



HAL
open science

A Log-Euclidean Statistical Analysis of DTI Brain Deformations

Andrew Sweet, Xavier Pennec

► **To cite this version:**

Andrew Sweet, Xavier Pennec. A Log-Euclidean Statistical Analysis of DTI Brain Deformations. MICCAI 2010 Workshop on Computational Diffusion MRI, Sep 2010, Beijing, China. hal-01525407

HAL Id: hal-01525407

<https://inria.hal.science/hal-01525407>

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.

A Log-Euclidean Statistical Analysis of DTI Brain Deformations

Andrew Sweet¹ and Xavier Pennec¹

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

Abstract. Diffusion tensor images (DTIs) provide information about deep white matter anatomy that structural magnetic resonance images typically fail to resolve. Non-linear registration of DTIs provides a way to capture the deformations of these structures that would otherwise go unobserved. Here we use an existing method that fully incorporates a useful vector space parameterization of diffeomorphisms, thereby allowing simple and well defined calculation of deformation statistics. An initial analysis of the statistics produced by registration of a group of 37 HIV/AIDS patients illustrates principal modes of deformation that are anatomically meaningful and that corroborate with previous findings. The registration method is developed by incorporating these modes into a statistical regularization criterion. Even though initial results suggest this new criterion over-constrains the registration method, we discuss plausible ways to address this.

1 Introduction and Motivation

Computational anatomy aims to use transformations, produced by non-linear registration, to compute deformation statistics of anatomical structures that can account for biological variability within a population [1]. A first requirement of the registration method is that the result should be easily usable in subsequent statistical analyses. A second requirement is that the images the method can register should contain information that is rich enough to correctly describe anatomically meaningful deformations. In this work, we specifically consider DTIs, which represent the diffusion of water in the brain using a second order symmetric tensor at each voxel [2]. DTI registration is of particular interest because it provides unique information about major deep white matter structures, the deformations of which we propose may be more significant than changes observed from scalar image registration. This is likely to be especially true in HIV/AIDS patients, where significant white matter changes have previously been reported [3,4].

Both these requirements are met by the log-domain diffeomorphic demons algorithm [5], which directly estimates a vector space parameterization of a diffeomorphism, and has been adapted for use on DTIs [6] as described in Sec. 2. We then explain how, in Sec. 3, first and second order statistics can be computed for inter-subject registration of a group subjects. Further developments to the

registration method are proposed in the same section, by describing how these statistics can be reintegrated into a regularization criterion. An initial analysis demonstrates, in Sec. 4, that the statistics describe anatomically meaningful modes of deformation. However, experiments detailed in this section show that the statistical regularization criterion over-constrains the registration in some areas of the brain. We conclude by discussing, in Sec. 5, how these current problems might be overcome and propose ways in which the statistics could be used to constrain registration of other modalities.

2 Log-domain Diffeomorphic Registration of DTIs

2.1 A Log-Euclidean Parameterization of Diffeomorphisms

Throughout this work, it is assumed there is a non-parametric spatial transformation s from a moving image, M , to a fixed image, F . When the images are of different subjects, it is desirable that s should be a *diffeomorphism* which respects an *inverse consistency* constraint [7]. The log-domain diffeomorphic demons method [5] provides an approach to find an s that respects these constraints by exploiting the log-Euclidean framework for diffeomorphisms [8]. This framework considers the space of diffeomorphisms that can be identified by tangent vectors at the identity transformation. Each tangent vector defines a smooth stationary velocity field \mathbf{v} , which is related to a diffeomorphism through the exponential map $s = \exp(\mathbf{v})$. The main advantage of this parameterization is that \mathbf{v} lies in a vector space, so that proper arithmetic and statistical analysis can be easily performed. Additionally, the negated velocity field provides simple access to the inverse transformation $s^{-1} = \exp(-\mathbf{v})$.

2.2 Log-Domain Diffeomorphic Demons Registration

The log-domain diffeomorphic demons registration method [5] works by attempting to iteratively minimize an energy

$$E(F, M, \mathbf{v}, \mathbf{w}) = \text{Sim}(F, M \circ \exp(\mathbf{w})) + \text{Dist}(\mathbf{v}, \mathbf{w}) + \text{Reg}(\mathbf{v}) \quad (1)$$

where $c = \exp(\mathbf{w})$ is a non-parametric spatial transformation that attempts to achieve point correspondences between F and M . The introduction of the hidden variable, \mathbf{w} , allows the energy to be split into two forms, each of which can be optimized alternately in the following scheme [9].

1. **Correspondence:** given the current \mathbf{v} , find the \mathbf{w} that minimizes

$$E_{\text{corr}}(F, M, \mathbf{v}, \mathbf{w}) = \text{Sim}(F, M, \exp(\mathbf{w})) + \text{Dist}(\mathbf{v}, \mathbf{w}). \quad (2)$$

2. **Regularization:** given the \mathbf{w} found from step 1, find the \mathbf{v} that minimizes

$$E_{\text{reg}}(\mathbf{v}, \mathbf{w}) = \text{Dist}(\mathbf{v}, \mathbf{w}) + \text{Reg}(\mathbf{v}). \quad (3)$$

Solving (2) is equivalent to finding a small update transformation $u = \exp(\mathbf{u})$ to compose with the current one such that $c = s \circ u$. One usually defines the similarity and distance criteria as $\text{Sim}(F, M, \exp(\mathbf{w})) = \sigma_i^{-2} \|F - M \circ \exp(\mathbf{w})\|^2$ and $\text{Dist}(s, c) = \sigma_d^{-2} \|\mathbf{v} - \mathbf{w}\|^2$ respectively, where σ_i weights the uncertainty of the images and is typically defined as $\|F - M \circ \exp(\mathbf{w})\|$, whereas σ_d weights the spatial uncertainty between \mathbf{v} and \mathbf{w} and therefore controls the size of the update \mathbf{u} . The inverse consistency constraint, which we use throughout this work, can be imposed by defining a similarity criterion that is independent of the image order $\text{Sim}(F, M, \exp(\mathbf{w})) = \sigma_i^{-2} (\|F - M \circ \exp(\mathbf{w})\|^2 + \|M - \exp(-\mathbf{w}) \circ F\|^2)$. These choices have the advantage that (2) has an approximate closed form solution that can be found independently at each point of an image.

In the absence of prior knowledge, a sensible choice for the regularization criterion would be one that penalizes the harmonic energy of the velocity field $\sigma_r^{-2} \|\nabla \mathbf{v}\|^2$, where σ_r^{-2} weights the spatial uncertainty of s alone. Instead, the demons method smooths the correspondence field to give the transformation velocity field $\mathbf{v} = \mathcal{G}[0, \sigma_d^{-2} \sigma_r^2 I] * \mathbf{w}$, which can be shown [10] to be the Tikhonov regularized solution of (3). This means that a harmonic energy criterion $\text{Reg}(\mathbf{v}) = \sigma_r^{-2} \|\nabla \mathbf{v}\|_K^2$ is only definable with respect to a Hilbert space K .

2.3 How to Account for Registration of DTIs

Like diffeomorphisms, symmetric second order tensors lie on a manifold, which makes properly defining the similarity criterion between two DTIs problematic. A definition can be made by exploiting the log-Euclidean framework for tensors [11], which achieves a good approximation of the Riemannian metric. In this case, a log-tensor, \mathbf{T} , which exists in the tangent space at the identity matrix, parameterizes a tensor T through the exponential map $T = \exp(\mathbf{T})$. This parameterization also allows proper linear interpolation of tensors images, which is required when resampling at the same regular points in both images after a discrete transformation has been applied to M .

A further difficulty of non-linear DTI registration is that non-rigid transformation of tensors causes their local orientation to be lost. There are two major reorientation schemes that attempt to correct for this: preservation of principal directions; and finite strain [12]. Here we only consider the finite strain approach because an analytic gradient of its effect can be used in the demons optimization and improves registration performance compared to a scheme where orientation is simply corrected after each standard update [13]. Furthermore, this method can be easily adapted to use the log-domain parameterization with little affect on registration performance [6].

3 Computation and Reuse of Deformation Statistics

3.1 Defining a Distribution of Group-Wise Deformations

While the log-domain parameterization allows desirable constraints to be respected, its most important benefit is that we can properly compute deformation

statistics by calculating vector statistics of the velocity fields. Imagine that we have a group of m subjects $\{S_1, \dots, S_m\}$ and want to find the set of $n = m(m-1)$ velocity fields $\mathbf{v}_{i,j}$ that represent the transformations from S_i to S_j . As the inverse consistency constraint is respected, only $n/2$ registrations need to be performed because each one gives access to the forwards and inverse transformations represented by \mathbf{v} and $-\mathbf{v}$ respectively. As this implies that the mean velocity field is guaranteed to be the zero field, the immediate topic of interest is how the n fields vary.

Consider that the number of variables in a field, p , is three times the number of voxels in an image. Ignoring background voxels means that $p \approx 10^6$ for a typical DTI with an isotropic spatial resolution of 2mm. We can usually expect DTI datasets to contain images from around $m \approx 10^2$ subjects, implying that $n \approx 10^4$. This suggests that estimating the $p(p+1)/2$ values of the covariance matrix $\Sigma_{\mathbf{v}}$ will be highly underdetermined. However, in the demons framework, the regularization criterion effectively imposes the velocity field to be drawn from $\mathcal{G}[0, \sigma_d^{-2} \sigma_r^2 I]$ subject to the constraints from the data. When $\sigma_d^{-2} \sigma_r^2$ is large enough, there should be a significant amount of spatial correlation in the field due to the locally imposed smoothness. This means that the effective number of variables, which we denote as p' , should be far less than p .

Even so, estimating the $p'(p'+1)/2$ values of $\Sigma_{\mathbf{v}}$ is still likely to be underdetermined and presents practical computation problems as well. In order to tackle both these difficulties, we perform principal components analysis (PCA) on the velocity fields to define a rank-reduced covariance matrix $\Sigma_{\mathbf{v}}^{(k)} = \mathbf{Y}^{(k)} \mathbf{\Lambda}^{(k)} \mathbf{Y}^{(k)\top}$, where the columns of $\mathbf{Y}^{(k)}$ are the k principal eigenvectors of $\Sigma_{\mathbf{v}}$, $\{\mathbf{y}_1, \dots, \mathbf{y}_k\}$, and the diagonal of $\mathbf{\Lambda}^{(k)}$ contains its k principal eigenvalues $\{\lambda_1, \dots, \lambda_k\}$. In this model, we can consider the projection of a velocity field \mathbf{v} onto the k major eigenvectors, which gives $\mathbf{v}^{(k)} = \sum_{i=1}^k (\mathbf{v}^\top \mathbf{y}_i) \mathbf{y}_i = \sum_{i=1}^k \alpha_i \mathbf{y}_i$. This assumes a Gaussian distribution on the fields where $P(\alpha_i) = \mathcal{G}[0, \lambda_i]$ and $P(\mathbf{v}^{(k)}) = \prod_{i=1}^k \mathcal{G}[0, \lambda_i]$.

Under the Baker-Campbell-Hausdorff (BCH) assumptions [5], the composition of two diffeomorphisms $\exp(\mathbf{a})$ and $\exp(\mathbf{b})$ can be approximated by the exponential of their addition $\exp(\mathbf{a}) \circ \exp(\mathbf{b}) \approx \exp(\mathbf{a} + \mathbf{b})$. Technically, this should only hold when \mathbf{b} has a small magnitude, but here, we investigate what happens empirically when this constraint is loosened. We do so because this means that the fields generated by registration from one particular subject to all $m-1$ others are approximate basis vectors for the space of all n fields (Fig. 1). Of course, if we choose a different reference subject, we will find a different set of approximate basis vectors. In this sense, performing PCA corresponds to defining an orthogonal basis that should not be biased towards any particular subject and that should reduce the effects of the BCH approximations in our model.

3.2 A Statistically Regularized Demons Method

Minimizing the regularization energy (3) in the demons framework can be seen as the equivalent of maximizing the posterior probability of \mathbf{v} given \mathbf{w} from

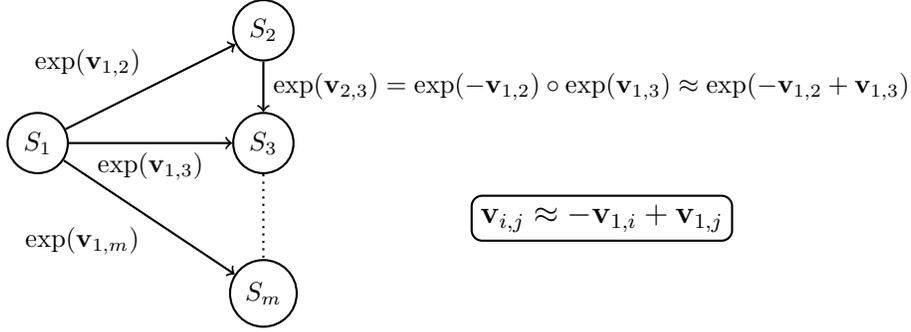


Fig. 1: Under our Baker-Campbell-Hausdorff approximations, the velocity fields that represent the transformations from one subject, here S_1 , to all other subjects, S_2, \dots, S_m , are basis vectors for the space of velocity fields that represent transformations between any two subjects.

the correspondence step. In this sense, the distance criterion is related to the likelihood, in that $\text{Dist}(\mathbf{w}, \mathbf{v}) \propto -\log(P(\mathbf{w}|\mathbf{v}))$, and the regularization criterion acts as a prior assumption on the distribution of \mathbf{v} , in that $\text{Reg}(\mathbf{v}) \propto -\log(P(\mathbf{v}))$. Under the usual assumptions, we have $P(\mathbf{w}|\mathbf{v}) = -\log(\mathcal{G}[\mathbf{v}, \sigma_d^2 \mathbf{I}])$ and $P(\mathbf{v}) = \mathcal{G}[\mathbf{0}, \sigma_r^2 (\nabla^* \nabla)^{-1}]_K$, the latter of which assumes no prior knowledge on where the field should be smooth or where there should be covariation.

Given our reduced-rank definition of $\Sigma_{\mathbf{v}}$, we can introduce such knowledge by replacing the usual prior on \mathbf{v} with a prior on the projection $P(\mathbf{v}^{(k)}) = \mathcal{G}[\mathbf{0}, \Sigma_{\mathbf{v}}^{(k)}]$. This yields a MAP estimate of

$$\mathbf{v}^{(k)} = \underset{\mathbf{v}^{(k)}}{\text{argmax}}[P(\mathbf{v}^{(k)}|\mathbf{w})] = \underset{\mathbf{v}^{(k)}}{\text{argmin}}[E_{\text{reg}}^{(k)}(\mathbf{v}^{(k)}, \mathbf{w})] \quad (4)$$

where

$$\begin{aligned} E_{\text{reg}}^{(k)}(\mathbf{v}^{(k)}, \mathbf{w}) &= -\log(\mathcal{G}[\mathbf{v}^{(k)}, \sigma_d^2 \mathbf{I}]) - \log\left(\prod_{i=1}^k \mathcal{G}[0, \lambda_i]\right) \\ &= \frac{1}{2} \sum_{i=1}^k \left(\sigma_d^{-2} (\alpha_i^2 - 2\alpha_i \beta_i) + \frac{\alpha_i^2}{\lambda_i} \right). \end{aligned} \quad (5)$$

Here $\beta_i = \mathbf{w}^T \mathbf{y}_i$ is the coordinate of the correspondence velocity field projected on the i^{th} eigenvector. The minimization of (5) can be performed independently on each projected coordinate to give $\alpha_i = (1 + \sigma_d^2 \lambda_i^{-1})^{-1} \beta_i$. Empirical results given in Sec. 4 show that when using a suitable value of k , it is the case that $\lambda_i \gg \sigma_d^2$, which means $\alpha_i \approx \beta_i$. Therefore, we approximately minimize (5) by simply projecting \mathbf{w} onto the eigenvectors of $\Sigma_{\mathbf{v}}$ to give $\mathbf{v}^{(k)}$.

4 Experiments and Results

4.1 Data, Pre-Processing and Registration Details

Diffusion weighted imaging data of a set of 37 HIV/AIDS patients 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. We use 30 subjects, denoted as group A, for statistical estimation, and reserve the other 7, denoted as group B, for validation. For each subject, a single unweighted ($b = 0$) image was acquired along with 23 gradient weighted ($b = 700 \text{ s/mm}^2$) images with spatial dimensions of $0.9375 \text{ mm} \times 0.9375 \text{ mm} \times 5.5 \text{ mm}$. Correction for subject motion and eddy currents is achieved by affinely registering the diffusion weighted images of each subject to their $b = 0$ image using FSL [14].

DTIs are reconstructed using a log-Gaussian noise model and non-positive tensors are replaced with a local mean [15]. The affine registration method described in [16] is used to register each subject's $b = 0$ image to that of the 2mm ICBM-DTI-81 template [17]. The affine transformations produced are applied to their corresponding DTIs, using finite strain reorientation, so that the DTIs lie in a common global coordinate system. Finally, the brain extraction toolkit [18] is used to generate a brain foreground mask from each affinely registered $b = 0$ image, which is applied to the affinely registered DTI to remove any tensors outside of the brain.

After pre-processing, non-linear registration is performed between each unique pair of subjects in groups A and B. All registrations use $\sigma_d = 1$ and are allowed to iterate ten times, which is enough to ensure reasonable convergence of the solutions. In order to generate results over a range of transformation regularities, registration is repeated using $\sigma_r = \{0.8, 1.2, \dots, 2\}$.

4.2 What Are the Major Modes of Deformation?

Eigenvectors and eigenvalues are estimated from the velocity fields generated in group A and we compute the i^{th} mode of deformation as $\exp(\sqrt{\lambda_i} \mathbf{y}_i)$. Figure 2 demonstrates that the first four major modes are anatomically meaningful in the sense that areas of significant displacement are concentrated in and around identifiable anatomical structures. We find that these modes are relatively consistent across the range of regularization parameters used. One unsurprising difference is that larger regularization produces smoother modes that are more spatially sparse. This suppression of probably noisy deformations swaps the order of the third and fourth modes at the extreme values of σ_r used.

The first mode exhibits large displacement in the area under the corpus callosum, which actually corresponds to an expansion in the ventricles, and correlates with a rotation around the cerebellum. The second mode represents an expansion in the cerebellum that causes a displacement of the brain stem, which both correlate with expansion in the cortical areas. The third mode, for $\sigma = 2$, exhibits fairly complicated deformation that is localized around the cerebellum and brain stem. The fourth, again for $\sigma = 2$, mostly accounts for an expansion of

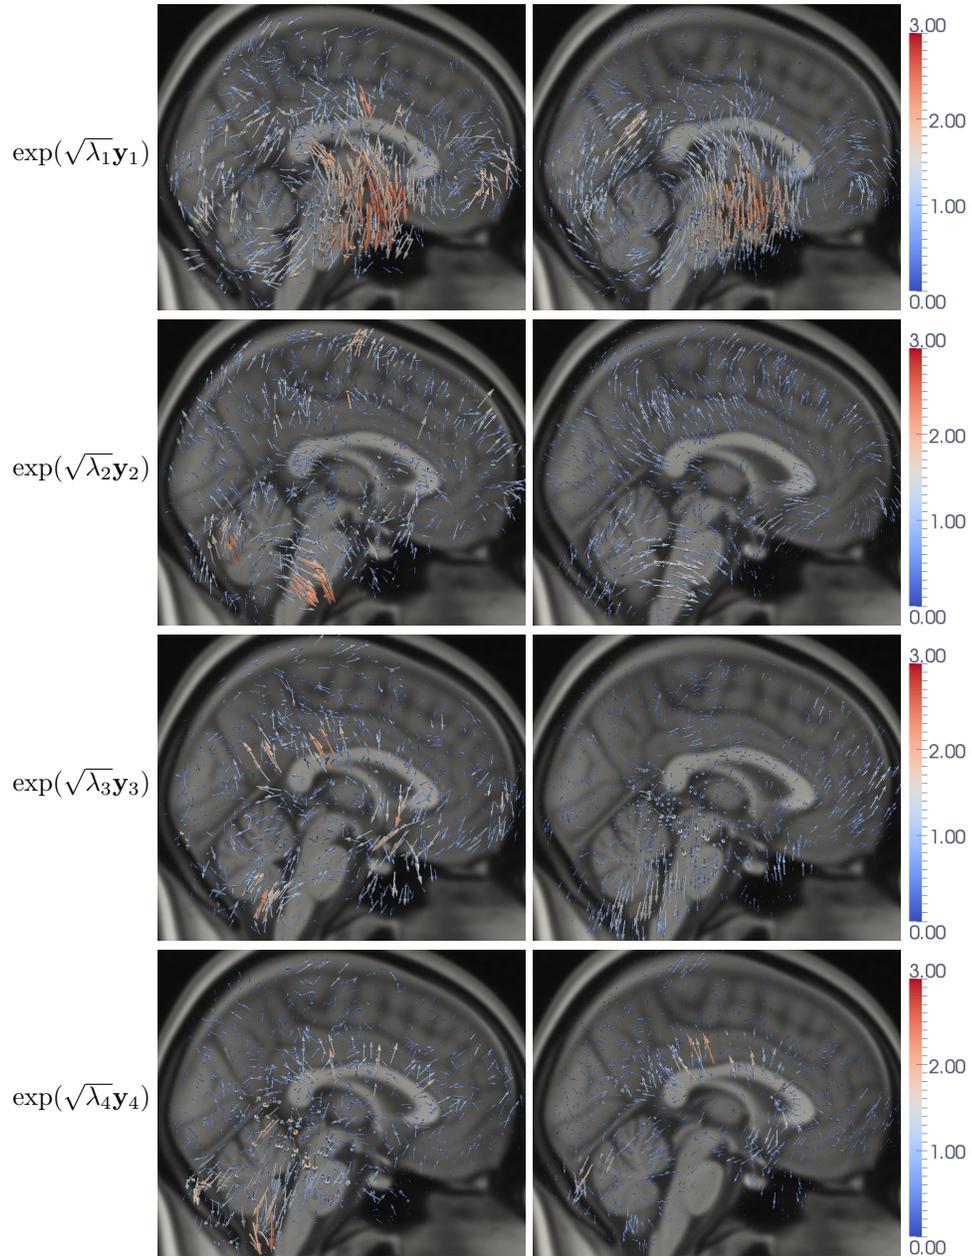


Fig. 2: A mid-sagittal view of the 4 major modes of displacement found from inter-subject DTI registration of group A using low (left) and high (right) values of $\sigma_r = 0.8, 2.0$. Arrows indicate displacement direction and color corresponds to displacement size in mm. The background is the T1 1mm MNI-152 template.

the ventricles that causes displacement above the corpus callosum, particularly close to the genu. It is of interest that there is large amount of deformation in and around the cerebellum because this is compatible with cerebellar changes previously found in HIV/AIDS [19].

Less regularization obviously produces more variance in the velocity fields, which causes larger eigenvalues (Fig. 3). Here, we see that the distribution of the first $m - 1$ eigenvalues is roughly exponential, but, as expected from our discussion in Sec. 3.1, the eigenvalues drop off sharply from here because the space of velocity fields is approximately $m - 1$ dimensional. The other $n - m + 1$ eigenvalues are roughly constant at all values of σ_r and are likely to be either noise or modes of deformation that are very specific to certain pairs of subjects. As σ_r is increased, we are better posing the problem of estimating the eigendecomposition given the number of data available. The small difference between distributions at $\sigma_r = 1.6$ and 2 suggests that using a value of around $\sigma_r = 1.6$ is suitable. However, even in this case, only around 80% of the variance in the data is accounted for by the first $m - 1$ eigenvalues. This may be due to the BCH approximations in our model or a non-Gaussian distribution of the velocity fields.

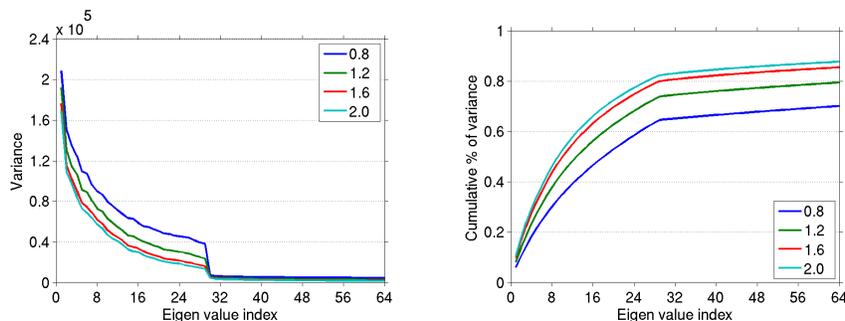


Fig. 3: The eigenvalues (left) and the cumulative percentage of variance they account for (right) from PCA performed on the 435 velocity fields found by inter-subject DTI registration of group A using $\sigma_r = \{0.8, 1.2, \dots, 2\}$.

In order to test whether or not the eigenvectors generalize well to other subjects, we calculate the mean square error between the velocity fields found from inter-subject registration of group B and the projections of these onto the eigenvectors $\frac{1}{n} \sum_{i=1}^n \mathbf{v}_i - \sum_{j=1}^k \alpha_j \mathbf{y}_j$. We see that as the number of eigenvectors is increased, this error approximately decreases linearly across all values of σ_r , which suggests that generalization is only reasonable (Fig. 4). Again, there is little difference between $\sigma_r = 1.6$ and 2. The localization of the errors for $\sigma_r = 1.6$ and $k = 29$, which is representative for all values of σ_r , is mostly in cortical regions rather than in major white matter structures. This suggests that the eigenvectors generalize well for deformations in deep central areas of the brain, but that deformations in outer areas are likely to be more subject specific.

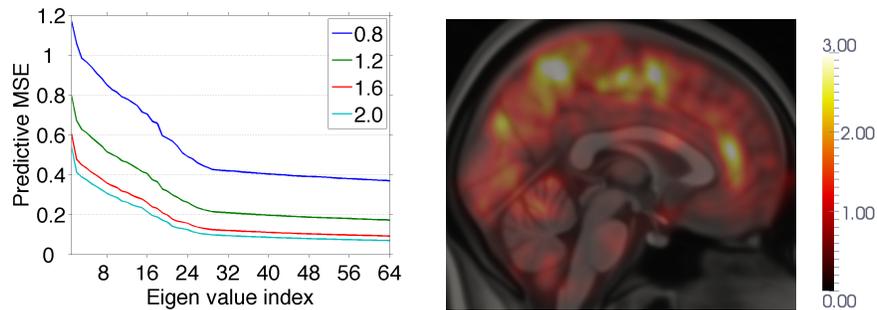


Fig. 4: The mean square error between the velocity fields in group B and the projections of these onto the eigenvectors estimated from group A using $\sigma_r = \{0.8, 1.2, \dots, 2\}$ (left). Also shown (right) is a mid-sagittal slice of the mean square error for $k = 29$ and $\sigma_r = 1.6$ overlaid on the T1 1mm MNI-152 template.

Using this approach, the error may be higher in cortical regions simply because the velocity is higher in these areas. To compensate for this potential bias, we consider the relative mean square error, where each error vector of each field is normalized by the square norm of the local velocity vector (Fig. 5). Although we observe some differences in how this error is distributed across σ_r , the localization of the error for $k = 29$ and $\sigma_r = 1.6$ is similar, in that high errors are mostly found close to the cortices and outside of major white matter structures.

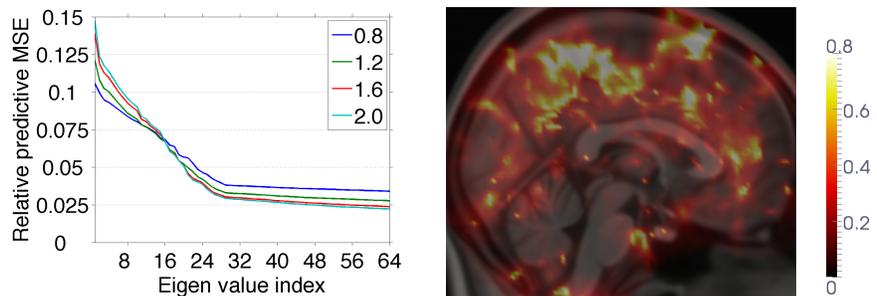


Fig. 5: The relative mean square error between the velocity fields in group B and the projections of these onto the eigenvectors estimated from group A using $\sigma_r = \{0.8, 1.2, \dots, 2\}$ (left). Also shown (right) is a mid-sagittal slice of the relative mean square error for $k = 29$ and $\sigma_r = 1.6$ overlaid on the T1 1mm MNI-152 template.

4.3 How Does Statistical Regularization Affect DTI Registration?

Registration is also performed on group B using the statistically regularized method with $k = \{4, 8, \dots, 48\}$ and the eigenvectors found from group A with $\sigma_r = 1.6$. Figure 6 shows that the mean square error between the registered DTIs using statistical regularization for all values of k is significantly higher than that using the original method, which suggests that the statistical prior knowledge is over-constraining the registration. The localization of this error (Fig. 7) shows that the statistically regularized method behaves similarly to the original method in some deep internal areas of the brain, such as the mid-body of the corpus callosum and the ventricles, but fails to account for the error in other areas, such as the genu, splenium and parts of the brain stem.

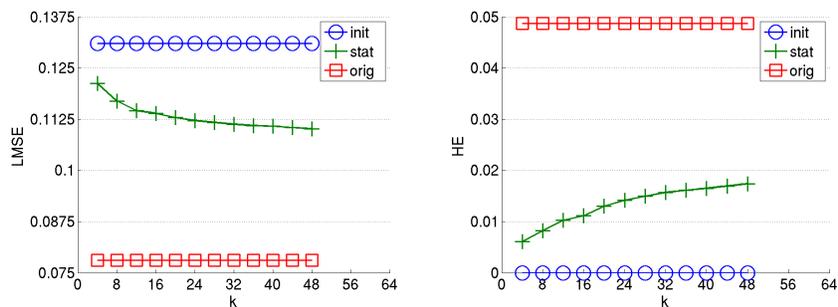


Fig. 6: The mean square error (left) between the registered DTIs and the harmonic energy of the transformation (right) using the statistically regularized method with $k = 4, 8, \dots, 48$ eigenvectors at $\sigma_r = 1.6$. For comparison the initial error and the final error for the original method at $\sigma_r = 1.6$ are also shown.

However, note that the harmonic energy of the statistically regularized transformation is considerably lower for all values of k (Fig. 6). It is possible that when compared at similar harmonic energies, the statistically regularized method may produce lower errors than the original method, although the current results cannot show this.

5 Conclusions and Further Work

Here, we present an investigation of the deformations between DTIs of a reasonably sized group of HIV/AIDS patients using log-domain diffeomorphic demons registration. The major principal components of the underlying transformations describe anatomically meaningful modes of deformation between deep structures in the brain, which generalize well to other patients. These encouraging initial results are not only interesting in their own right, but also help to validate the development of log-domain methods that directly parameterize diffeomorphic

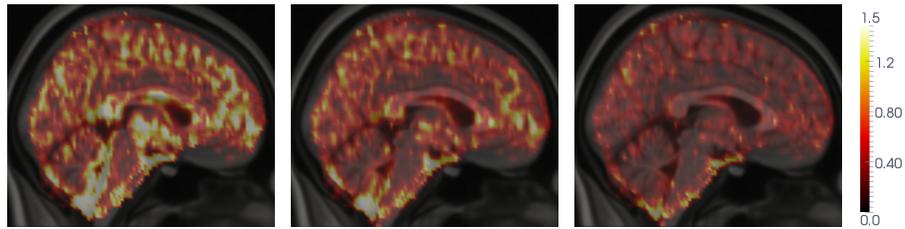


Fig. 7: The mean square error between the registered DTIs using the statistically regularized method with $k = 28$ eigenvectors at $\sigma_r = 1.6$ (middle) and the original method using $\sigma_r = 1.6$ (right). For comparison, the initial error is also shown (left). All errors are overlaid on T1 1mm MNI-152 template.

transformations. Although the incorporation of these modes into the statistical regularization criterion is perhaps less successful, we do observe that they reliably capture deformation in deep areas of the brain. As the formulation of the statistically regularized method is independent of the image modality, we could use DTI deformation statistics to better constrain T1 registration, which can prove difficult in these deep areas due to the lack of signal and contrast.

From a methodological point of view, performing PCA on the velocity fields amounts to computing the eigensystem from the sample covariance matrix, which is known to be unstable in high dimensions. This instability may explain why the statistical regularization method proposed here is overly restrictive, implying that alternatives to PCA should be considered. Alternatively, rather than project the correspondence field at each iteration of the demons algorithm, it may make more sense to directly optimize the coordinates of the velocity field $\{\alpha_1, \dots, \alpha_k\}$ on the eigenvectors. It could also be the case that the methodology is well defined, but that the registration cannot reliably capture deformations from DTIs in some areas of brain, such as the cortices. Further investigation should focus on exploring how stable the eigendecomposition as the size of the subject group is increased, although this may present problems in terms of computation time.

Acknowledgements This work was partially supported by the ANR Blanc Karametria program ANR-09-BLAN-0332 from the French Agence Nationale pour la Recherche.

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. Basser, P.J., Mattiello, J., LeBihan, D.: MR diffusion tensor spectroscopy and imaging. *Biophysical journal* **66**(1) (1994) 259–267
3. Sullivan, E.V., Pfefferbaum, A.: Diffusion tensor imaging in normal aging and neuropsychiatric disorders. *European Journal of Radiology* **45**(3) (2003) 244–255
4. Leporé, N., Brun, C., Chou, Y.Y., Chiang, M.C., Dutton, R.A., Hayashi, K.M., Luders, E., Lopez, O.L., Aizenstein, H.J., Toga, A.W., Becker, J.T., Thompson, P.M.: Generalized tensor-based morphometry of HIV/AIDS using multivariate statistics on deformation tensors. *IEEE Transactions on Medical Imaging* **27**(1) (2008) 129–141

5. 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
6. Sweet, A., Pennec, X.: Log-domain diffeomorphic registration of diffusion tensor images. In: Proc. of Workshop on Biomedical Image Registration 2010. Volume 6204 of LNCS., Lübeck, Germany (July 2010) 198–209
7. 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
8. 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
9. Cachier, P., Ayache, N.: Isotropic energies, filters and splines for vectorial regularization. *Journal of Mathematical Imaging and Vision* **20**(3) (2004) 251–265
10. Nielsen, M., Florack, L., Deriche, R.: Regularization, scale-space, and edge detection filters. *Journal of Mathematical Imaging and Vision* (1997)
11. 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
12. 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
13. 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
14. Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., Luca, M.D., Drobnjak, I., Flitney, D.E., Niazy, R.K., Saunders, J., Vickers, J., Zhang, Y., Stefano, N.D., Brady, J.M., Matthews, P.M.: Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* **23**(Supplement 1) (2004) S208 – S219 Mathematics in Brain Imaging.
15. 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
16. 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
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. Smith, S.M.: Fast robust automated brain extraction. *Human Brain Mapping* **17**(3) (2002) 143–155
19. Klunder, A.D., Chiang, M.C., Dutton, R.A., Lee, S.E., Toga, A.W., Lopez, O.L., Aizenstein, H.J., Becker, J.T., Thompson, P.M.: Mapping cerebellar degeneration in HIV/AIDS. *Neuroreport* **19**(17) (2008) 1655–1659