

A machine learning approach to Diffusion MRI partial volume estimation

No Author Given

No Institute Given

Abstract. Tissue-type partial volume modelling is generally an ill-posed problem in single-shell dMRI. On the other hand, T1w images, typically acquired along with the diffusion data, allow for an accurate estimation of the tissue partial volume fractions (PVFs).

We propose in this paper to use a data driven approach to predict the T1w-derived PVFs from the diffusion data. The aim is to alleviate the within subject mis-registration between the two modalities in order to use the predicted CSF volume fraction as prior information during the fitting of a two-compartment diffusion model.

Introduction

The presence of free water (FW) in a white matter (WM) is known to confound the interpretation of the diffusion tensor indices [10]. However, fitting a multi-compartment diffusion MRI model that includes a FW compartment is an ill-posed problem, as it often requires spatial regularisation or the introduction of prior knowledge on the tissue compartment [11] [13]. These regularisation schemes potentially reduce the sensitivity of the tensor indices to WM alteration (e.g., in pathology). Additionally, the voxel-wise fitting of the FW model estimates a non-zero CSF volume fraction in deep WM, where it is expected to be zero [12].

A T1w image is usually acquired along with the diffusion sequence in a typical group study. T1w images generally have high resolution, good SNR and a high contrast between tissues. This allows for an accurate and robust estimation of the tissue partial volume fractions (PVFs). We propose here to use the CSF PVF estimated from the T1w image as a prior in order to fit a two-compartment model on the diffusion data. We hypothesize that the extra information can improve the stability of the fitting and allow to relax the constraint on the MD.

A similar approach was previously proposed by [5]. However the crucial issue of mis-registration between the T1w image and the diffusion data was not addressed. In a typical group study, susceptibility-induced distortions associated with the echo planar imaging (EPI) and the eddy current induced distortions limit the accuracy of the co-registration. Consequently, using directly co-registered T1w-derived tissue PVFs in a diffusion model may result in significant deviations in the estimation of the tensor FA or MD.

In order to avoid the co-registration issue, we propose to infer the PVFs of tissues directly from the diffusion data. We propose to use a data-driven approach

to learn a mapping between diffusion features and the tissues' PVFs estimated from the T1w images. Supervised machine learning methods have already been successfully applied in diffusion MRI. For example, it was shown that the NODDI model can be successfully fitted on single shell data by applying a random forest regression method [1].

In this work, we compare four supervised machine learning approaches in their ability to predict the tissue PVFs. We use a range of diffusion features estimated from models applicable to low b-value single-shell data. We also evaluate the robustness of our approach to the error in registration between the T1w and diffusion images. Finally, we use the CSF PVF predicted by our method as a prior while fitting the FW model and assess how the additional information relaxes the constraints on the diffusion tensor parameters.

Materials and methods

Data acquisition and pre-processing We ran our experiments on two similar datasets for two separate subjects. Data was acquired on a 3T Siemens Verio scanner. Diffusion data consisted of a standard DWI acquisition with $TR=9600ms$, $TE=76ms$ and $b=1000s.mm^{-2}$, with 60 non-collinear diffusion gradient directions in addition to 6 non-weighted $b=0$ images. The voxel size was 2mm isotropic, with a parallel imaging acceleration factor of 2, for a total scan time of 12 minutes. The structural *T1w* image was acquired at a resolution of 1mm isotropic for a scan time of 6 minutes.

Another set of diffusion data was also acquired, consisting of an Inversion Recovery (IR) sequence where the CSF signal was nulled. The acquisition parameters were: $TR=20.6s$, $TI=2300ms$, $b=1000s.mm^{-2}$, with 60 non-collinear diffusion gradient directions in addition to 6 non-weighted $b=0$ images. This acquisition allows retrieving tensor indices that are not contaminated by the CSF contribution to the signal.

The diffusion data was preprocessed to correct for susceptibility-induced distortions with TOPUP [2] and for movement and eddy currents distortions with EDDY [3]. The diffusion features consisted of 17 parameters resulting from: the DWIs (mean b_0 , mean $b1000$), the diffusion tensor (FA, MD, tensor mode and the 3 tensor eigenvalues), Ball&Stick model [4] (Volume fractions of the 3 fibres, total fibre volume fraction, isotropic diffusivity and the dispersion of the 3 fibres orientations) and finally probabilistic tractography (visitation maps for full brain tractography). In order to add some information from the neighbouring voxels, we added a smoothed version of each feature map ($\sigma = 6mm$) as well as its gradient, for a total of 51 features.

The tissue PVFs were estimated on the T1w image using FAST [15]. These PVFs were then registered to the diffusion data with FLIRT BBR [7].

Predicting tissues volume fraction Our method relies on voxel-wise regression, that is to learn a mapping between the features derived from the diffusion

data and the corresponding tissue PVFs computed on the T1w image. A supervised machine learning (ML) approach requires a training and a testing dataset. In our case, training is performed on one subject and testing on another subject.

Comparison of the machine learning approaches We compare the performance of 4 machine learning approaches: random forest (RF), support vector machine regression (SVR), Gaussian process (GP) and neural networks (NN). The methods hyper-parameters were optimised using grid search for the RF and SVR, gradient descent for the Gaussian process and an empirical manner for the NN. The optimisation of the 4 ML approaches (RF, SVR, GP and NN) was based on the correlation between the predictions and the T1 segmentation, except for the GP, where the parameters are optimised only on the training set. Each of the 4 methods was optimised on the same training set, comprising 20000 samples.

The performance of each ML approach was similarly assessed with the correlation between their PVF predictions and the PVF estimated from the T1, for each tissue (WM, GM and CSF). The correlation was computed on the full brain of the test subject. It ensured that the performance measure includes false positives, e.g. non-zero white matter volume fraction in the CSF as well as false negatives, e.g. zero WM volume fraction in the middle of the white matter. The correlation coefficient is not expected to reach 1 because of residual mis-registration between the DWI and T1 images.

Mis-registration between T1 and diffusion space The main motivations to use a ML approach is the mis-registration between T1 and diffusion space. Therefore, we evaluated how the registration error affects the performance of the learning process. We introduced a registration error between structural and diffusion data on the learning dataset, and computed the performance of the segmentation on the test dataset was impacted.

Application to FW modelling In this experiment, we fitted the FW model described in [13]. We used the predicted CSF volume fraction as a prior for the FW compartment. The mean of the prior is equal to the volume fraction predicted, with a correction factor to account for the different T1 and T2 relaxation times of the tissues, with $FWF = \frac{\hat{S}_{fw}}{\hat{S}_{fw} + \hat{S}_T}$, where $S_X = \hat{S}_X(1 - \exp[-TR/T_{1,X}]) \exp[-TE/T_{2,X}]$. The standard deviation of the tissue PVF prior is set to 0.1, which approximately corresponds to the standard deviation of the error between our model prediction and the T1 segmentation. Compared to [13], we relaxed the prior on the MD, with its standard deviation to $0.5 \times 10^{-3} \text{ mm}^2.s^{-1}$.

Results

Machine learning approaches optimisation and comparison In figure 1, we show the performance obtained with the 4 ML approaches for the 3 tissue types (CSF, GM and WM). Overall, the RF exhibits the highest correlation between the predictions and the T1 segmentation, for the 3 tissue types ($r=0.83$, 0.88 and 0.93 for the CSF, GM and WM respectively). It is closely followed by the NN. The SVR and GP perform less well. In addition to its high accuracy, the random forest method was also the fastest. The training phase was performed in a few minutes, whereas the other methods required a few hours.




Fig. 1. Comparison of the four ML approaches. Correlation between T1w-derived and predicted (from diffusion data) partial volume segmentation using random forest (RF), Gaussian process (GP) and support vector regression (SVR)

Figure 2 shows the segmentation maps predicted by the RF method, as well as the T1w image segmentation. For a given voxel, the tissue with the highest PVF is represented. Visually, the RF segmentation maps that are very similar to the T1 segmentation. The RF still gives high WM PVF prediction in the crossing areas (> 0.95).

Mis-registration between T1 and diffusion space Figure 3 shows the correlation and MSE obtained on the test dataset as a function of the mis-registration (translation along the x, y and z axes) artificially introduced in the training dataset. To compute the scores, we used the RF regression. The plots show that the correlation is at its maximum when there is no mis-registration, reaching 0.943, with an MSE of 0.017. The correlation gradually decreases (and conversely, the MSE increases) as the registration error increases. Interestingly, we see that a mis-registration along the x axis (left-right orientation) has more impact on the performance than in the y and z axes.

Here, the key result is that the method is very robust to misregistration of the training dataset. Indeed, the correlation and the MSE obtained on the test dataset stays almost constant when the training error remains under 2mm,




Fig. 2. Segmentation maps predicted with the RF method. Tissues' partial volume fraction obtained from the T1 and RF method in diffusion space.




Fig. 3. Impact of the mis-registration on the RF performance. Correlation and MSE between predicted and ground truth PVF on the test data as a function of the misregistration between diffusion and structural space on the training data.

which corresponds to the voxel size of our diffusion data. The robustness of our approach validates its applicability to datasets where EPI distortions are significant and cannot be corrected.

Application to FW modelling We show here how the prior (derived from the RF predictions) on the CSF compartment affects the fitting of the FW model to the diffusion data. We used the RF predictions, as we previously showed that they are the most accurate. As a reference, we also included the results obtained with the T1-derived CSF PVFs, which are prone to registration error. Finally, we show the results obtained with the same model, without the prior on the CSF volume fraction.

In Table 1, we show the FA and MD obtained in the regions where partial voluming between the CSF and the WM occurs. These regions are defined by a WM PVF ≥ 0.5 , a CSF PVF ≥ 0.1 , and a GM PVF = 0. We compare our fitting with the inversion recovery acquisition where the CSF signal was nulled. We indicated the correlation as well as the MSE, the latter being more sensitive to biased estimates. The table shows that the best results (in bold) are obtained

when we used the CSF PVFs predicted by the random forest as a prior on the FW volume fraction.

Method	r (FA)	RMSE (FA)	r (MD)	RMSE (MD in $mm^2.s^{-1}$)
RF prior	0.66	0.15	0.51	0.27
T1 prior	0.57	0.19	0.35	0.31
np prior	0.62	0.17	0.49	0.34

Table 1. Correlation coefficient (r) and RMSE between the FA and MD predicted with the different priors and the inversion recovery ground truth.

Importantly, when we used the RF CSF prior, the CSF PVF estimated in the deep WM was constant with a value of 0. When the prior is not used, the model estimates a non-zero FWF (on average 0.1) in crossing fibres areas, which affects the tissue MD and FA.

These results indicate that the fitting benefits from the extra information inferred with the RF approach. Importantly, we showed that the PVFs estimated from the T1 cannot be directly used when fitting the model on the diffusion data.

Discussion

Machine learning approaches Amongst the 4 machine learning techniques we evaluated, we found that the random forest gives the most accurate predictions. It is closely followed by the NN method, while the SVR and GP are both underperforming. This result is not unexpected, as the RF approach is known to be very efficient, scalable, and effective for a wide range of supervised learning problems. The neural network approach has been shown to be promising as well, although it requires a more tedious tuning and is more computationally demanding than the RF. Additionally, the NN can be naturally extended to use the spatial neighbourhood with the convolutional neural networks (CNN). We tried to adapt an existing 3D CNN method that was originally designed for lesion segmentation [9]. Unfortunately, we faced implementation issues, as the optimisation of the weights in the context of regression proved to be difficult.

Diffusion features In this work, we employed diffusion features derived from well known models (diffusion tensor and ball and stick), as well as the tractography visitation map. One of the motivations was to make the method applicable to acquisitions with different parameters (e.g. number and orientation of diffusion gradients, resolution, number of b-values). Another motivation was that the features are rotationally independent, and therefore they do not require the training and testing datasets to be aligned. A model free approach was used in [6]. The signal intensities from each diffusion gradient were directly used as features in order to infer NODDI and DKI parameters. While their approach

appears to be promising, it is restricted to the case where the same diffusion protocol is used for an entire study.

Diffusion data segmentation We successfully applied our ML approaches to partial volume segmentation. The segmentation maps appeared homogeneous, even though no spatial regularisation was used. However, it heavily relies on the accuracy of the T1 segmentation, which was inaccurate, notably in the sub-cortical grey matter structures. Using a sub-cortical GM segmentation tool such as FIRST can provide good segmentations. Our method was unfortunately unable to properly segment those sub-cortical structures, as a significant WM PVF was still found.

Application to FW modelling The first motivation of our work was to improve the fitting of a two-compartment model that accounts for CSF contamination. By adding an informative prior on the CSF volume fraction, we were able to prevent the model from fitting a CSF compartment in deep white matter. As a result, the micro-structural properties estimated in the tissue are not biased either by the MD prior or by the presence of a spurious CSF compartment. Secondly, the CSF PVF prior enables to relax the prior on the MD. We showed that the extra information provided by the RF reduces the bias on the MD estimated by the model in areas where partial volume effects are present. We note that our work could easily be extended to a more complex diffusion model with a CSF compartment, such as NODDI [14], or models with a GM compartment, such as multi-tissue spherical deconvolution [8].

Implications for the future The analysis of diffusion data was until recently exclusively based on an explicit model that often require advanced diffusion acquisition protocols. Our results showed that a machine learning approach can provide extra-information that is useful for solving ill-posed problems. A corollary is that the diffusion data contains redundant information. It was indeed demonstrated that a machine learning approach (neural networks) can reduce the amount of data needed to fit NODDI model by a factor of 12 [6]. In another study, it was shown that using an ML based approach can infer NODDI micro-structural parameters more accurately than a direct fitting on the data [12]. They showed that their method is robust to mesoscopic changes of the white matter (i.e., the geometric padding of the fibres).

Conclusion

In this Chapter, we showed the potential of ML approaches to retrieve quantities from the diffusion that cannot be derived with an explicit model. On top of ensuring anatomical accuracy, our method also provides useful information that allows to retrieve unbiased tensor indices throughout the entire WM.

References

1. Alexander, D.C., Zikic, D., Ghosh, A., Tanno, R., Wottschel, V., Zhang, J., Kaden, E., Dyrby, T.B., Sotiropoulos, S.N., Zhang, H., Criminisi, A.: Image quality transfer and applications in diffusion MRI. *NeuroImage* **152** (May 2017) 283–298
2. Andersson, J.L.R., Skare, S., Ashburner, J.: How to correct susceptibility distortions in spin-echo echo-planar images: application to diffusion tensor imaging. *NeuroImage* **20**(2) (October 2003) 870–888
3. Andersson, J.L., Sotiropoulos, S.N.: An integrated approach to correction for off-resonance effects and subject movement in diffusion MR imaging. *Neuroimage* **125** (January 2016) 1063–1078
4. Behrens, T.E.J., Berg, H.J., Jbabdi, S., Rushworth, M.F.S., Woolrich, M.W.: Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? *NeuroImage* **34**(1) (January 2007) 144–155
5. Eaton-Rosen, Z., Melbourne, A., Cardoso, M.J., Marlow, N., Ourselin, S.: Beyond the Resolution Limit: Diffusion Parameter Estimation in Partial Volume. In: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016. Lecture Notes in Computer Science*, Springer, Cham (October 2016) 605–612
6. Golkov, V., Dosovitskiy, A., Sperl, J.I., Menzel, M.I., Csisch, M., Smann, P., Brox, T., Cremers, D.: q-Space Deep Learning: Twelve-Fold Shorter and Model-Free Diffusion MRI Scans. *IEEE Transactions on Medical Imaging* **35**(5) (May 2016) 1344–1351
7. Greve, D.N., Fischl, B.: Accurate and robust brain image alignment using boundary-based registration. *NeuroImage* **48**(1) (October 2009) 63–72
8. Jeurissen, B., Tournier, J.D., Dhollander, T., Connelly, A., Sijbers, J.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage* **103** (December 2014) 411–426
9. Kamnitsas, K., Ledig, C., Newcombe, V.F.J., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B.: Efficient multi-scale 3d CNN with fully connected CRF for accurate brain lesion segmentation. *Medical Image Analysis* **36** (February 2017) 61–78
10. Metzler-Baddeley, C., O’Sullivan, M.J., Bells, S., Pasternak, O., Jones, D.K.: How and how not to correct for CSF-contamination in diffusion MRI. *NeuroImage* **59**(2) (January 2012) 1394–1403
11. Pasternak, O., Sochen, N., Gur, Y., Intrator, N., Assaf, Y.: Free water elimination and mapping from diffusion MRI. *Magnetic Resonance in Medicine* **62**(3) (September 2009) 717–730
12. Reisert, M., Kellner, E., Dhital, B., Hennig, J., Kiselev, V.G.: Disentangling micro from mesostructure by diffusion MRI: A Bayesian approach. *NeuroImage* **147** (February 2017) 964–975
13. Vallee, E., Douaud, G., Monsch, A., Gass, A., Wu, W., Smith, S., Jbabdi, S.: Modelling free water in diffusion MRI. In: *Proc. Intl. Soc. Mag. Reson. Med.* (2015) 0474
14. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: Practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4) (July 2012) 1000–1016
15. Zhang, Y., Brady, M., Smith, S.: Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging* **20**(1) (January 2001) 45–57