–573