A spatial variation model of white matter microstructure

Gemma L. Morgan[1], Hui Zhang[1], Brandon Whitcher[2], and Daniel C. Alexander[1]

[1] Centre for Medical Image Computing, University College London, London, UK
g.morgan@cs.ucl.ac.uk
[2] Clinical Imaging Centre, GlaxoSmithKline, London, UK

Abstract. In this study, we introduce a new technique to model the variation of microstructural parameters across specific brain regions. We use a simple model of diffusion in each voxel, but model the variation of parameters across the region using penalised splines. We fit the whole region model directly to the diffusion-weighted signals. We test the technique on the mid-sagittal section of the corpus callosum (CC) using a diffusion MRI data set with distinct age groups. The method detects differences and separates the groups.

1 Introduction

Diffusion MRI is a powerful, non-invasive imaging tool which measures the displacement of water molecules in vivo. Because the paths of water molecules are heavily influenced by the shape and structure of the environment in which they move, diffusion MRI is a sensitive probe for measuring tissue microstructure. It has been used extensively to study white matter in the brain. From a clinical perspective, this is an invaluable technique as we can use diffusion MRI to infer microstructural tissue changes due to pathology, potentially improving our understanding and treatment of neurological diseases. However, to do this we need a biomarker that is both sensitive and specific to underlying microstructure changes. The most commonly used white matter biomarkers are Fractional Anisotropy (FA) and Mean Diffusivity (MD), but they are difficult to relate to specific tissue features. More complex models exist with parameters that correspond to specific microstructure changes [4–6] but as model complexity increases accurate parameter estimation is more susceptible to noise.

In this study, we present a new technique which uses a spatial model within an ROI or tract to describe the variation of tissue diffusion properties. The method fits a spatial model directly to the diffusion-weighted signals, which should reduce the effect of noise on parameter estimation and better capture the underlying variation of parameters. The method combines with diffusion models of varying complexity, allowing a variety of biomarkers to be studied. Because the model is fit within a shaped-based normalisation framework, this technique could be used in group comparison studies. The spatial model can be compared between populations, providing localised information about tissue changes.

The paper is organised as follows. In Section 2 we discuss diffusion models and review previously proposed group study techniques. In Section 3 we introduce our Spatial Variation framework and in Section 4 we apply it to diffusion MRI data with distinct age groups. We discuss our results and conclude in Section 5.

2 Background

In this section we describe some of the most commonly used models of water diffusion in tissue and provide an overview of current methods used to perform group studies. We also introduce the concept of continuous medial representation which we make use of in our Spatial Variation model in Section 3.

2.1 Diffusion Models

To date, most diffusion MRI group studies use the diffusion tensor (DT) [1] to relate diffusion-weighted signals to the diffusion properties of tissue. The DT model assumes that the displacement of water particles is Gaussian and fits a tensor to the diffusion-weighted signals that represents the amount of diffusion along different directions. Scalar indices such as FA and MD can be derived from the elements of the DT [2] and used as biomarkers in group studies. However, water diffusion in the brain is not Gaussian (particularly in white matter where water is restricted by myelinated axons) and FA and MD cannot be directly related to the underlying tissue microstructure. Recently, Panagiotaki et al [3] showed that simple two-compartment models of diffusion, such as a ball and stick [4] or tensor and stick, provide a better fit to diffusion MR data than the DT model. These models mimic the structure of white matter more closely than the DT and the model parameters are potentially more physiologically relevant.

2.2 Group Studies

Group studies can be broadly divided into three categories - region of interest (ROI), voxel or tract based. In ROI studies, the anatomical feature is segmented in each subject and the parameter of interest (usually FA or MD) is averaged over the ROI to provide one value for each subject. However, the effect of averaging could mask small but potentially significant differences in spatial variation between subjects. Unless the ROIs are small, it is not possible to localise the tissue changes. Voxel-based approaches such as Voxel Based Morphometry (VBM) [7] take a very different approach and compare all voxels in the brain between subjects. To do this effectively, the data must be very carefully registered as misalignment and excessive data smoothing can both introduce false positives. Tract-based spatial statistics (TBSS) [8] has been proposed to circumvent some of these problems by projecting each data set onto a template skeleton and performing statistics only on the skeleton. Whilst it does avoid the need for smoothing and perfect alignment, reducing the data down to a skeleton could reduce the statistical significance of the results. Finally, tract based methods use

fibre tractography algorithms to segment and study individual tracts of interest. Tracts are analysed either by averaging parameter values along the whole tract or, more recently, sampling parameter values along the arc length of the tract to create a function [9, 10]. These functions can be analysed either by using pointwise [9] or functional analysis [10]. Like TBSS, these techniques consider only peak parameter values from the centre of the tract, again discarding potentially useful information. These techniques work well for tract with tubular geometries; however it is also possible to derive skeletons for tracts with sheet-like structure using techniques such as tract-specific analysis (TSA) [11, 14].

2.3 Continuous medial representation

Continuous medial representation (cm-rep) is a shape analysis technique which has been shown to be suitable for analysing white matter tracts using medial axes for two-dimensional structures such as the mid-sagittal cross-section of the CC [13] and medial surfaces for three-dimensional sheet-like tracts such as the corticospinal tract [11, 14]. It allows for comparison between populations using shape-based normalisation which aligns objects based on global shape and is particularly appropriate for white matter tracts which have homogeneous interiors.

3 Methods

In this Section we introduce the Spatial Variation framework and describe a simple implementation. We discuss some statistical techniques which can be used to assess differences in spatial models between groups.

3.1 Spatial Variation framework

The Spatial Variation framework estimates microstructure parameters across a whole tract using prior spatial information to constrain the fitting. Rather than fitting in each voxel individually, it uses a forward model that predicts the diffusion MR signals within the whole tract using a set of spatial functions, the spatial model, that control the regional variation of microstructure parameters across the tract. The key components of the method are the spatial model, which predicts the diffusion model parameters in every voxel, and the local diffusion model, which calculates the diffusion MR signals in every voxel from the diffusion model parameters. The optimal spatial model, which should be fully described by a small set of parameters, is found by minimising an objective function based on differences between predicted and measured diffusion MR signals. This process is illustrated in Figure 1, which gives an overview of the pipeline.

The framework itself is independent of the choice of spatial model, diffusion model and objective function. In the next sections, we describe one possible implementation using simple spatial and diffusion models, and show how the cm-rep of tracts can be used to bridge the gap between global and local parameters.

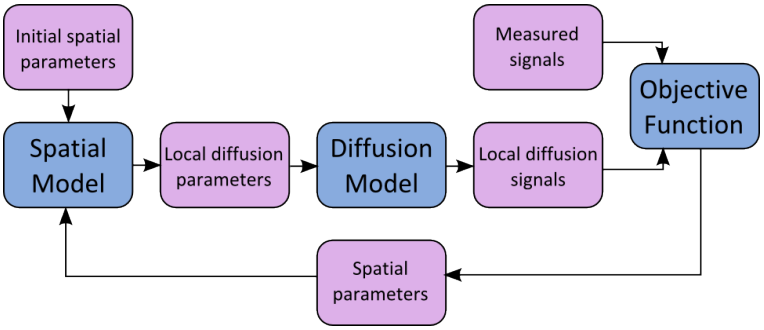

Fig. 1: The shape model comprises a set of curves describing the variation of diffusion model parameters across the tract medial axis. These curves predict local diffusion parameters in every voxel of the tract. The diffusion model predicts the MR signals, given the predicted diffusion parameters. The shape model is controlled by a small set of parameters, and its goodness of fit can be calculated using an objective function based on predicted and measured MR signals. The spatial model is optimised by iteratively minimising the objective function.

Spatial Model We choose to model the variation of the diffusion model parameters using B-splines, a natural choice of basis functions since they effectively capture local features of non-periodic data. Specifically, we follow the penalised B-spline (P-spline) approach of Eilers and Marx [16] which uses a large number N of equally spaced knots to specify the m^{th} order basis functions B. The basis functions are penalised during optimisation to reduce excessive local variation. Each diffusion model parameter $\boldsymbol{p}$=(p_1, p_2, ... p_k) is written as a function of position along the medial axis s using these basis functions as

$$p_k(s, \mathbf{a}) = \sum_{n=1}^{N-m} a_{k,n} B_n(s), \tag{1}$$

where $\mathbf{a}$ are the relative weights of the basis functions. We optimise $\mathbf{a}$ so the set of functions $\boldsymbol{p}$ best reflect the true variation of the diffusion parameters across the tract. This requires us to compare the observed MR signals, measured in voxels with discrete spatial positions, with MR signals predicted from our continuous model. The cm-rep provides a convenient way to convert our predictions from continuous to discrete as it outputs a list of points lying on the medial axis. We use a simple nearest neighbour approach to assign a distance along the medial axis to all voxels of interest. Once we know the positions of all voxels along the medial axis, we can predict the diffusion parameters, and thus the MR signals, at these locations only.

Diffusion Model A modified ball and stick model is used to describe water diffusion in tissue. The standard ball and stick model [4] is a simple two-compartment model that models the total diffusion MR signal as a mixture of

signals due to the restricted intra-axonal water (the stick) with volume fraction f and orientation $\mathbf{e}$ defined by angles θ and ϕ and the hindered extra-axonal water (the ball). Both compartments have diffusivity d. Our modification involves a third compartment to account for cerebrospinal fluid (CSF) contamination with volume fraction g and diffusivity d_{CSF}.

The spatial model outlined in Section 3.1 assumes that the white matter microstructure parameters $\boldsymbol{p}$=(f, d, θ, ϕ), may be written as smoothly varying functions parameterised by coefficients $\mathbf{a}$=($\mathbf{a}_f$, $\mathbf{a}_d$, $\mathbf{a}_\theta$, $\mathbf{a}_\phi$) as shown in Equation 1; however we cannot assume the same thing of the partial volume parameters, g and d_{CSF}. Prior to fitting the spatial model, we estimate g in all voxels of interest from the b=0 images. We write the total signal S in the b=0 image as a mixture of white matter signal S_{WM} and CSF signal S_{CSF}

$$S = gS_{CSF} + (1-g)S_{WM}. \tag{2}$$

We estimate average values for S_{CSF} and S_{WM} from voxels that are manually segmented as 'pure' CSF or white matter voxels in the b=0 image and solve for g. We set d_{CSF} to 3.0×10^{-9} $\mathrm{m^2s^{-1}}$, a typical value for CSF.

The model for the i^{th} diffusion MR signal in the j^{th} voxel may be written as

$$\tilde{A}_{ij}(\boldsymbol{p}_j(\mathbf{a}); g_j, b, \hat{\mathbf{G}}_i) = g_j \mathrm{e}^{-bd_{CSF}} + (1-g_j)\Big(f_j \mathrm{e}^{-bd_j(\mathbf{e}(\theta_j,\phi_j)\cdot\hat{\mathbf{G}}_i)^2)} + (1-f_j)\mathrm{e}^{-bd_j)}\Big) \tag{3}$$

where b is the diffusion-weighting factor and $\hat{\mathbf{G}}_i$ is the gradient direction.

Optimisation We estimate the spatial model coefficients $\mathbf{a}$ that best describe the variation of the model parameters $\boldsymbol{p}$ across the whole region. We do this by minimising the sum of squared differences between $\tilde{A}_{ij}$ predicted from Equation 3 and the observed signals A_{ij} in all J voxels, subject to regularisation. Our objective function is

$$\arg\min\Big(\sum_{i=1}^{I}\sum_{j=1}^{J}(A_{ij} - \tilde{A}_{ij}(\boldsymbol{p}_j(\mathbf{a}); g_j, b, \mathbf{G}_i))^2 + \sum_{k\in(f,d,\theta,\phi)} \lambda_k \sum_{n=1}^{N-m} (\Delta^2 a_{k,n})^2\Big). \tag{4}$$

For the regularisation term, we use a second order difference penalty on $\mathbf{a}_k$, which enforces the smoothness of our solutions [16]. We minimise the objective function iteratively using a Levenberg-Marquardt algorithm.

3.2 Statistical analysis

After fitting the spatial model for each subject, we can test for group differences using functional data analysis [15]. The mean curve for each parameter $\bar{p}_k(s)$ for each group can be written as

$$\bar{p}_k(s) = \sum_{q=1}^{Q} p_{k,q}(s) \tag{5}$$

where $p_{k,i}(s)$ are the individual curves for the Q subjects. We test for group diffences in these mean parameter curves at l points sampled along the medial axis using Student's t test. To correct for multiple comparisons we use a Westfall-Young randomisation method proposed by Cox and Lee [17], specifically designed for functional data. Unlike traditional multiple comparison corrections, which lose significance as l increases, p-values obtained using Westfall-Young randomisation converge to the continuum limit as $l \to \infty$.

4 Experiments and Results

We demonstrate this method on the mid-sagittal section of the CC. We choose this region as the relatively constant fibre orientation reduces the number of parameters to estimate. Histology studies [18] have also shown that underlying microstructural indices such as axon radius and density vary smoothly over the anterior-posterior direction of the CC, which we hypothesise will manifest as smooth variation in parameters such as f and d.

4.1 Data acquisition and pre-processing

In this study, our diffusion MRI data sets are drawn from the large IXI database (freely available at: www.brain-development.org). We use data from 30 subjects who divide into two distinct groups: 20-29 years old (9 female, 6 male) and 60-69 years old (9 female, 6 male). All data were acquired on a 3T scanner using 15 gradient directions (b=1000 s mm^{-2}) and 1 b=0 s mm^{-2} measurement.

The key pre-processing steps required are to identify the mid-sagittal slice of the CC and to extract the medial axis. First the brain is extracted using FSL's BET tool [19] and FA maps are calculated. From these FA maps we identify the mid-sagittal slice [20]. Thresholding (FA >0.35) and connected component analysis isolate the CC on this slice. ITK-SNAP [21] is used to make minor manual adjustments to the segmentation, for example in cases where the fornix is misclassified as part of the CC. The medial axis of the CC can then be extracted using cm-rep [13].

4.2 Spline fitting

Due to the coherent orientation of fibres within the CC, we fit constants for θ and ϕ across the whole ROI. The splines modelling the variation of f and d across the CC are fit by dividing the normalized medial axis length into 20 intervals described by 21 cubic B-splines. In total, we fit 44 parameters for each data set. Using Generalised Cross Validation (GCV) across all data sets, we set the regularisation parameters for f and d to be λ_d=λ_f=10.

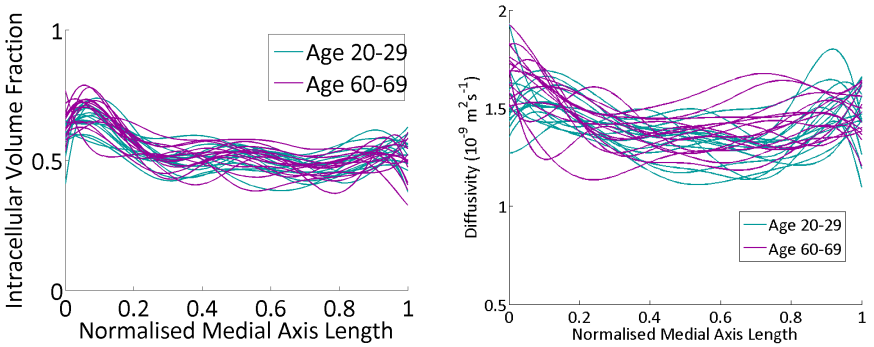

(a) Volume fraction curves (b) Diffusivity curves

Fig. 2: The fitted curves for (a) volume fraction and (b) diffusivity are shown for all subjects. Volume fraction is consistent amongst all subjects, but we see bigger group differences in diffusivity.

4.3 Results

Figure 2 shows the volume fraction and diffusivity curves for all subjects. The volume fraction curves have consistent shapes for all subjects, whereas the diffusivity curves show greater differences between the groups.

Figure 3 shows the mean volume fraction and diffusivity curves calculated using Equation 5. Again, this highlights the similarities in volume fraction between the groups and the increases in diffusivity in the anterior genu and anterior splenium in the older group.

Using the statistical techniques discussed in Section 3.2, we calculate p-values for the mean group differences in both volume fraction and diffusivity at 100 positions along the medial axis length (Figure 4). After correcting for multiple comparisons, we see that at the tip of the genu (s=0–0.4), there is significantly higher diffusivity in the older group than the younger, as well as a trend towards decreased volume fraction (s=0.02). Additionally, there is a trend towards higher diffusivity in the older group in the anterior splenium (s=0.71–0.76), although this is not accompanied by a corresponding change in volume fraction.

5 Discussion

We have presented a new technique that models the variation of microstructure parameters across white matter tracts. We have shown that it can identify and localise group differences due to age in the mid-sagittal section of the CC. Specifically, we have found a significant increase in diffusivity in the anterior genu with age combined with a trend towards reduced volume fraction. We also observe a trend towards increased diffusivity with age in the anterior splenium. Although it is commonly known that ventricle size increases with age [22], we

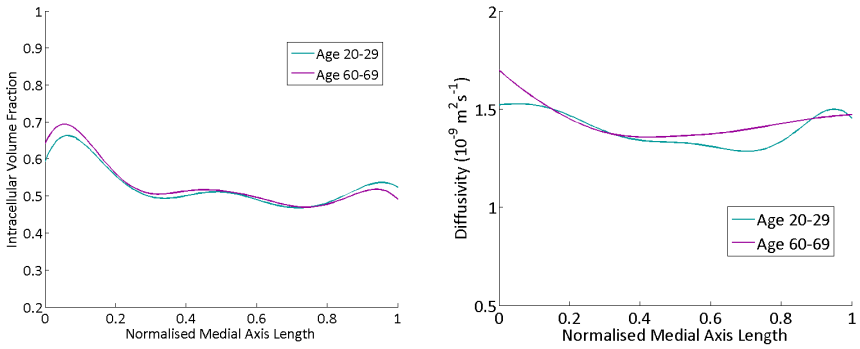

(a) mean volume fraction curve (b) mean diffusivity curve

Fig. 3: The mean curves for (a) volume fraction and (b) diffusivity are shown for all subjects.

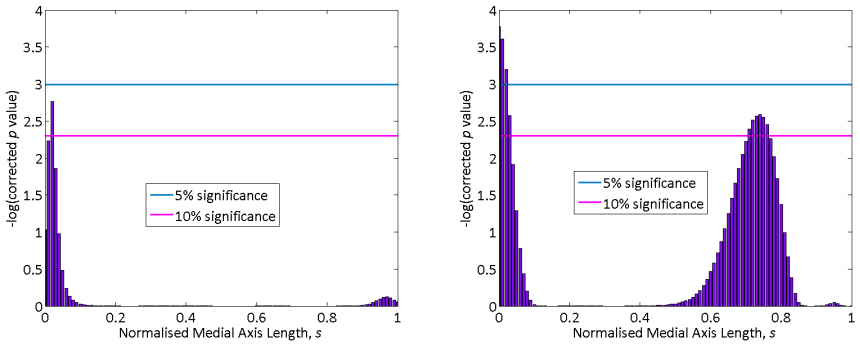

(a) (b)

Fig. 4: Plots of p-values as a function the medial axis position showing group differences in (a) volume fraction and (b) diffusivity. There is a significant increase in diffusivity and trend towards decreased volume fraction in the anterior genu in the older group. We also see a trend towards increased diffusivity with age in the anterior splenium.

think it is unlikely that the change in diffusivity in the genu is due to this effect as our diffusion model explicitly accounts for partial volumes. Previous studies [23] have shown significant changes due to age in the genu, which supports our results. Even though we have been able to localise tissue differences due to age, we cannot specify what causes these changes, e.g. larger axon diameters or lower axon densities, due to the simplicity of our diffusion model.

Future work will use this method to measure small but significant changes in tissue microstructure due to pathology that other standard techniques might miss. This is because data is pooled across the whole ROI or tract to reduce the effect of noise on parameter estimation. However several improvements need to be made. First, we will extend the method to three dimensions so that we can model microstructure variation across whole tracts rather than single slices. We will also exploit the analytical form of the spline curves further to see if the spline derivatives can be used to identify microstructural group differences. Finally, we will replace our simple model of diffusion with more complex models in order to investigate more physiologically relevant microstructure parameters such as axon radius and density. When using these complex models with parameters that are difficult to estimate, the real power of this method to pool data and reduce the effect of noise will be apparent.

6 Acknowledgements

We would like to thank Dr Hui Sun and Prof Paul Yushkevich, University of Pennsylvania, for invaluable assistance with the continous-medial representation code. This work is funded by GlaxoSmithKline and EPSRC.

References

1. Basser, P.J., Matiello, J., Le Bihan,D.: MR diffusion tensor spectroscopy and imaging. Biophysical Journal , 259–267 (1994)
2. Basser, P.J., and Pierpaoli, C.: Microstructural and Physiological Features of Tissues Elucidated by Quantitative-Diffusion-Tensor MRI. Journal of Magnetic Resonance, Series B, 111, 209–219 (1996)
3. Panagiotaki, E., Fonteijn, H., Siow, B., Hall, M.G., Price, A., Lythgoe, M.F., Alexander, D.C.: Two-compartment models of the diusion MR signal in brain white matter. Medical Image Computing and Computer Assisted Intervention (MICCAI) London:Springer, 1, 329–336 (2009)
4. Behrens, T.E.J., et al. Characterization and Propagation of Uncertainty in Diffusion Weighted MR Imaging. Magnetic Resonance in Medicine, 50, 1077–1088 (2003)
5. Assaf, Y. et al.: New Modeling and Experimental Framework to Characterize Hindered and Restricted Water Diffusion in Brain White Matter. Magnetic Resonance in Medicine, 52, 965–978 (2004)
6. Alexander, D.C.: A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. Magnetic Resonance in Medicine, 60, 439–448 (2008)

7. Ashburner, J., and Friston, K.J.: Voxel-Based Morphometry - The Methods. NeuroImage, 11, 805–821 (2000)
8. Smith, S., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., and Behrens, T.E.J.: Tract-based spatial statistics: Voxelwise analysis of multi-subject diffusion data. NeuroImage, 31, 1487–1505 (2006)
9. Maddah, M., et al Findings in Schizophrenia by Tract-Oriented DT-MRI Analysis. Medical Image Computing and Computer Assisted Intervention (MICCAI), Springer, 1, 917–924 (2008)
10. Goodlett, C.B., Fletcher, P.T., Gilmore, J.H., and Gerig, G.: Group analysis of DTI fiber tract statistics with application to neurodevelopment. NeuroImage, 45, S133–S142 (2009)
11. Yushkevich, P.A., Zhang, H., Simon, T.J, and Gee, J.C.: Structure-specific statistical mapping of white matter tracts. NeuroImage, 41, 448–461 (2008)
12. Yushkevich, P.A., Zhang, H., Gee, J.C.: Continuous medial representation for anatomical objects. IEEE Trans Med Imaging, 25, 1547–1564 (2006)
13. Sun, H., Yushkevich, P.A., Zhang, H., Cook, P.A., Duda, J.T., Simon, T.J., Gee, J.C.: Shape-Based Normalization of the Corpus Callosum for DTI Connectivity Analysis. IEEE Transactions on Medical Imaging, 26, 1166–1178 (2007)
14. Zhang, H., Awate, S.P., Das, S.R., Woo, J.H., Melhem, E.R., Gee, J.C. and Yushkevich, P.A.: A tract-specific framework for white matter morphometry combining macroscopic and microscopic tract features. Medial Image Analysis, In press (2010)
15. Ramsay, J.O. and Silverman, B.W.: Functional Data Analyis. Springer (2005)
16. Eilers, P.H.C. and Marx, B.D.: Flexible smoothing with B-splines and penalties. Statistical Science, 11, 89–121 (1996)
17. Cox, D.D. and Lee, J.S.: Pointwise testing with functional data using the Westfall-Young randomization method. Biometrika, 95, 621–634 (2008)
18. Aboitiz, F., Scheibel, A.B., Fisher, R.S. and Zaidel, E.: Fiber composition of the human corpus callosum. Brain Research, 598, 143–153 (1992)
19. Smith, S.M.: Fast robust automated brain extraction. Human Brain Mapping, 17, 143–155 (2002)
20. Ardekani, B.A., Kershaw, J., Braun, M. and Kanno, I.: Automatic detection of the mid-sagittal plane in 3-d brain images. IEEE Trans. Med. Imaging, 16, 947–952 (1997)
21. Yushkevich, P.A., Piven, J., Hazlett, H.C., Smith, R.G., Ho, S., Gee, J.C., and Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. NeuroImage, 31, 1116–1128 (2006)
22. Haug, G.: Age and sex dependence of the size of normal ventricles on computed tomography. Neuroradiology, 14, 201–204 (1977)
23. Otaa, M., Obataa, T., Akinea, Y., Itoa, H., Ikehirac, H., Asadad, T., and Suharaa, T.: Age-related degeneration of corpus callosum measured with diffusion tensor imaging. NeuroImage, 31, 1445–1452 (2006)