

Change Detection in Diffusion MRI Using Multivariate Statistical Testing on Tensors

Antoine Grigis^{1,2}, Vincent Noblet¹, Félix Renard^{1,2}, Fabrice Heitz¹,
Jean-Paul Armspach², and Lucien Rumbach²

¹ University of Strasbourg, CNRS, UMR 7005, LSIIT, France

² University of Strasbourg, CNRS, FRE 3280, LINC-IPB, France
grigis@unistra.fr

Abstract. This paper presents a longitudinal change detection framework for detecting relevant modifications in diffusion MRI, with application to Multiple Sclerosis (MS). The proposed method is based on multivariate statistical testings which were initially introduced for tensor population comparison. We use these methods in the context of longitudinal change detection by considering several strategies to build sets of tensors characterizing the variability of each voxel. These testing tools have been considered either for the comparison of tensor eigenvalues or eigenvectors, thus enabling to differentiate orientation and diffusivity changes. Results on simulated MS lesion evolutions and on real data are presented. Interestingly, experiments on an MS patient highlight the ability of the proposed approach to detect changes in non evolving lesions (according to conventional MRI) and around lesions (in the normal appearing white matter), which might open promising perspectives for the follow-up of the MS pathology.

1 Introduction

The automated detection of relevant changes in longitudinal Magnetic Resonance Imaging (MRI) sequences is crucial for medical diagnosis, follow-up and prognosis. The core problem is to identify image regions that are significantly different between two successive images. Because of its widespread interest and the large number of applications in diverse domains, change detection has been the topic of much interest. A review of the literature can be found in [1]. In conventional MRI, various methods have already been proposed to detect changes [2], [3], [4]. But few works have addressed change detection in Diffusion Tensor Imaging (DTI). Some previous works have addressed change detection in scalar images, characterizing diffusion properties, such as the Fractional Anisotropy or the Mean Diffusivity. These works rely either on statistical parametric testing, using the Generalized Likelihood Ratio Test (GLRT) [5], or on non parametric testing, for instance permutation testing [6]. A comparison of standard statistical testing on Diffusion Weighted (DW-) images is presented in [7]. In [8], the GLRT has been extended to diffusion tensor images, but the approach does not take into account the positive definite nature of matrices. This constraint has been

considered in [9] where tensor test statistics have been developed for population comparison. We propose to use this method for longitudinal change detection. To this end, we consider strategies to build sets of tensors characterizing the *variability* of each voxel. These strategies are based on the variability existing in the DW-images, or in the spatial neighborhood of the considered voxel, or a combination of these two. Based on the tensor model, we derive testing tools for the comparison of tensor eigenvalues and eigenvectors, thus enabling to differentiate orientation and diffusivity changes. The proposed framework is based on the following processing pipeline: 1- Preprocessing of the two DTI acquisitions, 2- Generation of the two tensor populations to be compared and 3- Multivariate statistical testing between the two sets of tensors.

2 Proposed Framework

2.1 Preprocessing

Eddy current distortion correction. Each set of DW-images may be misaligned due to eddy current distortions. Thus, a correction is achieved by affinely registering each slice of the DWI data set onto the corresponding slice of the T2 image (without diffusion weighting).

Registration. Change detection methods generally require the two images to be accurately registered. Thus, an affine transformation is estimated between the two DWI sets, considering the mutual information between the T2 images. Since an affine transformation may not always be sufficient to compensate for all the undesired global differences that may be present between the two acquisitions, the registration is then refined with the deformable method described in [10].

Interpolation and reorientation. An interpolation method is required to re-sample the warped tensor image. It has been shown in [11] that interpolation methods for tensor image can lead to significantly different results according to the chosen metric. In particular, the Log-Euclidean metric seems to be preferred to the Euclidean metric because of the swelling effect induced by the latter. Therefore, we use a third order B-Spline interpolation method in the Log-Euclidean space. Then, the Preservation of Principal Direction (PPD) reorientation strategy described in [12] is applied to preserve the orientation and shape information of the warped tensor image.

2.2 Generation of the Tensor Sets Reflecting the Local Variability

In this section, we investigate several strategies to generate a set of tensors characterizing the variability of each voxel, based on the variability existing in the DW-images, or in the spatial neighborhood of the considered voxel, or a combination of these two. Tensors are estimated using the standard least squares approach. Since this procedure does not guarantee the positive definiteness of the tensors (which is a required property for using Log-Euclidean metrics), negative eigenvalues are set to an arbitrary small positive value.

Local Bootstrap. DTI acquisitions are composed of P gradient directions, P being usually greater than six, thus leading to an overdetermined linear system to estimate the tensor. Using a bootstrap strategy enables us to learn the distribution reflecting tensor estimation uncertainty based on the DW-images [13]. To this end, we generate at each voxel N bootstrap samples by randomly selecting P DW-signal with replacement, and preserving the gradient directions [13] (Fig.1a). The random draw with replacement of the gradient directions amounts actually to associate different weights to each gradient direction for tensor estimation. Using the local bootstrap on the two registered data sets enables to generate two populations of tensors.

Spatial Neighborhood. An alternative idea is to learn the tensor distribution at each voxel by considering all the tensors in a surrounding user-defined spatial neighborhood. By this way, we make the implicit assumption that tensors follow a constant model on this neighborhood and that the observed variability is the consequence of random effects. Learning tensor distribution at each voxel by considering all the tensors in a surrounding spatial neighborhood is based on the commonly made assumption of a constant piecewise model [1], [2]. The limitations of this model are well known, in particular here, at the interface between different tissues, but the model is known to be robust and efficient. This method has the advantage to be computationally cheap. However, it does not take into account the information carried out by the DW-images.

Spatial Bootstrap. The spatial bootstrap is a combination of the local bootstrap and the spatial neighborhood methods. From a surrounding box with a user-defined size, we generate at each voxel N bootstrap samples by drawing for each gradient direction one DW-signal of the neighborhood (Fig.1b).

2.3 Multivariate Statistical Testing on Tensors

Log-Euclidean metric. Using Euclidean metrics, such as the Frobenius distance, may raise some problem when considering symmetric positive definite




Fig. 1. Schematic representation of the local and spatial ($3 \times 3 \times 3$ neighborhood) bootstrap for P gradient directions. 1₂₇ represents the DW-signal of the 27th voxel of the neighborhood according to the first gradient direction.

(SPD) matrices since it does not define a vector space on SPD matrices. A simple and computationally efficient way to build a vector space on SPD matrices is to consider the Log-Euclidean metric [11]. The Log-Euclidean metric corresponds to the Euclidean metric on the logarithm of matrices. By this way, the matrices with null or negative eigenvalues are at an infinite distance to any SPD matrices. The logarithm of an SPD matrix D is obtained as $L = \text{Log}(D) = U \text{Log}(\Lambda) U^T$, where Λ and U are the matrices derived from the standard spectral decomposition, containing respectively the eigenvalues and the eigenvectors of D . Then, the Log-Euclidean distance between two SPD matrices D_1 and D_2 can be defined as the Euclidean distance between their logarithms:

$$d^2(D_1, D_2) = \|\text{Log}(D_1) - \text{Log}(D_2)\|^2 \quad (1)$$

According to this metric, an estimator of the mean \bar{D} of a set of N tensors D_i is given by the exponential of the arithmetic log-tensors mean, i.e. $\bar{D} = \exp(\bar{L})$, with:

$$\bar{L} = \underset{\Sigma}{\text{argmin}} \sum_{i=1}^N \|\text{Log}(D_i) - \Sigma\|^2 = \frac{1}{N} \sum_{i=1}^N \text{Log}(D_i) \quad (2)$$

Statistical tests. Many statistical tests rely on the normal distribution. Considering the multivariate normal distribution for the SPD matrices has the drawback to associate matrices with negative or null eigenvalues with a non null probability. To circumvent this limitation, Schwartzman [9] suggests to model the matrix logarithms with the multivariate normal distribution, which comes to model the SPD matrices with a Log-normal distribution. Based on this model, it is possible to derive statistical tests on tensors eigenvalues and eigenvectors. We consider two populations of N_1 and N_2 tensors respectively. Under the assumption that the tensor logarithms of the two populations follow the normal distributions $L_1 \sim \mathcal{N}(M_1, \sigma^2 Id)$ and $L_2 \sim \mathcal{N}(M_2, \sigma^2 Id)$, the Maximum Likelihood Estimates of M_1 , M_2 , and σ^2 are respectively \bar{L}_1 , and \bar{L}_2 computed according to Eq. 2, and:

$$\hat{\sigma}^2 = \frac{1}{6(N_1 + N_2 - 2)} \left[\sum_{i=1}^{N_1} \text{tr} (L_{1i} - \bar{L}_1)^2 + \sum_{i=1}^{N_2} \text{tr} (L_{2i} - \bar{L}_2)^2 \right] \quad (3)$$

We consider a test [9] that evaluates whether the two populations of diffusion tensors have similar eigenvalues, but possibly different eigenvectors. Let D_1 , U_1 and D_2 , U_2 be the matrices derived from standard spectral decomposition, and containing respectively the eigenvalues and eigenvectors of M_1 and M_2 . The test, based on the log-likelihood ratio under hypotheses $H_0 : D_1 = D_2$ vs $H_1 : D_1 \neq D_2$ is:

$$T = \frac{(N_1 + N_2 - 2)N_1N_2}{3(N_1 + N_2)^2\hat{\sigma}^2} \text{tr} [(A_1 - A_2)^2] \quad (4)$$

with A_1 and A_2 the eigenvalue matrices of \bar{L}_1 and \bar{L}_2 , respectively.

Another case of interest is to test whether the two population of diffusion tensors have similar eigenvectors, i.e the same orientation. When we combine the two populations, we derive the mean tensor \bar{L} of the new population from Eq. 2. We compute then the eigen-decomposition of this latter, and note Λ the eigenvalue matrix. The proposed test, based on the log-likelihood ratio under hypotheses $H0 : U_1 = U_2$ vs $H1 : U_1 \neq U_2$, treating $D_1 = D_2 = D$ as a nuisance parameter is:

$$T = \frac{2(N_1 + N_2) - 3}{6\hat{\sigma}^2} \text{tr} \left[\left(\frac{N_1\Lambda_1 - N_2\Lambda_2}{N_1 + N_2} \right)^2 - \Lambda^2 \right] \quad (5)$$

3 Results

Experiments on synthetic lesions. We considered two successive DTI acquisitions of the same subject acquired on a 3.0T MRI scanner with 30 encoding gradients (b-value of 1000 s/mm^2). By this way, the differences between the two scans are only due to the acquisition noise and distortion. The image dimensions are $128 \times 128 \times 40$ and the spatial resolution is $1.8 \times 1.8 \times 3.5 \text{ mm}^3$. A synthetic lesion is simulated in one of these scans as follows. We consider a lesion mask, located in the white matter, and delineated by a physician to ensure its location and shape to be clinically relevant. Inside the lesion mask, we modify uniformly either the diffusivity in the principal direction, i.e. the principal eigenvalue (application of a multiplicative factor $k \in]1, 2[$), or the tensor orientations (rotation of the ellipsoid canonical xyz system by $\alpha \in]0, \pi/8[$). The criterion used to compare the different methods is the area under Receiver Operating Characteristic (ROC) curves. In a ROC curve, the true positive rate (or sensitivity) is plotted in function of the false positive rate (or 1-specificity) for different cut-off points. A test that allows a perfect discrimination is characterized by a ROC plot that passes through the upper left corner (i.e. an area of one). The Table 1 summarizes the results for the three investigated strategies: S (Spatial), LB (Local Bootstrap), and SB (Spatial Bootstrap). For all the experiments, we consider a $3 \times 3 \times 3$ spatial neighborhood for the SB and S strategies, and $N = 100$ bootstrap samples for LB and SB strategies. For both diffusivity and orientation modifications, the two methods based on the spatial neighborhood stand out, thus pointing out the prominent part of the spatial information for reliable change detection.

Table 1. Areas under the ROC curve for eigenvalues and eigenvectors tests (MS simulation: see text)

	k	S	LB	SB		α	S	LB	SB
eigenvalues:	1.2	0.713	0.202	0.724	eigenvectors:	$\pi/32$	0.814	0.590	0.726
	1.4	0.902	0.417	0.927		$2 * \pi/32$	0.871	0.756	0.844
	1.6	0.960	0.743	0.982		$3 * \pi/32$	0.870	0.822	0.850
	1.8	0.976	0.787	0.991		$4 * \pi/32$	0.871	0.877	0.873

Experiments on an MS patient. We considered two brain DTI acquisitions of an MS patient acquired on a 3.0T MRI scanner with 33 encoding gradients (b-value of 1000 s/mm^2) at two different times ($t_2 - t_1 = 15 \text{ months}$). The images dimensions are $256 \times 256 \times 34$ and the spatial resolution is $1 \times 1 \times 3 \text{ mm}^3$. Examples of detection maps obtained with the three proposed approaches are presented in Fig.2. Each potential approach gives different results for the variance (Eq.3), and mean tensors estimation (Eq.2). For instance (Fig.2: var-LB), the DW-signal variability is greater in areas of high diffusivity (e.g. cerebrospinal fluid), and much lower in highly structured areas (e.g. white matter). On average, considering the spatial neighborhood gives larger values for the variance except in the cerebrospinal fluid (Fig.2: var-S). As a consequence, the bootstrap yields less false detections in the ventricles. The local bootstrap generates change detection maps sensitive to small variations since the variance values at each voxel are smaller (Fig.2: eig-LB). Directly using the spatial information after the tensors estimation, achieves satisfactory results for a low computational burden. A spatial bootstrap estimation benefits from the interesting properties of the local bootstrap (i.e., direct modeling of the DW-signal), and reduces false detections thanks to the spatial neighborhood information (Fig.2: SB).

A visual inspection of the results by an expert helps us to verify that anatomical changes are also observable in diffusion imaging. Then other changes are observed, that may provide information of a different kind. In Fig.2: eig, two




Fig. 2. Variance (var), and detection maps obtained for the eigenvalues (eig) and eigenvectors (vec) test statistics with the three different strategies for tensor population generation

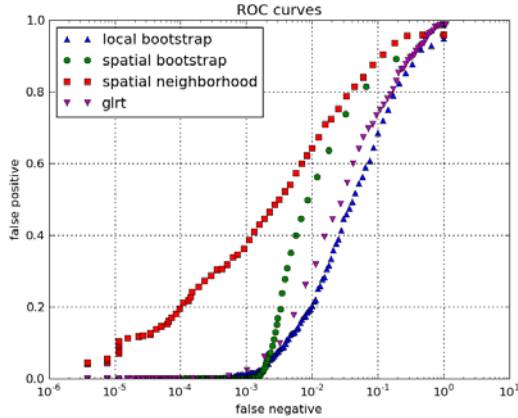


Fig. 3. ROC curves (semilogx plot): longitudinal change detection results for the three proposed generation of tensor sets and the GLRT

prominent lesions are detected. Both detections correspond to existing lesions (in anatomical images), but according to the expert, only the bottom one is changing. The top lesion is of widespread interest, since considering the anatomical images, it does not seem to evolve during the time elapsed between the two acquisitions. Then in Fig.2: vec, if it was obvious that for MS patients we have to detect variations in the eigenvalues map (demyelination Fig.2: eig), we also highlight the fact that the orientation is significantly changing. Notice also that the proposed framework detects white matter modification near the ventricles corresponding to physiological process known as leukoaraiosis.

ROC curves are used to quantitatively compare the three strategies, Fig.3. We compare the results with a reference detection map obtained by a manual segmentation of evolving lesions on the T2 images by an expert. Thus, ROC plots should be analyzed carefully since the segmentation from the expert cannot really be considered as a ground truth (changes may occur in diffusion imaging without being visible in anatomical images). The proposed framework is compared with the GLRT framework introduced in [8], which differs in the following points. During the preprocessing, we directly resample the DTI rather than going through the DW-signal. On top of that, the GLRT does not take into account the positive definite nature of matrices. From Fig.3, we highlight the interesting behavior of the proposed test statistic compared to the GLRT. The separability property (eigenvalues/eigenvectors) is of much interest.

4 Conclusion

Clinical disability for MS patients are localized in the white matter, and are driven by a significant modification of the perpendicular diffusivity. As expected, the eigenvalues test statistic is relevant to detect such modifications. The com-

parison of three strategies for the generation of the two tensor sets to be compared has demonstrated the superiority of methods based on a spatial neighborhood. Nonetheless, the spatial information introduces a smoothing effect on the detection, while the bootstrap based approaches enable a better spatial localization. We also notice significant orientation modifications during longitudinal analysis. In the future we want to investigate if some correlation exists between MS disease evolution and these pathological modifications, and try to combine both information as a predictive tool for MS.

References

1. Radke, R.J., Andra, S., Al-Kofahi, O., Roysam, B.: Image change detection algorithms: a systematic survey. *IEEE Trans. Image Process* 14(3), 294–307 (2005)
2. Bosc, M., Heitz, F., Arispach, J.-P., Namer, I., Gounot, D., Rumbach, L.: Automatic change detection in multimodal serial MRI: application to multiple sclerosis lesion evolution. *NeuroImage* 20, 643–656 (2003)
3. Rey, D., Subsol, G., Delingette, H., Ayache, N.: Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis. In: Kuba, A., Sámal, M., Todd-Pokropek, A. (eds.) *IPMI 1999*. LNCS, vol. 1613, pp. 154–157. Springer, Heidelberg (1999)
4. Thirion, J.-P., Calmon, G.: Deformation analysis to detect and quantify active lesions in three-dimensional medical image sequences. *IEEE Trans. Med. Imaging* 18(5), 429–441 (1999)
5. Boisgontier, H., Noblet, V., Heitz, F., Rumbach, L., Arispach, J.-P.: An automatic method for change detection in serial DTI-derived scalar images. In: *Workshop MIAMS - MICCAI (2008)*
6. Chung, S., Pelletier, D., Sdika, M., Lu, Y., Berman, J.I., Henry, R.G.: Whole brain voxel-wise analysis of single-subject serial DTI by permutation testing. *Neuroimage* 39(4), 1693–1705 (2008)
7. Boisgontier, H., Noblet, V., Heitz, F., Rumbach, L., Arispach, J.-P.: Statistical detection of longitudinal changes between apparent diffusion coefficient images. Application to multiple sclerosis. In: Yang, G.-Z., Hawkes, D., Rueckert, D., Noble, A., Taylor, C. (eds.) *MICCAI 2009*. LNCS, vol. 5761, pp. 959–966. Springer, Heidelberg (2009)
8. Boisgontier, H., Noblet, V., Heitz, F., Rumbach, L., Arispach, J.-P.: Generalized likelihood ratio tests for change detection in diffusion tensor images. In: *ISBI*, pp. 811–814 (2009)
9. Schwartzman, A.: Random ellipsoids and false discovery rates: statistics for diffusion tensor imaging data. PhD thesis, Stanford University (June 2006)
10. Noblet, V., Heinrich, C., Heitz, F., Arispach, J.-P.: Retrospective evaluation of a topology preserving non-rigid registration method. *Medical Image Analysis* 10(3), 366–384 (2006)
11. Arsigny, V., Fillard, P., Pennec, X., Ayache, N.: Geometric means in a novel vector space structure on symmetric positive-definite matrices. *SIAM Journal on Matrix Analysis and Applications* 29(1), 328–347 (2007)
12. Alexander, D.C., Pierpaoli, C., Basser, P.J., Gee, J.C.: Spatial transformations of diffusion tensor magnetic resonance images. *IEEE Trans. Med. Imaging* 20(11), 1131–1139 (2001)
13. Pajevic, S., Basser, P.J.: Parametric and non-parametric statistical analysis of DT-MRI data. *Journal of Magnetic Resonance* 161(1), 1–14 (2002)