Beyond the resolution limit: parameter estimation in partial volume

Zach Eaton-Rosen[1], Andrew Melbourne[1], M. Jorge Cardoso[1], Neil Marlow[2], Sebastien Ourselin[1]

[1]Translational Imaging Group, University College London, UK
[2]Academic Neonatology, EGA UCL Institute for Women's Health, London, UK

Abstract. Diffusion MRI is a frequently-used imaging modality that is used to infer microstructural properties of tissue, down to the scale of microns. For single-compartment models, such as the diffusion tensor, the interpretation of the models depends on the voxels having homogeneous tissue. This limitation makes it difficult to directly measure diffusion parameters for small structures such as the fornix, which may have partial volume in every voxel. In this work, we use a segmentation from a structural scan to calculate the tissue composition for each voxel in diffusion space. We model the signal from a voxel as a linear combination of the signals from these components, and thus fit parameters on a per-region basis. Fitting diffusion parameters to all regions, simultaneously, bolsters the data for small regions, which allows accurate estimation of the diffusion parameters. We test the proposed method by using diffusion data from the HCP. We downsample the HCP data, and show that our method returns parameter estimates that are closer to the high-resolution ground truths than region-of-interest methods. We apply our technique to the diffusion parameters in the fornix for adults born extremely preterm and matched controls. We show that our method estimates diffusion parameters accurately for structures that are small, compared to a typical diffusion MRI resolution.

1 Introduction

Diffusion imaging is a vital tool for probing the microstructure of *in-vivo* tissue. Parametric models of diffusion offer an informative way to summarise the information from many different b-values and gradient directions. The model parameters are often averaged over a region, under the reasonable assumption that tissue within a structure will have similar diffusion properties. This approach works well in large regions, where we can erode a probabilistic segmentation to obtain voxels that are fully within the tissue. But, the diffusion parameters within structures such as the fornix - a narrow white matter structure surrounded by cerebrospinal fluid - may not be measured by this approach [1]. Because of the large scale of diffusion MRI voxels relative to the fornix a majority of, and perhaps all, voxels will contain partial volume. This partial volume affects the measurements for fitting diffusion parameter models, especially in regions with

a spatial extent on the order of a voxel width (for example, the fornix). If we selected a threshold-based segmentation, we may have to adjust the threshold probability over time if the fornix grew or atrophied, biasing the experiment. If we chose a low percentage threshold throughout, we would be biasing the experiment by including more or less partial volume depending on the structure's size, when we are more interested in its microstructure.

In this work, we extend the calculation of a region's parameters to include information from all voxels in the region *during* the model-fitting, instead of fitting voxel-wise and then averaging. While there has been work on eliminating the contribution of free water to diffusion parameter estimates [2], the proposed approach directly estimates the diffusion parameters of all tissue types within the image, without relying on *a priori* diffusion models or values. In the proposed framework, we use a probabilistic segmentation as weights for canonical diffusion signals, optimised for each segmentation class (henceforth referred to as *tissues*). The modelled signal in a voxel is a linearly weighted sum of each tissue present in the voxel, where the weights are given by the segmentation probabilities.

We first validate the method on *in-vivo* diffusion data from the Human Connectome Project [3]. By using such high-resolution data, we can measure diffusion parameters in the fornix directly, using hand-drawn regions of interest. By downsampling this data, we simulate a more typical diffusion acquisition, and are able to test whether our approach retrieves the correct parameter values. After validating our approach, we apply it to adults born extremely preterm, comparing the diffusion within the fornix to term-born controls. The comparison is interesting as the patient group has pervasive differences in brain morphology and function, including memory (associated with the fornix).

This framework presented is similar to [4] in its use of multi-modal imaging to make a diffusion mixture model. This work differs in that there is no requirement of multiple shells of diffusion data, an important advantage for using this method in older data.

2 Methods

2.1 Theory

In this work we attempt to measure diffusion parameters from below the resolution at which they were obtained. If we imagine a voxel at a higher resolution (for example in the T_1-weighted scan) being downsampled to a lower resolution (the diffusion scans), the proportion of the tissue in a voxel of diffusion space will be reduced. Even in a best case scenario, the probability of there being at least a threshold T% of the tissue within a voxel depends on the position of the tissue relative to the voxel borders. Our approach eliminates any dependence of measured parameters on the precise voxel boundaries, by using all diffusion information within the region of interest.

Varying b-values and gradient directions during a diffusion scan establishes how water diffuses in each voxel, which is summarised with a mathematical

model. Within a given voxel, the water diffusion from several microstructural environments is measured together. A voxel's signal, S, in the DT model, is given by:

$$\frac{S}{S_0} = e^{-b\mathbf{g}^T\mathbf{D}\mathbf{g}} \tag{1}$$

where S_0 is the diffusion signal with zero diffusion weighting, b are the b-values, $\mathbf{g}$ are the gradient directions and $\mathbf{D}$ is the second-rank diffusion tensor.

In our approach, we are trying to obtain $\mathbf{D}$ for each of the k tissue classes (in the case of the fornix, white matter, grey matter and CSF). In a given voxel, we model the signal as being represented as a weighted sum of each of the tissue classes that are present:

$$S = S_0 \sum_{j=1}^{k} p_j e^{-b\mathbf{g}^T\mathbf{D}_j\mathbf{g}} \tag{2}$$

$\mathbf{D_j}$ is now the diffusion tensor for a given region. The p_j are non-negative, and constrained between 0 and 1. This approach considers the components of the $\mathbf{D_j}$ as unknown parameters that are optimised to best fit the data. This is a mixture-model approach, that generates the diffusion parameter estimates for the entire volume simultaneously, instead of per voxel. For the case of two tissue classes, this reduces to the signal model in [1].

In order for a single diffusion tensor to represent the diffusion properties in different voxels, we must account for different orientation in different parts of the same tissue. In the conventional approach, we would use the same b-matrix for each voxel in the image. However, we are mainly interested in orientationally-independent measurements, such as the fractional anisotropy (FA). In this work, we redefine the gradient-directions for each voxel, so that the principal directions of all voxels in the image align. The gradient directions at each voxel are calculated by first, performing a tensor-fit to the voxel and establishing V_1 and V_2, the first and second eigenvectors of $\mathbf{D}$. We then calculate the rotation matrix R such that V_1 and V_2 align with [1,0,0] and [0,1,0]. Our vector for the i^{th} voxel then becomes $\mathbf{g}_i = R\mathbf{g}$.

After calculating principal diffusion directions in every voxel, and the S_0, with a weighted-least-squares tensor fit, we initialise a $3 \times k$ matrix with identical diffusivities in each of the tissue classes. At each iteration of the optimisation, we calculate the signal for the entire volume simultaneously, before the $3k$ diffusion parameters are updated. We fit using Matlab 2014b, using non-linear optimisation [5].

2.2 Data

To test the proposed approach, we use data from the Human Connectome Project (HCP). This diffusion data has a resolution of 1.25^3mm^3, with 108 volumes with $b \leq 1200s.mm^{-2}$. The T$_1$-weighted MRI is at resolution 0.70^3mm^3.

For the experiments on adult subjects, we collected MRI data at 19 years of age from 15 adolescents. Eight (4 Male) of these were born extremely preterm

(fewer than 26 weeks completed gestation) and seven (3 Male) were recruited as matched controls. We acquired 3D T1-weighted volume at 1 mm isotropic resolution (TR/TE = 6.78/3.06 ms) for segmentation and diffusion MRI with the following characteristics: Diffusion-weighted data was acquired across four b-values at b = $(0, 300, 700, 2000)s.mm^{-2}$ with n = (4, 8, 16, 32) directions respectively at TE=70ms (2.5x2.5x3.0mm). For the fitting, we discarded the highest shell of b-values. All data was acquired using a Philips 3T Achieva. For the segmentations, we manually drew the fornix on a T_1-weighted segmentation and also labelled the surrounding tissue using multi-atlas label propagation and fusion [6] based on the Neuromorphometrics, Inc. labels.

3   Experiments and Results

3.1   Validating method using HCP data.

We used the high resolution of the HCP data to determine pseudo ground-truth values for diffusion parameters in the column, the crus and the body of the fornix. Each of these regions is hand-drawn onto the subject's T_1-weighted MRI. We downsample the segmentation from categorical labels into a probabilistic diffusion segmentation, where the probabilities represent fractions of the tissue in that diffusion voxel. We varied the downsampling to achieve voxels of isotropic dimension from 1.25mm to 3.5mm. In order to use HCP data as a model for a more typical diffusion acquisition, we adjust it in the following ways. We added rician noise to the downsampled data, to bring the data to a clinically realistic SNR. We exclude any volumes corresponding to b values of over 1200, to ignore effects that are not modelled with the DT. We also used a subset of the 108 available diffusion volumes. We tested the performance of our algorithm with varying numbers of these subsets, using between 12 and 60 readings. We compare three approaches for analysing average parameter values:

M1 For each region, we identify voxels where the membership to that region is above the threshold and average their values.
M2 We resample the downsampled DWI to high-resolution HCP space before fitting the DT model and, again, averaging the values for each region. For this, we use 7th-order b-splines, as recommended in [7].
M3 (proposed): We calculate parameter values for each region, explicitly accounting for partial volume. The p are given by downsampling labels from the T_1-weighted segmentation into diffusion space.

3.2   Results

In Figure 1 we test approach M1. As the threshold changes, so do the results for the classical approach. With a 90% threshold, at a resolution of 2mm isotropic, all fornix tissue has partial volume, and so even at this good resolution, we would be unable to proceed with this threshold. However, as we decrease the threshold, increasing the resolution results in decreasing FA, as CSF partial
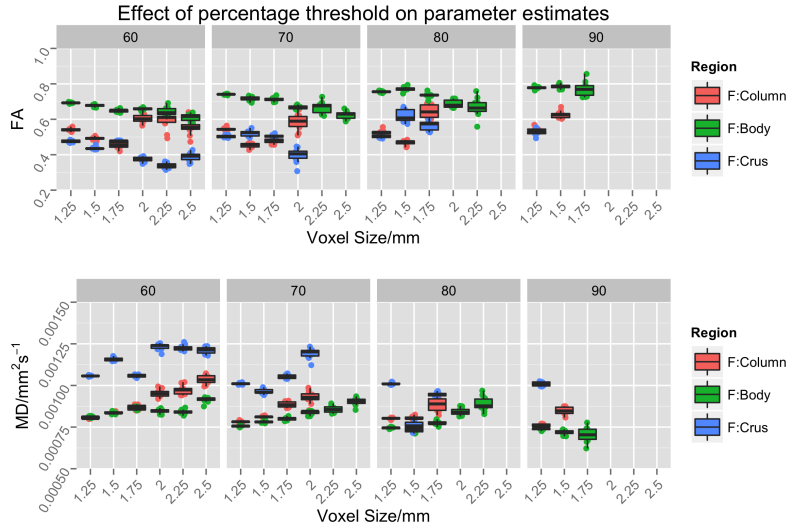
Fig. 1. In this graph, we see the effect of downsampling the resolution (x axis) on parameter estimates from a threshold-based approach. As the voxel dimension increases, the parameter estimation is less reliable and at some point stops, as there are no more supra-threshold voxels to sample. The choice of the threshold will influence the measured parameter value.

volume contaminates the estimates. For the body of the fornix, the measured FA decreases by up to 13% by changing the thresholding, and when downsampled to 2.5mm, the measured FA is up to 25% lower than the pseudo-gold-standard.

In approach M2, we use the segmentation at the HCP resolution. After downsampling the data and adding noise, we interpolate the diffusion data back to the HCP resolution in order to fit the diffusion tensor and average the results over the ROI. These results are displayed in Figure 2. The measured diffusion parameters diverge from their 'true' values as we interpolate data of lower resolution. With no downsampling, the values of nearby white matter, the column and the crus of the fornix are similar. However, as the downsampling increases, the FA estimates decrease due to the partial volume. This means that parameter values that should be similar are diverging because of local surroundings.

With the proposed method, the FA in the column and crus of the fornix is constant (Figure 2). The body of the fornix has an increasing FA. The mean diffusivities are more constant in the proposed method than with the classical. Downsampling the data affects the accuracy of all the measured parameters using the classical approach.

3.3 Comparison of preterm-born and term-born young adults.

We compare fornix DTI parameters as calculated with M2 (upsampling DWI to T_1-space and fitting the tensor) and M3 (proposed) in Figure 3. The MD in the
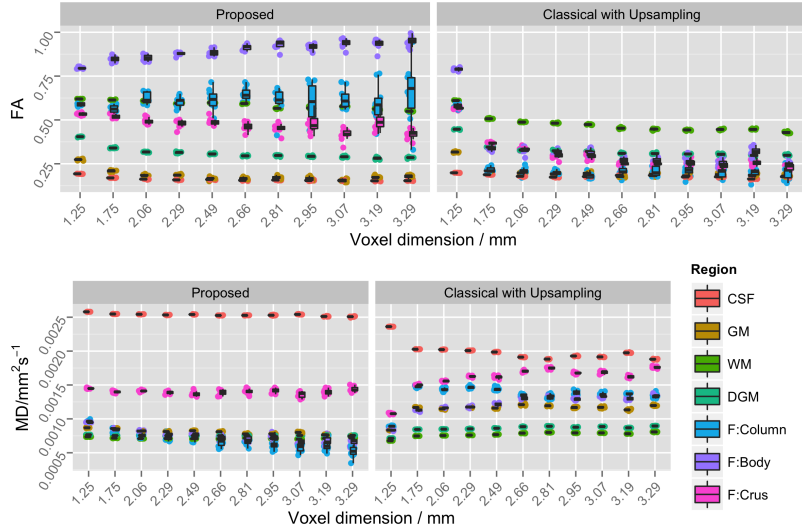
Fig. 2. In our method (left), the diffusion parameters in the fornix are fairly consistent with downsampling. While the results for larger regions match the classical approach, we improve for the fornix. We present the results for a thresholding approach *with* prior upsampling of the data (right). The results here, for diffusion parameters of the fornix, show a divergence of the diffusion parameter readings depending on their surrounding tissue. The scale factor is the factor by which we've downsampled the volume. In these experiments, we used 12 diffusion readings, 2 of which were reference volumes.

fornix is higher in general for the classical approach compared to ours. In the classical approach, there is a significant difference in the MD with the subject group having higher MD (p $\leq$ 0.0005). Both approaches measure higher mean FA in the control group, but neither is significant when accounting for multiple comparisons. The optimisation took less than a minute for each subject.

4 Discussion

Our method achieves consistent and accurate parameter estimates for small regions in partial volume. Although interpolating data reveals some details that are hidden at low resolution [7], interpolation of downsampled HCP data biased the results of the measured diffusion parameters in the fornix. FA values in all parts of the fornix tended to be underestimated and diverged from FA estimates in other white matter regions.This means that the local surroundings of the fornix biased the diffusion results, which our method was able to address.

There is promise in using this approach in subject groups, such as the preterm-born young adults in this study. Our approach reduces the impact of partial volume on measuring the properties of the fornix. The lower MD we measured for subjects and controls is in accord with this. The higher FA values suggest that we are able to measure the diffusion in the highly-anisotropic
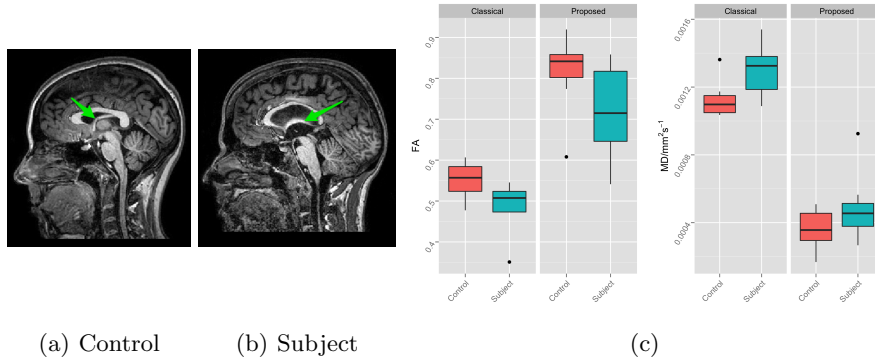
(a) Control   (b) Subject   (c)

Fig. 3. In a-b the fornix is highlighted with an arrow in a control and a preterm subject. The preterm-born subject has noticeable abnormalities in the corpus callosum, and enlarged ventricles. In c, we display the measured parameters using our proposed approach vs the classical.

region of the fornix with less impact from the surrounding cerebrospinal fluid. While this is not a conclusive result, due to the small number of subjects, it is a promising sign. There is some evidence from both methods that preterm-born adults have lower FA in the fornix than controls, which is congruent with the general consequences of being born prematurely.

While there are a range of biophysical compartment models in use in diffusion imaging, most of these rely on multi-shell data to fit compartments in each voxel, or else have to heavily restrict the available parameters. We circumvent this by fitting on a per-tissue basis, by using information from a structural segmentation. This means our method has the advantage that it is completely determined by the data.

Another way to calculate diffusion parameters would be investigating super-resolution techniques. Alexander et al [8] showed that a machine-learning approach could be used to effectively super-resolve the diffusion data from DTI parameters. For this particular method, it is unclear how generalisable the approach is without high-resolution training data from each scanner in use. Our validation only used 12 diffusion volumes, and no training data, which renders the method suitable to past datasets.

We show that it is feasible and possible to estimate diffusion parameters for regions that are small on the scale of diffusion MRI. In large, contiguous regions we achieved the same results as for the classical approach, of thresholding and averaging. We used the fornix as a region of interest to show that our approach was able to recover diffusion parameter estimates consistently, when the classical approach failed. Although results were good in the fornix, the model would have to be extended significantly to cope with geometry such as crossing fibres.

The presented approach achieves close-to gold-standard results with minimal processing time and requirements for the diffusion acquisition. This is because we aggregate data from all voxels in which a particular region is present, even

in part. This type of approach fits conceptually with more sophisticated, multi-compartment models - we plan to extend our investigation to using multi-modal segmentations to determine volume fractions for certain tissue types. In this work, we proposed a method to extract diffusion tensor parameters from tissue that has partial volume. We have validated the method using high-quality data from the HCP, and applied it in a new cohort of clinical interest.

References

1. 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) (2012) 1394–1403
2. Pasternak, O., Sochen, N., Gur, Y., Intrator, N., Assaf, Y.: Free water elimination and mapping from diffusion MRI. Magnetic Resonance in Medicine **62**(3) (2009) 717–730
3. Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E.J., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., Della Penna, S., Feinberg, D., Glasser, M.F., Harel, N., Heath, A.C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., Oostenveld, R., Petersen, S.E., Prior, F., Schlaggar, B.L., Smith, S.M., Snyder, A.Z., Xu, J., Yacoub, E.: The Human Connectome Project: A data acquisition perspective. NeuroImage **62**(4) (2012) 2222–2231
4. Eaton-Rosen, Z., Cardoso, M.J., Melbourne, A., Orasanu, E., Bainbridge, A., Kendall, G.S., Robertson, N.J., Marlow, N., Ourselin, S.: Fitting parametric models of diffusion MRI in regions of partial volume. In: Proceedings of SPIE, Medical Imaging 2016 : Image Processing. Volume 9784. (2016)
5. Coleman, T.F., Li, Y.: An Interior Trust Region Approach for Nonlinear Minimization Subject to Bounds. SIAM Journal on Optimization **6**(2) (1996) 418–445
6. Cardoso, M.J., Modat, M., Wolz, R., Melbourne, A., Cash, D., Rueckert, D., Ourselin, S.: Geodesic Information Flows: Spatially-Variant Graphs and Their Application to Segmentation and Fusion. IEEE Transactions on Medical Imaging (2015) 1–1
7. Dyrby, T.B., Lundell, H., Burke, M.W., Reislev, N.L., Paulson, O.B., Ptito, M., Siebner, H.R.: Interpolation of diffusion weighted imaging datasets. NeuroImage **103** (2014) 202–213
8. Alexander, D.C., Zikic, D., Zhang, J., Zhang, H., Criminisi, A.: Image quality transfer via random forest regression: Applications in diffusion MRI. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) **8675 LNCS**(PART 3) (2014) 225–232