

Machine learning based compartment models with permeability for white matter microstructure imaging

Gemma L. Nedjati-Gilani^{a,b,1}, Torben Schneider^{d,1}, Matt G. Hall^c, Niamh Cawley^d, Ioana Hill^a, Olga Ciccarelli^d, Ivana Drobnjak^{a,*,2}, Claudia A.M. Gandini Wheeler-Kingshott^{d,e,f,2}, Daniel C. Alexander^{a,2}

^a Centre for Medical Image Computing and Dept of Computer Science, University College London, Gower Street, London WC1E 6BT, UK

^b MRC Centre for Outbreak Analysis and Modelling, Department of Infectious Disease Epidemiology, Imperial College London, London, UK

^c Institute of Child Health, University College London, London, UK

^d Department of Neuroinflammation, Institute of Neurology, University College London, London, UK

^e Brain MRI 3T Mondino Research Center, C. Mondino National Neurological Institute, Pavia, Italy

^f Department of Brain and Behavioural Sciences, University of Pavia, Pavia, Italy

ABSTRACT

Some microstructure parameters, such as permeability, remain elusive because mathematical models that express their relationship to the MR signal accurately are intractable. Here, we propose to use computational models learned from simulations to estimate these parameters. We demonstrate the approach in an example which estimates water residence time in brain white matter. The residence time τ_i of water inside axons is a potentially important biomarker for white matter pathologies of the human central nervous system, as myelin damage is hypothesised to affect axonal permeability, and thus τ_i . We construct a computational model using Monte Carlo simulations and machine learning (specifically here a random forest regressor) in order to learn a mapping between features derived from diffusion weighted MR signals and ground truth microstructure parameters, including τ_i . We test our numerical model using simulated and in vivo human brain data. Simulation results show that estimated parameters have strong correlations with the ground truth parameters ($R^2 = \{0.88, 0.95, 0.82, 0.99\}$) for volume fraction, residence time, axon radius and diffusivity respectively), and provide a marked improvement over the most widely used Kärger model ($R^2 = \{0.75, 0.60, 0.11, 0.99\}$). The trained model also estimates sensible microstructure parameters from in vivo human brain data acquired from healthy controls, matching values found in literature, and provides better reproducibility than the Kärger model on both the voxel and ROI level. Finally, we acquire data from two Multiple Sclerosis (MS) patients and compare to the values in healthy subjects. We find that in the splenium of corpus callosum (CC-S) the estimate of the residence time is 0.57 ± 0.05 s for the healthy subjects, while in the MS patient with a lesion in CC-S it is 0.33 ± 0.12 s in the normal appearing white matter (NAWM) and 0.19 ± 0.11 s in the lesion. In the corticospinal tracts (CST) the estimate of the residence time is 0.52 ± 0.09 s for the healthy subjects, while in the MS patient with a lesion in CST it is 0.56 ± 0.05 s in the NAWM and 0.13 ± 0.09 s in the lesion. These results agree with our expectations that the residence time in lesions would be lower than in NAWM because the loss of myelin should increase permeability. Overall, we find parameter estimates in the two MS patients consistent with expectations from the pathology of MS lesions demonstrating the clinical potential of this new technique.

1. Introduction

Techniques such as AxCaliber (Assaf et al., 2008) and ActiveAx (Alexander et al., 2010) use computational models to provide estimates of tissue microstructure properties, such as cell size or packing density,

from diffusion-weighted (DW) MR data. These models use simple geometries such as cylinders and spheres to represent axons and other cells, so that closed form mathematical expressions can be derived that closely approximate the expected MR signal. However, some tissue properties such as permeability or fibre undulation, have previously

* Corresponding author.

E-mail address: i.drobnjak@ucl.ac.uk (I. Drobnjak).

¹ These authors contributed equally to this work.

² These authors contributed equally to this work.

remained elusive because mathematical models that express their relationship to the MR signal accurately are intractable. Numerical simulations (Hall and Alexander, 2009) provide flexible forward models that can capture such effects, however they are computationally expensive to use in solving the corresponding inverse problem to estimate parameters from data. Here we propose to use a machine learning approach to learn the mapping between the underlying model parameters and the diffusion signal using numerical simulations, and then use that mapping to estimate the parameters. We demonstrate this idea on the residence time τ_i of water molecules inside the white matter axonal fibres.

There is a widespread interest in developing imaging biomarkers based on the permeability associated with the intra-axonal water residence time τ_i . Permeability is an important property of axonal membranes, and its changes have been associated with the condition of the insulating myelin sheet (Nilsson et al., 2013) and with various neurological conditions, such as Multiple Sclerosis (MS) or Parkinson's disease (Volles et al., 2001). DW MRI is potentially able to estimate τ_i as it is sensitive to the motion of water molecules within tissue. However, progress has been limited by lack of sufficiently accurate mathematical models relating τ_i to the MR signal.

The analytical expressions used to calculate signals of white matter tissue assume that the water molecules are completely restricted and thus ignore or use simplistic models of permeability (Van Gelderen et al., 1994; Callaghan, 1997; Codd and Callaghan, 1999). A new approach by Grebenkov et al. (2014), assuming high gradients and narrow pulses, may help to address these issues but current results are only applicable to compartment sizes substantially larger than human cells.

Other mathematical models, such as the Kärger model and apparent exchange rate (AXR) imaging, explicitly incorporate τ_i by considering the effect of exchange among water pools. Of the two approaches, the Kärger model (Kärger et al., 1988) is most commonly used (Stanisz et al., 1997; Lätt et al., 2009; Nilsson et al., 2010) as it is compatible with data acquired using widely available pulsed gradient spin echo (PGSE) and stimulated echo (STE) imaging sequences. It accounts for inter-compartmental water exchange by coupling the DW MR signals due to the separate compartments via τ_i . However, it relies on the assumption that the two individual pools of water are well mixed and it does not model restriction which in white matter tissue is prominent due to the presence of axonal membranes which restrict the motion of water molecules. The Kärger model may still provide a reasonable description of the diffusion-weighted signal in the long time limit, however the cell membranes have to be close to impermeable, and in the case of fast exchange the model fails (Fieremans et al., 2010). Even though these limitations have been known for over 20 years, there have been no improvements to the model due to the mathematical intractability of the problem. AXR imaging (Lasič et al., 2011; Nilsson et al., 2013) has recently been introduced as an alternative to estimating τ_i via Kärger model; however it requires specialised double diffusion encoding imaging sequences (DDE) (Shemesh et al., 2016) and it again relies on strong assumptions about the compartmentation of water into a 'fast' and 'slow' pool. The estimated AXR parameter also conflates τ_i with intra-axonal volume fraction f , making it difficult to disentangle the origin of any measured change. Furthermore, AXR is based on Gaussian compartments, with time-independent Gaussian diffusion in each compartment, and allows for only two compartments otherwise further conflation occurs (Lasič et al., 2016).

Given the inherent difficulties involved in deriving accurate analytical models of exchange with permeable axonal membranes, there have been alternative approaches which bypass mathematical models altogether and use simulations to learn how permeability affects measured DW MRI signals. Nilsson et al. (2010) generate libraries of microstructure parameters and their corresponding DW MR signals from Monte Carlo (MC) simulations, and use them to find the nearest-neighbour microstructure parameters that matched unseen data, i.e.

combinations of tissue parameter values not explored in the training set. However, nearest-neighbour matching typically has poor generalisation. Furthermore, using the raw signals to perform the matching is also potentially inefficient, as it requires new libraries to be generated for each acquisition protocol used and each library entry is specific to a particular fibre orientation.

Our machine learning approach constructs a mapping between microstructural parameters of interest and orientationally invariant features derived from the DW MR data. We use MC simulations to generate synthetic signals from a library of histologically relevant microstructure indices. A random forest regressor then learns the mapping between features derived from DW MR signals and ground truth microstructure parameters using the synthetic data, providing an efficient and accurate method for estimating microstructure parameters, including τ_i , that generalises smoothly between training examples. We compare our approach to the Kärger model using simulated data and demonstrate that the trained random forest can be used to estimate sensible and consistent estimates of microstructure indices from *in vivo* healthy human brain white matter. Finally, we demonstrate the clinical relevance of this approach using data acquired from two patients with MS.

2. Methods and materials

2.1. Imaging protocol

We use an orientationally invariant, DW-PGSTE protocol with “echoplanar imaging” (EPI) sequence, optimised (Alexander, 2008) for a two-compartment model with exchange, assuming a maximum imaging time of approximately 30 minutes. PGSTE sequences are particularly well suited for estimating τ_i , as they allow us to probe long diffusion times by exploiting the longer T_1 compared to T_2 relaxation rates. **Here we use an orientationally invariant protocol, which does not depend on the orientation of the fibres and is optimised assuming that the fibre orientation is unknown.** The protocol is optimised for the following biophysically plausible tissue parameters, assuming that axons are modelled as infinitely long, randomly packed parallel cylinders: intra-axonal volume fraction $f=0.7$, intrinsic diffusivity $d = 2 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$, apparent diffusivity in the perpendicular direction in the extracellular space $d_{\perp} = 0.7 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$, axon radius $R=1 \mu\text{m}$, intracellular residence times $\tau_i \in \{0.05, 0.1, 0.2, 0.4, 1\} \text{ s}$, $T_1=0.832 \text{ s}$ (Stanisz et al., 2005). Exchange is incorporated in the optimisation via the Kärger formalism. This choice of model and parameter values in the protocol optimisation are not critical, here we design it to find a combination of measurements broadly sensitive to the white matter parameters of interest: f , d_{\perp} , d , R and τ_i . For example, axon radius of $1 \mu\text{m}$ is relatively large for human brain (Aboitiz et al., 1992), however since the clinical scanners are not sensitive to small diameters (Drobnjak et al., 2016), choosing a smaller R would not make any difference in the optimal sequence. The resulting protocol contains 108 measurements divided into 4 equal shells, with Δ ranging from 0.102 to 0.412 s. The final protocol, with nominal b values, is shown in Table 1. The additional crusher and slice select gradients used in the PGSTE sequence result in altered diffusion weighting, which we account for in the data generation (Alexander and Dyrby, 2013; Lundell et al., 2014).

2.2. Monte Carlo simulations

The first step in constructing our computational model is to use Monte Carlo simulations (Hall and Alexander, 2009) in combination with the imaging protocol in Table 1 to generate synthetic diffusion MR signals from digital phantoms with a wide range of plausible ground truth microstructure indices.

Each phantom is characterised as a unique combination of five parameters: the mean μ_R and standard deviation σ_R of the axon radius

Table 1
PGSTE protocol parameters, optimised for Kärger model parameter estimation.

Shell	# $b=0$ (s mm ⁻²)	# gradient directions	b (s mm ⁻²)	$ G $ (mT m ⁻¹)	Δ (s)	δ (s)	TM (s)
1	4	23	1622	62	0.102	0.0077	0.07
2	4	23	1718	62	0.412	0.0039	0.375
3	4	23	3611	62	0.406	0.0057	0.37
4	4	23	4031	62	0.169	0.0094	0.135

distribution, the intrinsic diffusivity of the spins d , the intracellular volume fraction f and the intracellular water residence time τ_i . Here, intracellular residence time τ_i is defined as the average time a water molecule spends inside the intracellular space represented by the white matter axonal fibres. This is connected to permeability through the expression: $k = \frac{R}{2\tau_i}$, where R is the axon radius and k is the permeability. **This, and the more general form of this expression are described in Jespersen et al. (2005)** The membrane permeability k is specified via a probability p of the water molecule stepping through a membrane when it encounters it. The molecule will be reflected back with a probability of $1-p$. A value of $p=0$ corresponds to impermeable membranes while a value of $p=1$ corresponds to fully permeable membranes. The relationship between the probability p and the membrane permeability k is given by the expression: $p = k \sqrt{6 \frac{dt}{d}}$, where dt is the simulation time increment and d is the intrinsic diffusivity.

To mimic the structure of in vivo human brain data, we model white matter as a collection of non-abutting parallel cylinders with radii drawn from a gamma distribution. We construct 12,500 unique white matter substrates, with substrate parameters randomly selected in the ranges: $\mu_R \in [0.2, 5] \mu\text{m}$, $\sigma_R \in [\min(0.1, \frac{\mu_R}{5}), \frac{\mu_R}{2}] \mu\text{m}$ (to ensure that the distributions have a non-zero mode, matching the distributions observed in histology (Aboitiz et al., 1992)), $f \in [0.4, 0.7]$, $\tau_i \in [0.02, 0.95] \text{ s}$, $d \in [0.8, 2.2] \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$. The cylinders are randomly packed in the substrates as described in Hall and Alexander (2009), with example substrates shown in Fig. 1.

All simulations are performed using 100,000 spins and 2000 time steps. We chose these values as they provide precision of 10^{-10} of the unweighted signal, which is several orders of magnitude smaller than realistic signal noise (Hall and Alexander, 2009). The diffusion step size calculated as in Hall and Alexander (2009) is between $0.5 \mu\text{m}$ and $0.9 \mu\text{m}$, depending on the diffusivity of each substrate. We also choose


100,000 non-abutting parallel cylinders in a substrate to avoid variation due to sampling of the gamma distribution (Hall et al., 2014).

We generate two sets of signals for each geometry: noise-free and noisy. As spins undergo T_1 relaxation during the mixing time TM between the second and third 90 degree RF pulses, and thus also between the two diffusion gradients, measurements made using longer Δ (and so longer TM) experience more relaxation leading to lower signal intensities and signal to noise ratios (SNR). For the noisy data set we scale the signals by $\exp(-TM/T_1)$ using $T_1 = 0.832 \text{ s}$ for white matter at 3 T (from Stanisz et al., 2005). We then add Rician noise, choosing the standard deviation of the noise σ so that the SNR of the $b=0$ images with $\Delta=0.102 \text{ s}$ is 20 which reflects our in-vivo data.

2.3. Data acquisition

We scan two healthy subjects and two MS patients. The first healthy volunteer is a 32 year old male, and the second healthy volunteer is a 30 year old female. Both healthy volunteers are scanned twice, with the rescan taking place within a month of the initial scan. The first MS patient is a 22 year old female with relapsing remitting multiple sclerosis (RRMS), a disease duration of 2 years and expanded disability status scale (EDSS) of 2. The second MS patient is a 63 year old male with secondary progressive multiple sclerosis (SPMS), a disease duration of 25 years and EDSS of 6. Both MS patients were scanned once. None of the MS patients recruited into the study experienced a relapse or a course of corticosteroids three months prior to imaging. We obtained written informed consent for all participants, and the study was approved by our local research ethics committee.

All data are acquired using the imaging protocol in Table 1 on a 3 T Philips Achieva scanner, using a 32-channel receive-only RF coil, SENSE acceleration factor of 2.5, full k-space acquisition, and receiver bandwidth of 1948 Hz/pixel. Additional imaging parameters are as



(a) $\mu_R = 2.69 \mu\text{m}$, $\sigma_R = 0.68 \mu\text{m}$, $f=0.44$, $\tau_i = 0.344 \text{ s}$. (b) $\mu_R = 0.39 \mu\text{m}$, $\sigma_R = 0.13 \mu\text{m}$, $f=0.68$, $\tau_i = 0.564 \text{ s}$.

Fig. 1. Example substrates for cylindrical geometries, showing variations in both cell size distribution and packing fraction.

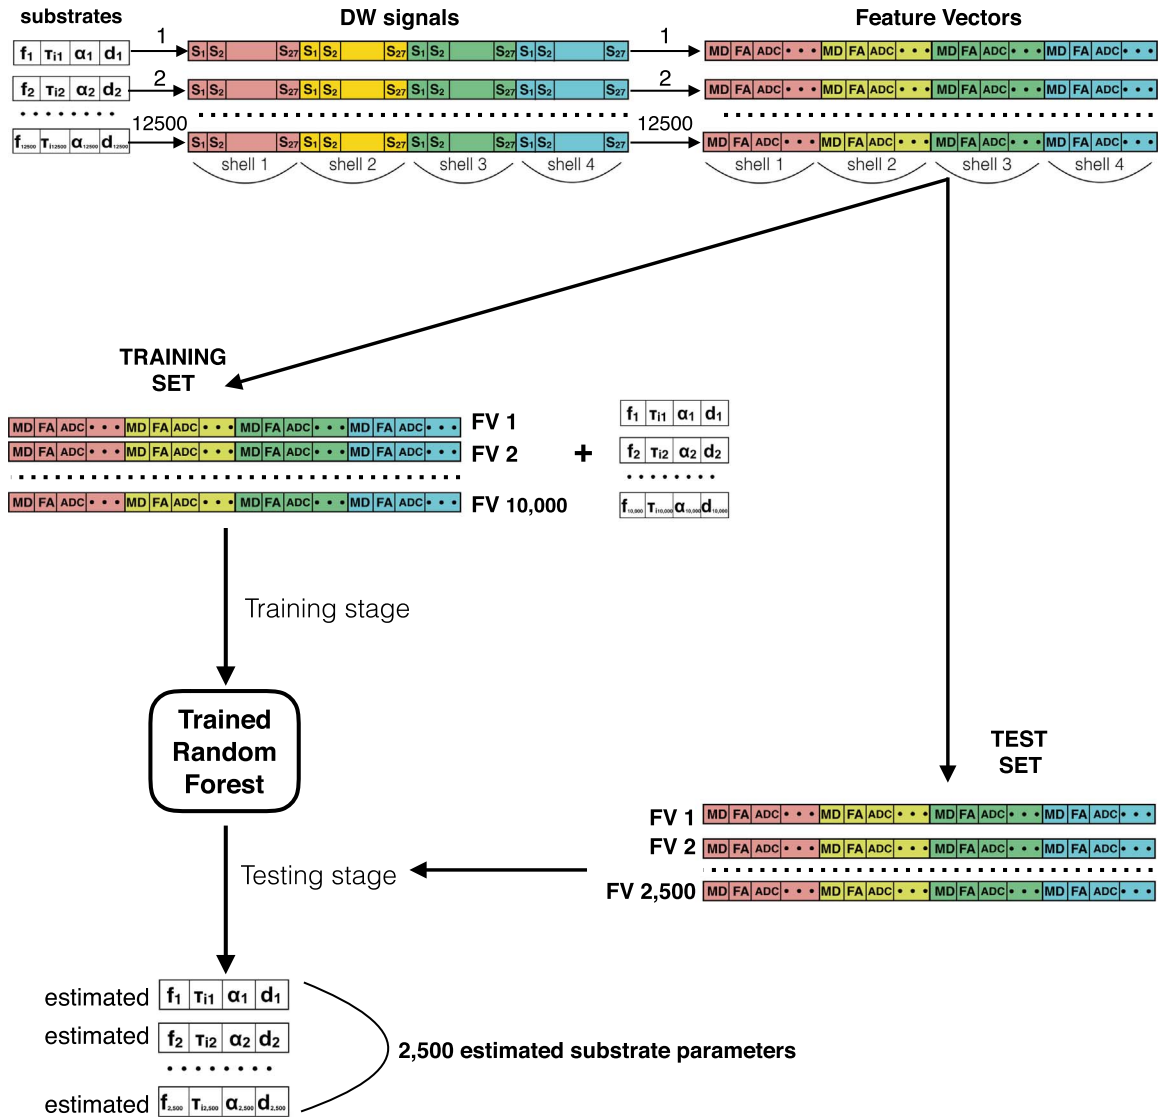


Fig. 2. A schematic overview of how testing, training and estimation is done with random forest regression. In the appendix we go into further details of the random forest regressor itself.

follows: $TE=0.068$ s, $TR=12$ s, $FOV=256$ mm \times 256 mm, matrix size=128 \times 128, slice thickness=4 mm, number of slices=40 (the first scan of the first volunteer has 30 slices). The total imaging time is approximately 40 minutes. For the MS patients, in addition to the diffusion-weighted data we also acquire a high resolution (1 mm \times 1 mm \times 3 mm) T_2 -weighted scan for the purposes of creating a lesion mask.

Following the data acquisition, some post-processing steps are performed. We perform eddy current and motion correction using the eddy tool in FSL (Smith et al., 2004). As the model that we learn here for white matter is specific to axons that resemble parallel cylinders, it is not applicable in regions containing CSF, grey matter or highly dispersed or crossing white matter fibres. We mask out these voxels by computing maps of linearity $C_L = \frac{\lambda_1 - \lambda_2}{\lambda_1}$ and planarity $C_P = \frac{\lambda_2 - \lambda_3}{\lambda_1}$ (Westin et al., 1999) from diffusion tensor (DT) fits to the lowest b -value shell, and select only those voxels with $C_L > 0.5$ and $C_P < 0.3$. The SNR of the selected white matter region in the $b=0$ image with $\Delta=0.102$ s is approximately 19 for all subjects.

2.4. Random forest regression

Random forest regression is a machine learning technique which

here aims to learn the mapping between the microstructure parameters and the features of the DW MRI signal. The regressor is a collection of decision trees trained on separate randomly chosen samples from the available training data (Criminisi et al., 2013). Each tree estimates a set of parameters of interest from an input feature vector, and the final estimation is obtained by averaging the estimations from all trees in the forest (Criminisi et al., 2013). In order to do this, an initial training step is required in which every tree in the forest learns a mapping from a vector of signal features to a set of ground-truth microstructure parameters. The trained forest can then be used to estimate the parameters given a previously unseen feature vector. Here, the MC simulations provide training data with known microstructure parameters (specified in the simulation); the simulation outputs DW-MR signals from which we derive the features that form input to the mapping.

Each tree in the forest is trained as follows: at the parent node of the tree, we perform an initial linear regression to find a relationship between the training features and parameters. We then search the feature space for a partition such that having separate regressions on either side of the partition improves the estimation. If such a partition exists the node is split, resulting in a pair of child nodes. We then continue this procedure for the child nodes, and all subsequent nodes, until there is no more improvement in the estimation. Randomness is

typically introduced into the forest in two key ways. The first method, termed bagging, trains each tree on a randomly selected subset of the training data (Breiman, 1996); the second uses a subset of feature space to search for the best partition at each node (Criminisi et al., 2013). In the appendix we provide a more detailed explanation of the method. In the following sections we provide more details on the features we use in training, as well as the training and testing procedures.

2.4.1. Feature extraction

The feature vector used in this study comprises rotationally invariant features derived from diffusion tensor and 4th order spherical harmonic (SH) fits to each shell of data separately. The features that we calculate from the DT are the eigenvectors $\lambda_1, \lambda_2, \lambda_3$, the mean diffusivity MD and the fractional anisotropy FA for each voxel. From the spherical harmonic fits, we calculate the mean and peak ADC, peak dispersion (i.e. the eigenvalues of the hessian matrix at the peak), anisotropy, skewness and kurtosis of the apparent diffusion coefficient profile, as well as simple combinations of the SH coefficients given by

$$I_k = \sum_{i=-k}^k |a_{k,i}|^2 \quad (1)$$

for $k = 0, 2, 4$ where $a_{k,i}$ is the coefficient of SH order k and index i . The skewness of the ADC profile f is $(\frac{\int_S (f(x) - F)^3}{\int_S f^3(x)})^{1/3}$ and the kurtosis is $(\frac{\int_S (f(x) - F)^4}{\int_S f^4(x)})^{1/4}$, where $F = \int_S f(x)$ and S is the unit sphere. This gives 15 features from each shell for a total of 60. When fitting the DT and SH models, we compensate for the additional STE gradients as described in Lundell et al. (2014). Although some of these features are not mutually independent, we include all of them in order to have a comprehensive set to ensure we can find the most informative split criteria.

2.4.2. Training, testing and estimation

In Fig. 2 we show a schematic overview of how training, testing and estimation are done using the random forest regression. We use a widely used and freely available random forest regressor in the sci-kit learn python toolkit (Pedregosa et al., 2011) with 100 trees and a maximum depth of 20. Accuracy generally increases with the number of trees and tree depths. Preliminary experiments found that we obtain diminishing returns in accuracy above these values, but computational complexity increases. We use bagging to inject randomness into the forest. In the Appendix we provide an overview of the random forest regression method we use. More in-depth details of the implementation are available on the scikit-learn website (<http://scikit-learn.org/>). We train separate regressors for the noisy and noise-free data sets resulting in two distinct forests. The forests are trained on 10,000 of the 12,500 feature vectors, with the remaining 2500 previously unseen feature vectors used for testing. This is repeated for 100 randomly generated test and training sets to investigate bias due to the choice of training data. The testing phase allows us to evaluate how well the regressors perform by directly comparing the estimations to the known ground truth values. When estimating from the noise-free test feature vectors we use the random forest trained on noise-free data, whereas we use the random forest trained on the noisy data to estimate from the noisy test feature vectors.

Following the training and testing stages, we use the noisy random forest regressor on the in vivo human data sets. We use the noisy rather than noise-free random forest regressor as it has similar noise characteristics to the actual MR data.

The parameters that we estimate during the random forest regression describe the underlying microstructure properties of the tissue. We combine μ_R and σ_R , which describe size distribution of cells, into a single volume-weighted radius index α (Alexander et al., 2010), so that

the final set of parameters we estimate are f, τ_i, α and d .

2.5. Kärger model fitting

The most commonly used mathematical model to estimate τ_i from diffusion-weighted PGSE or PGSTE data is the Kärger model (Stanisz et al., 1997; Lätt et al., 2009; Nilsson et al., 2010). Therefore, in order to compare our approach to the current state of the art method, we fit a two-compartment Kärger model to the 2,500 noise-free and noisy test data sets and the masked white matter voxels from the in vivo human data sets.

The intracellular compartment, with volume fraction f , is modelled using randomly packed, parallel cylinders characterised by a single radius index α and with an intra-axonal water residence time τ_i . The extracellular space is modelled as a cylindrically symmetric DT with diffusivities, d and d_\perp . d is assumed to be the same in both compartments, and d_\perp is estimated from d and f using the tortuosity model in Szafer et al. (1995).

Prior to model fitting, each measurement is normalised by the mean $b=0$ measurement with the same TM to eliminate T_1 effects. For all data sets the model is fit using Markov Chain Monte Carlo (MCMC) (Gilks et al., 1994) with an offset Gaussian noise model (Panagiotaki et al., 2012) (assuming different noise standard deviations σ for each shell of data, which we estimate a priori) to sample from the posterior distribution over the model parameters. The burn-in phase for the MCMC consists of 10,000 steps, after which we collect 1000 samples at an interval of 100 steps.

2.6. Data analysis

2.6.1. Simulated data

To assess how well the random forest estimates and Kärger model estimates match the known ground truth values in the test set, we use Bland-Altman plots and calculate for each parameter. Bland-Altman plots show the means of the ground truth and estimated values plotted against their differences, allowing us to assess the bias in our estimates. Points on the Bland-Altman plots are colour-coded with colour bars showing the percentage error. Due to better visibility we limit the colour bar to -50 to 50% , with points outside of this range taking the same colour as the two maximum range points. The supplementary material includes additional visualisation of the agreement using scatter plots.

To determine whether the choice of training set used for the random forest regression introduces bias, we calculate the mean and standard deviation of the correlation coefficients for each parameter over all 100 training/test sets.

2.6.2. Healthy subjects

We use the scan-rescan data from the two healthy volunteer subjects to investigate intra-subject reproducibility on both the region of interest (ROI) and voxel level.

For the ROI analysis, we manually define regions in the splenium (CC-S) and genu (CC-G) of the corpus callosum, the left (ALIC-L) and right (ALIC-R) anterior limbs of the internal capsule and the left (CST-L) and right (CST-R) corticospinal tracts on the scan and rescan FA maps for both volunteers. We calculate the mean parameter estimates from the random forest and Kärger model in each ROI for both the scan and rescan data and plot scatter plots of scan versus rescan estimates and Bland-Altman plots for f, τ_i, α and d individually. Bland-Altman plots show the mean of the scan and rescan parameter estimates against their difference; for highly reproducible parameters the points on a Bland-Altman plot should be closely clustered about the zero difference line. We also investigate the regional variability of each model parameter by calculating the coefficient of variation (CoV), defined as the ratio of the standard deviation to the mean, i.e. $\frac{\sigma}{\mu}$ and

represents the variability of a parameter in relation to its mean value over the population. To do this, for each parameter, we pool together the scan and rescan parameter estimates for both volunteers in each region.

We also investigate intra-subject reproducibility on the voxel level. For each subject, we warp the parameter maps estimated from the rescan data into the space of the scan data using the FLIRT tool in FSL (Jenkinson et al., 2002) in order to obtain the correspondence between voxels. To assess the similarities between the scan and rescan, we again use Bland-Altman plots and calculate correlation coefficients for each parameter. The supplementary material includes additional visualisation of the agreement using scatter plots.

2.6.3. MS subjects

To investigate the sensitivity of the random forest and the Kärger model to tissue damage in MS lesions, an expert clinical researcher (NC) marked the lesions on the high resolution T_2 -weighted images. We also mark additional ROIs in the contralateral normal appearing white matter (NAWM) for comparison. The ROI and lesion masks are then registered to the mean $b=0$ image from shell 1 of the diffusion-weighted data using FLIRT (Jenkinson et al., 2002). We use this shell for registration as it has the shortest diffusion time and so provides the best contrast.

In the first subject, in the early stages of MS, two of the marked lesions overlap the white matter mask completely and are investigated further. For the second subject, in the late stage of MS, the lesions are much more widespread, overlapping most of the white matter mask used to select voxels for analysis. However for this subject, our analysis is limited by the availability of contralateral NAWM for comparison and thus we only use one lesion in the CST. We then calculate the mean parameter estimates from both models in the lesions and control ROIs and compare. We then used t-test to find out if there is a statistically significant difference between parameter values in the lesions compared to those in NAWM.

3. Results

3.1. Simulations

Fig. 3 shows Bland-Altman plots of f , α , τ_i and d for both the random forest regressor and the Kärger model for noise-free data generated from cylinder substrates. The data points are colour-coded

according to how close the estimates are to the actual values and the percentage error is shown on the colour bars. The solid black line indicates the mean difference between ground truth and estimated parameters and the dashed lines show the 95% limits of agreement. For the random forest model, points are clustered about the zero difference line, indicating low bias. However, despite the absence of noise the recovery of the parameters is not perfect. In particular, for f , τ_i and α , there is some bias in the estimated values which depends on the ground truth value, for example, large values of τ_i are consistently underestimated. However, the noise-free performance provides a benchmark for the best estimation we can achieve given the data available. Correlation coefficients are also strong for all parameters (f : $R^2=0.88$, τ_i : $R^2=0.95$, α : $R^2=0.82$, d : $R^2=0.99$).

Estimates obtained from the Kärger model are generally less accurate and more biased. In particular, τ_i is consistently underestimated for all ground truth values, and although low diffusivities are well-recovered, the model underestimates high diffusivities. Correlations between ground truth and estimated values of f ($R^2=0.75$) and d ($R^2=0.99$) are good, however the model struggles to recover τ_i ($R^2=0.60$) and α ($R^2=0.11$). It consistently underestimates τ_i , and provides no correlation between ground truth and estimated α .

Fig. 4 shows similar plots, but for simulated data with SNR=20. For the random forest model, the results for all parameters are consistent with those obtained from the noise-free data, although the 95% limits of agreement are slightly wider. Although the mean difference lines are mostly close to zero, again we see that there is some bias in our estimates, which depends on the ground truth values. For f ($R^2=0.70$), we see that the larger volume fractions tend to be underestimated slightly whereas low f are slightly overestimated. This is also the case for τ_i ($R^2=0.70$). Residence times of up to approximately 0.6 s are estimated well, after which the estimates level off. The estimations of axon radius index are weaker than of the other parameters ($R^2=0.48$), but we still see a positive correlation, indicating that the method is still able to distinguish small axons from large axons even in the presence of noise. Diffusivities are again very well estimated ($R^2=0.98$).

The addition of noise weakens the performance of the Kärger model still further. The 95% limits of agreement are much wider than for the random forest model, and again there is more bias in τ_i and d . Similar to the random forest model, there is some bias in f ($R^2=0.65$), although the Kärger model tends to underestimate low values and slightly overestimate higher values. Once again, the Kärger model struggles to estimate high d , despite having a high correlation coefficient

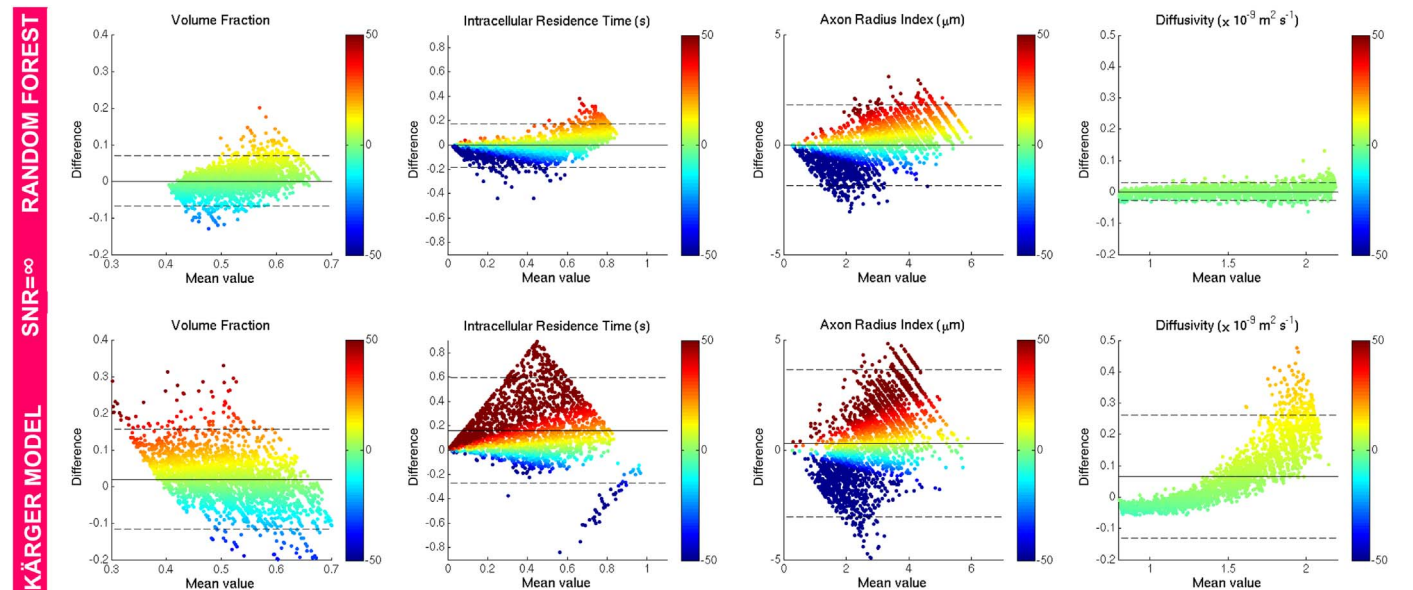


Fig. 3. Bland-Altman plots for f , τ_i , α , d for both the random forest and the Kärger model for noise-free data simulated from cylinder substrates.

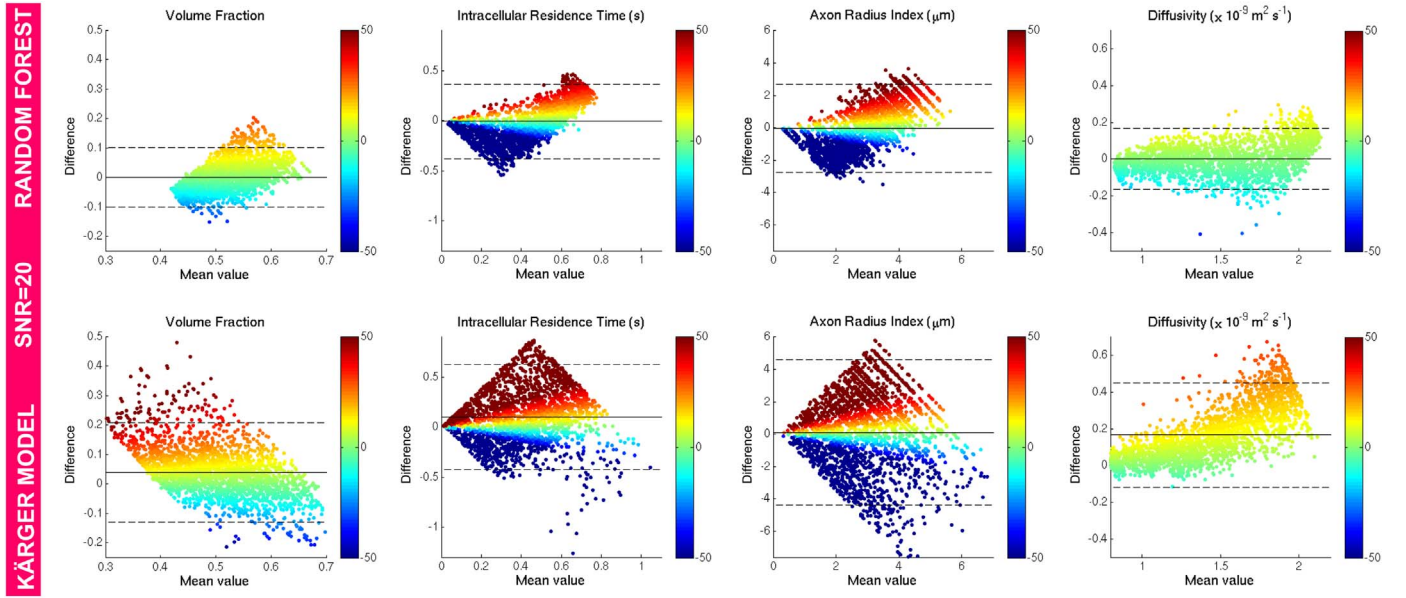


Fig. 4. Bland-Altman plots of f , τ_i , α , d for both the random forest and the Kärger model for data simulated from cylinder substrates with SNR=20.

($R^2=0.96$). The Kärger model is again much worse at estimating τ_i ($R^2=0.40$) and α ($R^2=0.02$). As can be seen from both the Bland-Altman plot and the correlation coefficients, estimations of α are now essentially noise.

Table 2 shows mean correlation coefficients for f , τ_i , α and d averaged over the 100 training/test data sets for both noise-free data and data with SNR=20. The correlations between the ground truth and estimated values are very consistent for all parameters and both noise levels, indicating that while the random forest model needs sufficient coverage of the parameter range, it is not biased by the choice of training data.

3.2. Healthy subjects

Fig. 5(a) shows scan and rescan parameter maps estimated using the random forest regressor across representative axial, coronal and sagittal slices for the first healthy volunteer. An initial visual inspection suggests good agreement between the scan and rescan parameters, with positive correlation coefficients (f : $R^2=0.57$, τ_i : $R^2=0.45$, α : $R^2=0.48$, d : $R^2=0.65$) and similar trends observed across all the major white matter tracts. The values estimated for all parameters are within plausible ranges. Estimates of volume fraction f are in the range 0.44–0.65. The upper bound is slightly lower than expected, but as suggested by the simulation results in Fig. 4 large f tends to be underestimated by the random forest when the data is noisy. However, we still see the expected high-low-high trend in f across the mid-sagittal CC. Estimates of τ_i are consistently in the range 0.4–0.5 s across the major tracts. The scan and rescan maps of d are also highly consistent, with estimations for most voxels in the range $1.4\text{--}1.8 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$.

Table 2


Means and standard deviations of the correlation coefficients between ground truth and estimated values using 100 different training and test data sets for the random forest model, for both SNR = ∞ and SNR=20.

Parameter	SNR = ∞		SNR=20	
	$\mu(R^2)$	$\sigma(R^2)$	$\mu(R^2)$	$\sigma(R^2)$
f	0.882	0.005	0.708	0.010
τ_i	0.946	0.002	0.694	0.009
α	0.817	0.006	0.491	0.015
d	0.999	0.000	0.978	0.001

Fig. 5(b) shows the scan and rescan parameter maps estimated using the Kärger model across the same slices. Estimates of volume fraction f are in the range 0.2–0.8. We see the high-low-high trend in f more distinctly in the Kärger model data than in the random forest data although estimated volume fractions in the midbody of the CC (0.25–0.3) are lower than expected, especially as the white matter masks were generated using a linearity threshold of 0.5. The scan and rescan maps of f are generally in good agreement though with $R^2=0.57$, which is the same correlation coefficient as that of the random forest estimate. Maps of τ_i are noisier and show less consistency between the scan and rescan data with $R^2=0.28$, lower than that of the random forest estimate. Moreover, the location of these areas varies from scan to rescan. For example there are several voxels in the splenium of the mid-sagittal CC which have low τ_i in the scan data but much higher τ_i in the rescan data (shown by the arrows in Fig. 5(b)). Both the scan and rescan maps of α are very noisy, with an extremely low correlation coefficient of $R^2=0.06$ and no discernible patterns in the data. Estimates of d however do appear reproducible ($R^2=0.66$), and in the range of $0.8\text{--}1.4 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$.

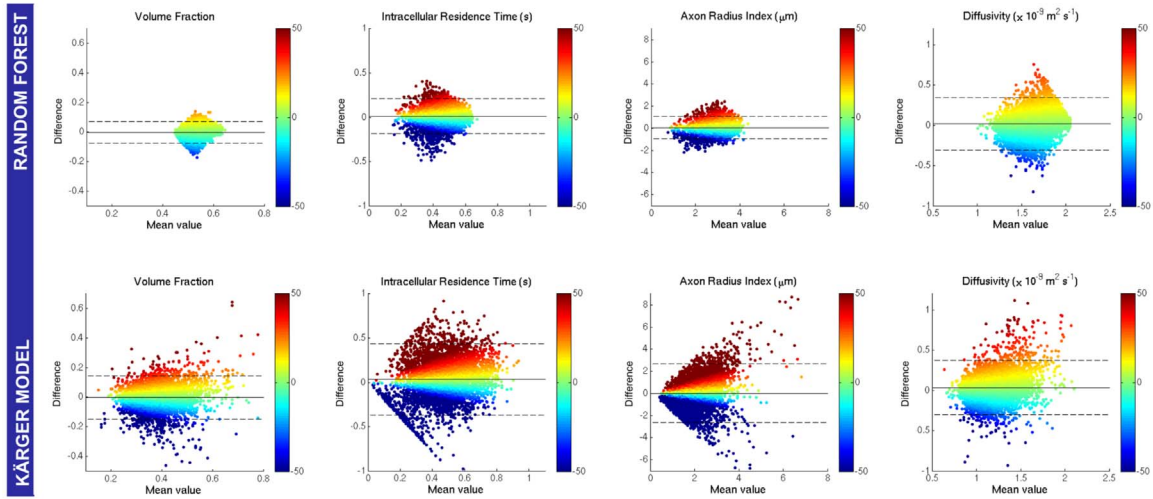
To provide a more quantitative analysis, we warp the parameters maps derived from the rescan data into the space of the scan data using the FLIRT tool in FSL (Jenkinson et al., 2002). Bland-Altman plots of scan versus rescan parameter estimates of f , τ_i , α and d across corresponding voxels for both the random forest and Kärger model are displayed in Fig. 5(c). The data points are colour-coded based on the difference between scan and rescan parameter estimates, with colour indicating percentage error. The colour scale is the same for the random forest and Kärger model for each parameter. Neither model shows significant bias for any parameter, and all points are clustered about the zero difference line. For f , τ_i and α , the random forest model has better reproducibility as suggested by the narrower 95% limit lines, but for d these limits are roughly the same for both methods. The correlation coefficients for f (random forest: $R^2=0.57$, Kärger model: $R^2=0.57$), d (random forest: $R^2=0.65$, Kärger model: $R^2=0.66$), τ_i (random forest: $R^2=0.45$, Kärger model: $R^2=0.28$) and α (random forest: $R^2=0.48$, Kärger model: $R^2=0.06$) also support these findings. Scatter plots showing these correlations for both subjects are included in the supplementary material.

Figs. 6(a)–(c) show equivalent parameter maps and scatter plots for volunteer 2. As for subject 1, the scan and rescan maps for the random forest model show good visual similarities for all parameters. We see the high-low-high trend in f and the low-high-low trend in α across the




(a) Parameter maps of f , τ_i , α and d , estimated using the random forest model, across representative axial, coronal and sagittal slices of both scan and rescan data for volunteer 1.

(b) Parameter maps of f , τ_i , α and d , estimated using the Kärger model, across representative axial, coronal and sagittal slices of both scan and rescan data for volunteer 1. Arrows indicate regions in which there are differences in the estimates of τ_i between the scan and rescan data.

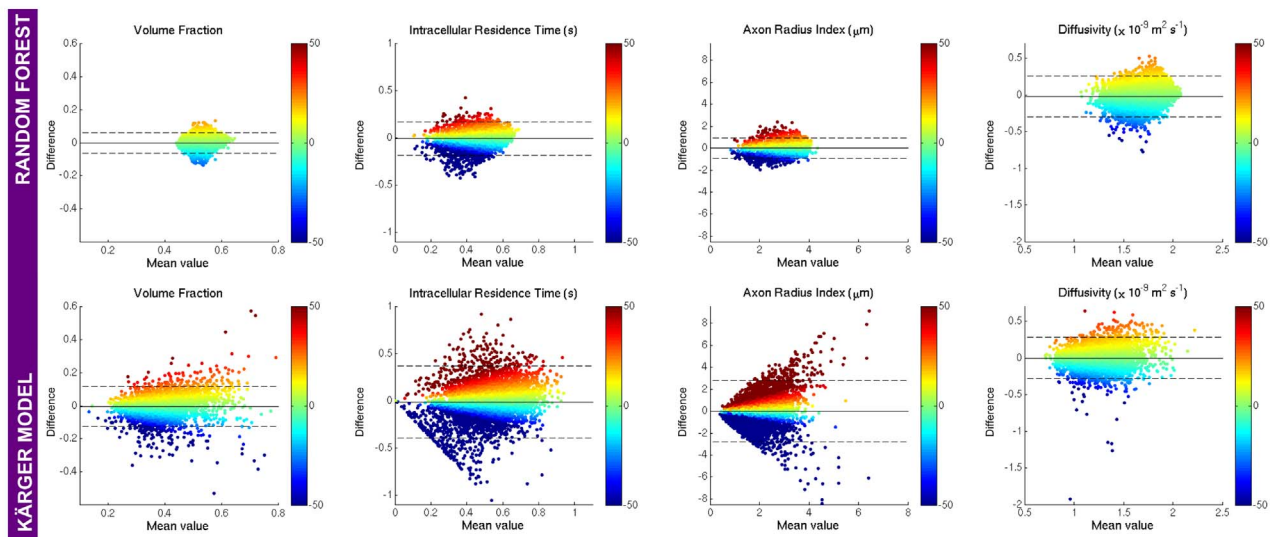


(c) Bland-Altman plots showing voxelwise scan-rescan reproducibility for f , τ_i , α and d estimated using both the random forest and Kärger model results for volunteer 1. The range of f estimated using the random forest method is narrower than the values estimated by the Kärger model, but it is informed by the training data which is representative of white matter tissue.

Fig. 5. Scan-rescan data for volunteer 1.



(a) Parameter maps of f , τ_i , α and d , estimated using the random forest model, across representative axial, coronal and sagittal slices of both scan and rescan data for volunteer 2. (b) Parameter maps of f , τ_i , α and d , estimated using the Kärger model, across representative axial, coronal and sagittal slices of both scan and rescan data for volunteer 2.



(c) Bland-Altman plots showing voxelwise scan-rescan reproducibility for f , τ_i , α and d estimated using both the random forest and the Kärger model for volunteer 2.

Fig. 6. Scan-rescan data for volunteer 2.

mid-sagittal CC and estimates of τ_i and d are consistent between scans as well as between the two subjects. For the Kärger model, reproducibility is again reasonable for f and d but the maps of τ_i and α are noisy. Fig. 6(c) shows voxelwise Bland-Altman plots of scan versus rescans parameters. As for the first volunteer, these plots show that there is no bias between scan and rescans estimates for either method, although again the random forest model has better reproducibility for f , τ_i and α as indicated by the narrower 95% limit of agreement lines. The correlation coefficients for the random forest and Kärger model are very similar for f (random forest: $R^2=0.73$, Kärger model: $R^2=0.73$) and d (random forest: $R^2=0.72$, Kärger model: $R^2=0.75$), but the reproducibility of residence time (random forest: $R^2=0.63$, Kärger model: $R^2=0.35$) and axon radius metrics (random forest: $R^2=0.56$, Kärger model: $R^2=0.02$) is much higher for the random forest than the Kärger model.

Following on from the voxelwise analysis of scan-rescan reproducibility, we analyse the reproducibility of the indices across regions of interest (ROIs). We manually define ROIs in the splenium (CC-S) and genu (CC-G) of the corpus callosum, the left (ALIC-L) and right (ALIC-R) anterior limbs of the internal capsule and the left (CST-L) and right (CST-R) corticospinal tracts on the scan and rescans data for both volunteers. Fig. 7 shows an example ROI mask overlaid on the FA image of the scan data for volunteer 2. We calculate the mean parameter estimates from the random forest and Kärger model in each ROI for all data sets. The top row of Fig. 7 shows scatter plots of scan versus rescans estimates of f , τ_i , α and d . Results from the Kärger model (blue) and the random forest (pink) are shown on the same plots. The random forest shows high reproducibility for all parameters at the ROI level. Estimates of d from the Kärger model show good reproducibility, matching that of the random forest. However, f , τ_i and in particular α show lower reproducibility. The bottom row of Fig. 7 shows a Bland-Altman plot, which plots the mean of the scan and rescans parameter estimates against their difference. Dashed lines indicate the 95% limits for each method. The results of the Bland-Altman analysis confirm the results of the correlation plots. For d , the data points from both the random forest and Kärger model show similar spread about the zero line, indicating that the methods have equivalent reproducibility. However, for the other three parameters the data points for the random forest are more closely clustered around the zero line than the data points for the Kärger model, indicating that the scan and rescans parameter estimates from the random forest model show better agreement.

Finally, we pool the scan and rescans parameter estimates for both subjects across the six regions of interest and calculate the mean μ , standard deviation σ and coefficient of variation (CoV) of each parameter in each region (as described in the Methods section). Tables 3 and 4 show the results for the random forest and Kärger model respectively. CoVs are lower for all parameters when estimated using the random forest as opposed to the Kärger model. The two methods perform similarly for d ; both methods result in similar standard deviations, but the CoVs for the Kärger model appear slightly higher than for the random forest due to the slightly lower mean estimates of d . However, for f , τ_i and α , CoVs calculated from random forest parameter estimates are substantially lower than for the Kärger model.

3.3. MS patients

In Fig. 8, we present the parameter maps of the first MS subject (early stage MS). First from the left is a $b=0$ image with the overlaid lesion mask (red areas in squares) and respective normal appearing white matter tissue (NAWM) (green areas in circles). The remaining four columns are parameter maps from the same slice estimated using the random forest model (top row) and Kärger model (bottom row). An initial visual inspection of parameter maps, suggests that neither the random forest nor the Kärger model detects any obvious parameter differences in the lesion in genu (top area in a square), compared to the respective NAWM tissue (top area in a circle). For the lesion in splenium though (bottom area in a square), there are differences compared to the NAWM (bottom area in a circle): τ_i estimated using the random forest is reduced relative to the NAWM area and the volume fraction estimated by Kärger model is lower compared to the NAWM. When comparing random forest and Kärger model parameter estimates, random forest provide higher values than Kärger model for all four parameter maps over all white matter voxels.

Equivalent parameter maps for the second subject with late stage MS are shown in Fig. 9. In this subject the lesions are much more widespread, overlapping most of the white matter mask used to select voxels for analysis, including the genu and the corticospinal tracts. An initial visual inspection of the whole white matter mask shows more dramatic changes in the parameter maps for both the random forest and Kärger model compared to the first MS subject in early stage MS. For the random forest model, there is a reduction in f and τ_i and an increase in α in the lesion area (red area in a square) compared to the

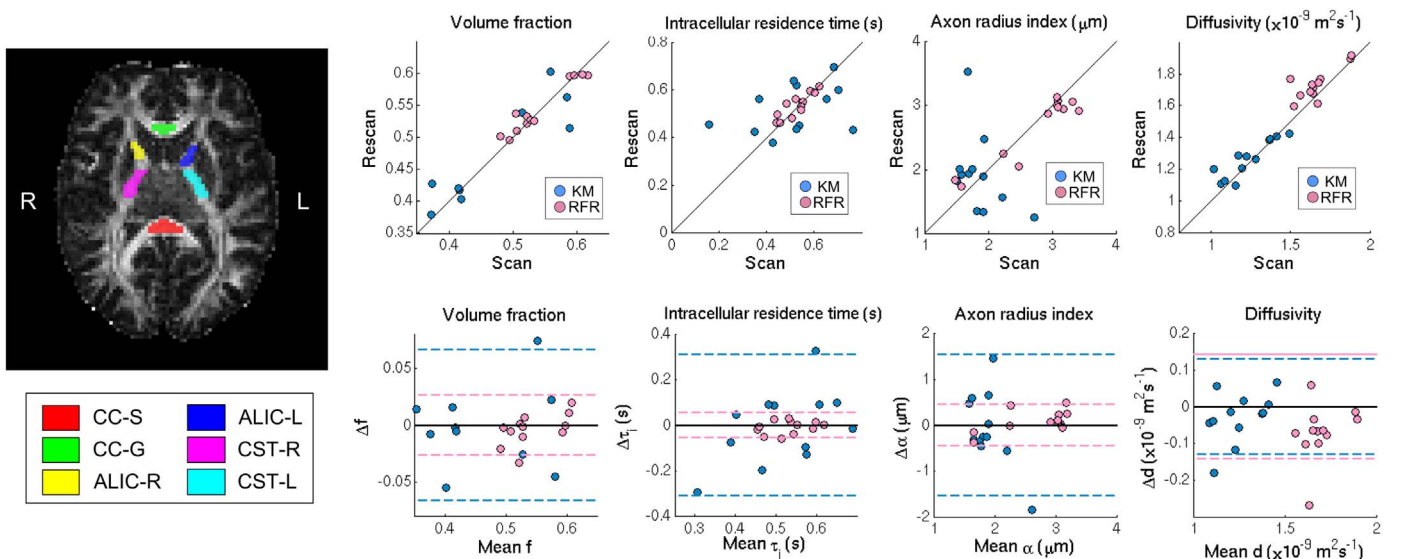


Fig. 7. The left hand panel shows example regions of interests used in the analysis overlaid on the FA map of subject 1. The top row shows the correlation of scan-rescan estimates of all parameters using both the random forest regression (RFR) and Kärger model (KM). The bottom row shows the equivalent Bland-Altman plots.

Table 3


Mean, standard deviation and coefficient of variation of all parameters estimated using the random forest model in all six regions of interest.

ROI	f		τ_i (s)		α (μm)		d ($\times 10^{-9} \text{ m}^2 \text{ s}^{-1}$)	
	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV
CC-S	0.61(0.02)	3.9	0.57(0.05)	9.2	1.65(0.58)	35.1	1.67(0.19)	11.3
CC-G	0.59(0.02)	3.2	0.60(0.04)	7.4	2.26(0.70)	31.1	1.89(0.09)	4.6
ALIC-L	0.51(0.04)	7.0	0.48(0.10)	20.1	3.05(0.43)	14.0	1.61(0.18)	11.2
ALIC-R	0.50(0.04)	8.3	0.46(0.11)	24.2	3.14(0.52)	16.6	1.63(0.18)	10.8
CST-L	0.53(0.04)	7.1	0.54(0.08)	14.4	3.06(0.32)	10.3	1.70(0.15)	8.9
CST-R	0.53(0.04)	7.1	0.52(0.09)	16.7	3.09(0.35)	11.4	1.66(0.14)	8.7

Table 4

Mean, standard deviation and coefficient of variation of all parameters estimated using the Kärger model in all six regions of interest.

ROI	f		τ_i (s)		α (μm)		d ($\times 10^{-9} \text{ m}^2 \text{ s}^{-1}$)	
	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV	$\mu(\pm\sigma)$	CoV
CC-S	0.56(0.09)	15.4	0.37(0.28)	75.9	1.79(1.52)	85.2	1.16(0.19)	16.0
CC-G	0.55(0.10)	17.5	0.59(0.24)	41.9	2.32(1.86)	80.1	1.41(0.14)	9.9
ALIC-L	0.37(0.06)	17.1	0.44(0.15)	33.9	1.77(0.96)	54.4	1.14(0.15)	13.6
ALIC-R	0.33(0.07)	20.9	0.45(0.15)	34.1	1.78(1.04)	58.3	1.12(0.13)	12.0
CST-L	0.42(0.05)	11.9	0.67(0.16)	23.8	1.78(0.95)	53.2	1.40(0.15)	11.0
CST-R	0.41(0.06)	14.1	0.59(0.16)	27.0	1.84(0.92)	50.3	1.26(0.13)	10.3

**Fig. 8.** Mean $b=0$ image and parameter maps showing f , τ_i , α and d estimated using the random forest and estimated using the Kärger model for the first MS subject. There are two lesions (red areas in squares), one in genu and one in splenium, and the mask for these is overlaid on the mean $b=0$ image. For each of the lesions, its respective contralateral normal appearing white matter (NAWM) mask is shown as well (green areas in circles).

contralateral NAWM (green area in a circle). For the Kärger model, f and τ_i are lower in the lesion areas compared to the contralateral NAWM, but there also appears to be increase in d .

To provide a more quantitative analysis, we calculate the mean parameter estimates from both models in the lesions we marked in Figs. 7 and 8 (red areas in squares), and their respective contralateral NAWM ROIs (green areas in circles), as shown in Table 5. In the lesion in the genu of subject 1, the diffusivity is higher than that in NAWM. The difference between the other model parameters is not statistically significant. For the lesion in the splenium of subject 1, the random forest estimates a lower τ_i of 0.19 s in the lesion compared to 0.33 s in the contralateral NAWM. The random forest also estimates a higher α in the lesion compared to the NAWM. The Kärger model does not show any differences in τ_i or α , but it does suggest a reduction in axonal volume fraction in the lesion, with f dropping from 0.35 in the NAWM

to 0.26.

For the lesion in the CST of subject 2, the quantitative analysis confirms the initial visual inspection. Using the random forest model, we find that f drops from 0.52 to 0.47 and τ_i is reduced from 0.56 s in the NAWM to 0.13 s in the lesion. There is also an increase in α from 3.3 μm to 4.3 μm , as well as an increase in d , although for d the change is the same order of magnitude as the standard deviation across the ROI. Estimates from the Kärger model show a larger decrease in f from 0.41 in NAWM to 0.18 in the lesion. τ_i also decreases from 0.71 s in the NAWM to 0.33 s in the lesion, although the standard deviation is much higher when using the Kärger model compared to the random forest. The intrinsic diffusivity increases from $1.50 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$ to $1.84 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$. There is also an increase in the mean value of α in the lesion compared to the NAWM, but the standard deviation is much larger than the apparent change. Lesion parameter values for which




Fig. 9. Mean $b=0$ image and parameter maps showing f , τ_i , α and d estimated using the random forest and estimated using the Kärger model for the second MS subject. Masks for the lesion and NAWM are overlaid on the mean $b=0$ image. The lesion is in CST and is coloured in red (in a square), and the NAWM is also in CST but contralateral to the lesion and is coloured in green (in a circle).

Table 5

Mean and standard deviation of random forest regression (RFR) and Kärger model (KM) parameters in MS lesions compared to contralateral NAWM in both MS subjects. Lesion parameter values which are statistically different ($p < 0.02$) to the values in NAWM are marked with stars.

		MS subject 1				MS subject 2	
		CC: Genu		CC: Splenium		CST	
		Lesion	NAWM	Lesion	NAWM	Lesion	NAWM
RFR	f	0.49 (0.03)	0.50 (0.04)	0.49 (0.02)	0.52 (0.05)	0.47 (0.01) *	0.52 (0.04)
	τ_i	0.48 (0.07)	0.45 (0.09)	0.19 (0.11) *	0.33 (0.12)	0.13 (0.09) *	0.56 (0.05)
	α	3.24 (0.39)	3.16 (0.41)	3.25 (0.77)	2.60 (0.73)	4.30 (0.39) *	3.30 (0.23)
	d	1.74 (0.13) *	1.51 (0.14)	1.69 (0.22)	1.76 (0.26)	1.96 (0.07) *	1.91 (0.06)
KM	f	0.35 (0.04)	0.33 (0.06)	0.26 (0.12)	0.35 (0.14)	0.18 (0.04) *	0.41 (0.05)
	τ_i	0.48 (0.09) *	0.36 (0.07)	0.22 (0.15)	0.22 (0.18)	0.33 (0.22) *	0.71 (0.12)
	α	1.82 (1.10)	2.28 (1.06)	1.74 (1.22)	1.74 (1.22)	2.39 (1.08)	1.86 (1.19)
	d	1.36 (0.11) *	1.06 (0.13)	1.24 (0.25)	1.22 (0.23)	1.84 (0.17) *	1.50 (0.08)

there is a statistically significant difference ($p < 0.02$) to the values in NAWM are marked with stars.

4. Discussion

In this study we have introduced a computational framework to learn a computational model that maps features derived from DW MR signals to microstructure parameters. This approach supports the estimation of microstructure parameters that have previously proven elusive due to lack of accurate mathematical models relating them to the MR signal. The key example we focus on here is the residence time of water between intra and extracellular spaces and we used random forest regression to learn the mapping. Simulations show that the parameter correlations from the new computational model approach

are much stronger than those obtained using the Kärger model, particularly for residence time τ_i and axon diameter index α . The trained random forest also estimates sensible microstructure parameters from in vivo human data, even for parameters such as α and τ_i which only weakly influence the DW MR signals available from human scanners (Drobnjak et al., 2016; Nilsson et al., 2013). From scan-rescan data of healthy subjects, we show that results from the random forest have better reproducibility than the Kärger model on both the voxel and ROI level. Finally, data acquired from two MS patients demonstrate the clinical potential of estimates of exchange-related parameters obtained this way. In both subjects, we observe reductions in f and τ_i and increases in α in lesions compared to NAWM when using the random forest model, which is consistent with expectations from the pathology of MS lesions. In particular we expect the axon diameter index to be higher in lesions than NAWM, because small axons are more vulnerable to MS pathology (DeLuca et al., 2006; De Luca et al., 2004; Huang et al., 2016); axonal loss corresponds to lower f ; and myelin damage can be expected to reduce τ_i .

4.1. Simulations

Our simulation results, which cover a wide range of ground truth parameters, including both situations in which the system is close to well-mixed (short τ_i ; long Δ) and not well-mixed (long τ_i), show that the Kärger model performs worse than the random forest approach as it is less able to recover the underlying ground truth parameters accurately. Kärger model consistently underestimates τ_i , even short τ_i , i.e. less than 0.1 s, where for the longest diffusion times of 0.406 s and 0.412 s in the protocol, the system approximates the well-mixed conditions for which the Kärger model is valid. However, the comparison is not totally fair, as the random forest is trained directly on simulated data, while Kärger model, although used to optimise the protocol for sensitivity to τ_i , is not. Nevertheless, the comparison is worthwhile, because, the study still provides a useful insight into the limitations of the Kärger model's assumptions, i.e. that the system is well-mixed, in situations where they are known to be violated, as in white matter. We show that for realistic white matter models Kärger model provides inaccurate estimates, which could be argued as expected as Kärger model is not meant to be valid under those assumptions. Nevertheless, it is worthwhile to do the comparison

because Kärger model provides the current standard model of exchange for PGSE and PGSTE data, even when diffusion is restricted, and the evaluation of its performance under these more realistic conditions is important.

To investigate the performance of the Kärger model further, we also consider the case in which the Kärger model is fitted to test data generated from the Kärger model, both with and without noise i.e. the situation in which the Kärger model should explain the underlying data exactly (results not shown here). For noise-free data, the Kärger model estimates f ($R^2=0.98$), τ_i ($R^2=0.76$) and d ($R^2=1$) well. Estimates of τ_i are still not exact as, even at an maximum diffusion time of 0.412 s, the signals generated from the Kärger model for any τ_i greater than ≈ 0.6 s are virtually indistinguishable. Estimates of α are still poor ($R^2=0.21$) but as discussed earlier the protocol was not optimised for maximum sensitivity to this parameter. For noisy data with SNR=20, estimations of f ($R^2=0.87$) and d ($R^2=0.96$) are still strong, but results for τ_i ($R^2=0.35$) and α ($R^2=0.09$) are poor. This implies that estimates of τ_i from the Kärger model are extremely sensitive to noise, and thus the model is unlikely to be able to provide meaningful estimates of intracellular water residence time from data obtained under realistic acquisition conditions.

4.2. Healthy subjects

The scan-rescan experiments in healthy subjects show that the random forest has good reproducibility compared to the Kärger model. However for all parameters the correlation coefficients are much lower than the very high reproducibility of DT-derived parameters such as FA, which typically has a CoV of 0.8–3.0% (Vollmar et al., 2010) or NODDI parameters (Tariq et al., 2013) which achieve CoV of around 5–6%. Difficulties in registering the scan and rescan volumes due to different numbers of slices in the data sets contribute, but mainly this arises because sensitivity of the signal to τ_i and α is weak. However, we also note that the CoVs for τ_i estimated using the random forest model with the current protocol compare favourably to values published in a recent reproducibility study of AXR imaging (Lampinen et al., 2014). Lampinen et al estimate the CoV of AXR in the CC, ALIC and CST and report values of 58%, 41% and 45% respectively. Their study differs from ours in that they combine AXR estimates from the left and right ALICs and CSTs and from the genu and splenium of the CC; if we combine our data in a similar manner, the CoVs for τ_i would be 8.9%, 22.3% and 15.6%. However, they also have a larger number of subjects and note that at least in the CST approximately half of the variance can be attributed to inter-subject differences.

Random forest estimates for the residence time and the axon diameter are closer to the reported values in literature than those of the Kärger model. There are numerous voxels in which Kärger model estimates τ_i to be between 0.01 and 0.06 s. These values are an order of magnitude smaller than the values reported in literature (Finkelstein, 1976; Quirk et al., 2003). Estimates of d using Kärger model are in the range $0.8\text{--}1.4 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$ which is lower than what we would typically expect in the in-vivo white matter tissue (Kaden et al., 2016). The ROI analysis shows that for f , τ_i and α , CoVs calculated from random forest parameter estimates are substantially lower than for the Kärger model, suggesting that the random forest may provide more consistent estimations of microstructure indices across a population. We also note that for the random forest method, the values of f that are estimated fall in a narrower range than those estimated using the Kärger model, which may artificially narrow the 95% limits slightly. This is due to the range of volume fractions present in the training set ($f=0.4\text{--}0.7$) which was chosen to be representative of white matter tissue. Also, the correlation coefficients are much higher for the second volunteer than the first volunteer, which may be due to a better registration between the data sets.

4.3. MS subjects

In MS subjects, we observe reductions in f and τ_i and increases in α in lesions compared to NAWM when using the random forest model. This is intuitive as the breakdown of myelin in the lesion is likely to make the residence time shorter compared to the NAWM, the damage of the tissue is likely to cause cell loss and the small axons are found to be preferentially damaged (DeLuca et al., 2006; De Luca et al., 2004; Huang et al., 2016). Analysing further and comparing the NAWM in MS subjects to the same ROIs in healthy subjects (presented in Table 3, CC:S, CC:G, CST-R) we observe that the trend continues, i.e. we observe reductions in f and τ_i and increases in α in patients' NAWM compared to healthy subjects. It has been shown previously that NAWM in MS subjects is somewhat damaged compared to the normal tissue (Filippi et al., 2012) and our results support this suggestion. These are just preliminary results because the number of both MS and healthy subjects is quite small. However, the trends are plausible and the estimated parameters sensible. We also found that, for one of the lesions in the first MS subject, almost no difference between the parameters in the lesion and the NAWM tissue was found. This could be because the lesion was relatively new compared to the other lesion and the damage has not yet fully developed. However, using the Kärger model for that lesion produces counterintuitive results as we found that the residence time was shorter in the NAWM than in the lesion itself. This is the opposite trend to what we expect, as damage to the myelin sheath is thought to increase permeability and thus reduce the residence time. We did not observe this when using the random forest model, and therefore believe that the difference is more likely due to noisy estimation of τ_i using the Kärger model. In future, we plan to translate the technique for estimating residence time τ_i we introduced here. We show some promising new results, however a larger study of MS patients is necessary in order to investigate whether the parameters of the random forest model, in particular τ_i , could act as suitable biomarkers for detecting and tracking changes in MS pathology.

4.4. Limitations and further work

One of the limitations of the work we present here is that the training uses a very simple model of white matter tissue. Specifically, the model uses long straight circular cylinders that mimic axon bundles and does not account for myelin water, curvature of axons, dispersion or crossing fibres. However, the general idea we present here, of computational models based on simulations and machine learning extends naturally to much more complex models. Future work will investigate incorporation of such effects, although these require further development of the simulation system we use to construct training data.

Another limitation of the current study is that the machine learning procedure is tested on the same type of substrates that it was trained on. We note however that although the model of the substrate is the same, the model parameters between the training and the test data are different so the results demonstrate generalization to some extent. Future work will explore a meaningful evaluation using more diverse test data with additional effects (e.g. undulation, dispersion or a histology-based-mash). However, this would require substantial further development of our simulation system and is outside of the scope of this paper. One important advance for the future is to use the random forest regression to obtain measures of uncertainty (as in Tanno et al., 2016) to highlight situations where the signals are unfamiliar so the model produces unreliable results.

The success of random forest regression is dependent on the sensitivity of the diffusion MR signal to the parameters of interest. This can be seen in the simulation results where for τ_i greater than the maximum diffusion time, the difference in signals (and thus in the features) due to differences in τ_i is within the standard deviation of the noise, and so cannot be differentiated. The estimations of axon radius

index are weaker than that of the other parameters ($R^2=0.48$), with errors going as high as 100% in some cases as can be seen in Figs. 3 and 4, which is simply due to a low sensitivity of the signal to axon diameter for the gradient strength used (Drobnjak et al., 2016). Nevertheless, for both residence time and axon diameter we still see a positive correlation between the estimates and the ground truth, indicating that the method is able to distinguish small axons from large axons and very permeable from impermeable axons even in the presence of noise. The sensitivity of the signal and hence the reproducibility of our parameters can potentially be improved by optimising the acquisition protocol further, e.g. by improving the angular resolution (subject to time constraints) or increasing the maximum diffusion time or gradient strength to provide better contrast (also subject to SNR constraints) or using the random forest model in the protocol optimisation instead of the Kärger model.

The results we obtained in-vivo are inherently difficult to validate. Accurate estimates of τ_i are not obtainable via histology, however, an existing study of intra- to extra-axonal water exchange across the whole in vivo rat brain using relaxometry and contrast agents suggests a mean τ_i of approximately 0.55 s (Quirk et al., 2003) aligning with our estimates. Furthermore, studies on sphingomyelin membranes (which are found in axonal membranes) estimate permeabilities of $\approx 1 \times 10^{-6} \text{ ms}^{-1}$ (Finkelstein, 1976), which correspond to residence times of 0.3–0.6 s for axons with radii of 0.5–1 μm . The maps of α are slightly noisier, but we observe the characteristic low-high-low trend across the CC (Aboitiz et al., 1992) in both scan and rescan maps. Although it is possible that the random forest distinguishes the small axons in the genu and splenium from the larger ones in the midbody and cortico-spinal tract, other effects not accounted for in the model (dispersion, undulation, etc.) mean that it does not accurately recover actual mean diameters. The scan and rescan maps of d are highly consistent. The estimations for most voxels lie in the range $1.4\text{--}1.8 \times 10^{-9} \text{ m}^2 \text{ s}^{-1}$, which is plausible for white matter tissue.

We used the same intrinsic diffusivity and the same MR parameters such as T_1 and T_2 inside and outside the axons. This assumption may be inaccurate, particularly in pathology, and departures would indeed disrupt the parameter estimates. The method of using machine learning to generate computational models extends naturally to situations where these parameters differ in different water pools. However, accurate estimations of such parameters would demand richer data sets with higher b-values - the computational models do not affect fundamental limitations on identifiability, as discussed for example in Jelescu et al. (2016).

One possible way to improve performance of the Kärger model in recovering ground truth parameters in simulation is to constrain its parameters to the range used in the simulation. We note that the random forest model is not constrained in this way either, as the linear models at each leaf node can extrapolate, although in practice the random forest does limit parameter estimates mostly within that range. Even so, constraining the Kärger model to the same range is unlikely to improve performance. For example, in the noise free case, 86% of the Kärger model's estimates of intracellular exchange time and 99% of the Kärger models estimates of axon radius index are within the same range as the random forest training values anyway. The intracellular exchange rate is constrained to be positive in the Kärger model, and of

the 14% of estimates that do lie outside of the random forest training set range most are simply slightly closer to zero. For these values, constraining Kärger model further will result in them converging to the lower limit, which will not improve the correlation with the ground truth values.

The random forest models that we use are non-linear although they do use a linear model at the leaf nodes. Linear models at the leaf nodes prove to be a good trade off between the segregation of the training set and the complexity of the models at each leaf node, however the future work may explore alternative leaf node models such as constant quadratic etc. Criminisi et al. (2013).

4.5. Conclusions

We propose to use machine learning to construct computational models that relate the MR signal to microstructure parameters. We demonstrate the idea by testing the ability to estimate residence time τ_i using clinical scanner settings. This can potentially be important as the axonal membrane permeability, which is associated with the residence time, is hypothesised to be correlated with the condition of the insulating myelin sheaths around the axonal fibre (Nilsson et al., 2013). Numerous white matter pathologies of the human central nervous system, such as MS, spinal cord injury and leukodystrophies, are characterised by damage to the myelin; thus sensitive and specific biomarkers, such as the residence time τ_i that we estimate here, could improve the diagnosis of and treatment for these conditions.

The computational modelling approach we use here opens doors to estimating a wide range of other parameters for which mathematical models are intractable: undulation; properties of the extra-cellular space; etc. The approach, as used here, also neatly avoids the need for any analytical model of tortuous extra-cellular diffusion for which only approximations are available (Burcaw et al., 2015). Given that these compartment models tend to overestimate key microstructure parameters such as the axon diameter (Alexander et al., 2010), including the effects of these intractable parameters not only gives us clinically useful parameters, but may also improve the estimation of other microstructure indices. Although the mapping we learn is specifically for randomly packed, parallel, non-abutting cylinders and a STE imaging sequence, the approach easily extends to other tissue configurations and pulse sequences. In future we plan to incorporate models of fibre dispersion into our MC simulations to extend our technique to dispersed white matter fibre regions as well as grey matter. We also plan to investigate more specialised pulse sequences, such as the AXR sequence (Nilsson et al., 2013) or optimised generalised gradients (Drobnjak et al., 2010, 2011), which may improve sensitivity as well as allowing us to compare other analytic models, or other kinds of machine learning technique, e.g. deep learning.

Acknowledgments

This work was supported by EPSRC (EP/G007748, EP/I027084/01, EP/L022680/1, EP/M020533/1), EPSRC-funded UCL Centre for Doctoral Training in Medical Imaging (EP/L016478/1), Spinal Research, UK MS Society and NIHR Biomedical Research Centre at University College London Hospitals.

A.1. Random forest regression

The random forest regressor used here was implemented using the open source scikit-learn python toolkit (Pedregosa et al., 2011), a widely used free software machine learning library for Python. The random forest regressor method has two stages. Firstly, during the training stage the forest learns a mapping between the features of the data (in our case the metrics derived from the diffusion MRI signals) and the parameters to be




Fig. A10. Random Forest Regressor. The regressor used here is formed of 100 trees and the final estimations of the tissue parameters are obtained by averaging over the estimations of each individual decision tree.

estimated (tissue microstructure parameters). Secondly, during the testing stage, the forest uses the mapping to estimate tissue parameters from previously unseen data.

A random forest regressor is formed of a collection of regression decision trees and its output is a weighted average of the estimate from each individual decision tree (see Fig. A10). A decision tree is formed of internal nodes (represented by grey circles in Fig. A11), terminal nodes (the blue rectangles) and edges which connect all the nodes in a hierarchical fashion. Each internal node has one incoming and two outgoing edges and stores a test (split function) that it applies to the incoming data, which is a multi-dimensional vector that, here, contains rotationally invariant features of the MRI signal corresponding to a substrate. After the test, the data is sent along one of the outgoing edges depending on the result of the test. The

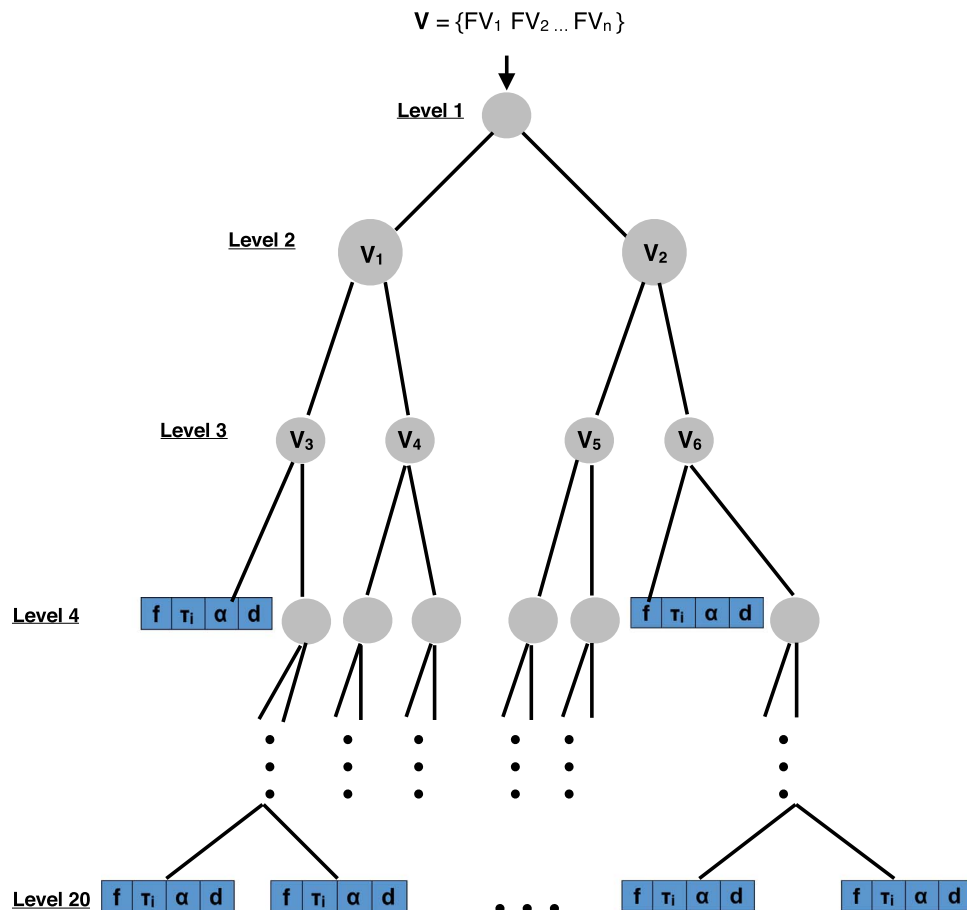


Fig. A11. Regression decision tree: the ensemble model representing the training stage. During the training, each tree receives a random subset of labelled feature vectors which are used to optimise the parameters of the tree in order to maximise the information gain and increase the confidence of the estimations. The terminal nodes store the predictor that relates the incoming data to the estimates of the microstructure parameters of interest: f , τ_i , α and d .

terminal nodes store the predictor that relates the incoming data to the final answer, which here is the estimate of the microstructure parameters of interest.

The training phase decides on three things: the internal structure of the tree, i.e. what internal and leaf nodes exist; the test at each internal node; and the regression function at each leaf node. It does that through a greedy splitting process: it starts with all the input training data V at the root node and constructs a linear mapping from V to the corresponding labels. Here V is a set of feature vectors $V = \{FV_1, FV_2, \dots, FV_n\}$ each derived from the diffusion MRI signals of n different substrates and the labels are the corresponding ground truth microstructure parameters for each substrate. The training phase then seeks a test that divides the training data into two subsets V_1 and V_2 , each with an independent linear mapping, in such a way as to maximise the information gain over the single mapping at the root node. If no pair of mappings increases the information over the single one, the root node remains a leaf node. Otherwise it splits into two siblings containing the two subsets and the process continues recursively until no information gain can be found. At the end of the training stage, each tree encodes a piece-wise linear mapping between the feature vectors (MRI signal features) and the parameters to be estimated (the four microstructure parameters f, τ_i, a, d).

For the testing stage, previously unseen feature vectors are fed into the random forest. Using the mapping learned during the training stage, the forest returns estimated tissue parameters for each of the previously unseen substrate feature vectors. Random forests repeat the same training process for each decision tree but introduce randomness in different ways. Here we use bagging, which trains each decision tree on a different, randomly selected subset of data. Randomisation is introduced into the training process as it has been shown to improve the performance of the forest and its robustness to noise (Criminisi et al., 2013).

A more in-depth discussion together with details of the implementation is available in scikit online documentation in <http://scikit-learn.org> and papers (Criminisi et al., 2013) and (Pedregosa et al., 2011).

Appendix B. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.neuroimage.2017.02.013>.

References

- Aboitiz, F., Scheibel, A.B., Fisher, R.S., Zaidel, E., 1992. Fiber composition of the human corpus callosum. *Brain Res.* 598 (1), 143–153.
- Alexander, D.C., 2008. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Mag. Res. Med.* 60 (2), 439–448.
- Alexander, D.C., Hubbard, P.L., Hall, M.G., Moore, E.A., Pitto, M., Parker, G.J., Dyrby, T.B., 2010. Orientationally invariant indices of axon diameter and density from diffusion MRI. *Neuroimage* 52 (4), 1374–1389.
- Alexander, D.C., Dyrby, T.B., 2013. Diffusion imaging with stimulated echoes: signal models and experiment design. *ArXiv:1305.7367v2*.
- Assaf, Y., et al., 2008. AxCaliber: a method for measuring axon diameter distribution from diffusion MRI. *Mag. Res. Med.* 59 (6), 1347–1354.
- Breiman, Leo., 1996. Heuristics of instability and stabilization in model selection. *Ann. Stat.* 24 (6), 2350–2383.
- Burcaw, L.M., Fieremans, E., Novikov, D.S., 2015. Mesoscopic structure of neuronal tracts from time-dependent diffusion. *NeuroImage* 114, 18–37.
- Callaghan, P.T., 1997. A simple matrix formalism for spin echo analysis of restricted diffusion under generalized gradient waveforms. *J. Magn. Reson.* 129 (1), 74–84.
- Codd, S.L., Callaghan, P.T., 1999. Spin echo analysis of restricted diffusion under generalized gradient waveforms: planar, cylindrical, and spherical pores with wall relaxivity. *J. Magn. Reson.* 137 (2), 358–372.
- Criminisi, A., et al., 2013. *Decision Forests for Computer Vision and Medical Image Analysis*. Springer.
- De Luca, G.C., Ebers, G.C., Esiri, M.M., 2004. Axonal loss in multiple sclerosis: a pathological survey of the corticospinal and sensory tracts. *Brain* 127, 1009–1018.
- DeLuca, G.C., Williams, K., Evangelou, N., Ebers, G.C., Esiri, M.M., 2006. The contribution of demyelination to axonal loss in multiple sclerosis. *Brain* 129 (6), 1507–1516.
- Drobnjak, I., Zhang, H., Ianus, A., Kaden, E., Alexander, D.C., 2016. PGSE, OGSE, and sensitivity to axon diameter in diffusion MRI: insight from a simulation study. *Magn. Reson. Med.* 75 (2), 688–700.
- Drobnjak, I., Siow, B., Alexander, D.C., 2010. Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted MR. *J. Magn. Reson.* 206 (1), 41–51.
- Drobnjak, I., Zhang, H., Hall, M.G., Alexander, D.C., 2011. The matrix formalism for generalised gradients with time-varying orientation in diffusion NMR. *J. Magn. Reson.* 210 (1), 151–157.
- Filippi, M., Rocca, M.A., Barkhof, F., Breck, W., Chen, J.T., Comi, G., DeLuca, G., De Stefano, N., Erickson, B.J., Evangelou, N., Fazekas, F., Geurts, J.J., Lucchinetti, C., Miller, D.H., Pelletier, D., Popescu, B.F., Lassmann, H., 2012. Association between pathological and MRI findings in multiple sclerosis. *Lancet Neurol.* 11 (4), 349–360.
- Fieremans, E., Novikov, D., Jensen, J.H., Helpert, J.A., 2010. Monte Carlo study of a two-compartment exchange model of diffusion NMR. *Biomedicine* 23 (7), 711–724.
- Finkelstein, A., 1976. Water and nonelectrolyte permeability of lipid bilayer membranes. *J. General. Physiol.* 68, 127–135.
- Gilks, W.R., Richardson, S., Spiegelhalter, D., 1994. *Markov Chain Monte Carlo in Practice*. Chapman and Hall.
- Grebenkov, D.S., Van Nguyen, D., Li, J.-R., 2014. Exploring diffusion across permeable barriers at high gradients: I. Narrow pulse approximation. *J. Magn. Reson.* 248, 153–163.
- Hall, M.G., Alexander, D.C., 2009. Convergence and parameter choice for Monte-Carlo simulations of diffusion MRI. *IEEE Trans. Med. Im.* 28 (9), 1354–1364.
- Hall, M.G., Nedjati-Gilani, G., Alexander, D.C., 2014. Monte-Carlo simulation of diffusion MRI with realistic voxel sizes. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*, 2463.
- Huang, S.Y., Tobyne, S.M., Nummenmaa, A., Witzel, T., Wald, L.L., McNab, J.A., Klawiter, E.C., 2016. Characterization of Axonal Disease in Patients with Multiple Sclerosis Using High-Gradient-Diffusion MR Imaging Radiology. Feb 09. <http://dx.doi.org/10.1148/radiol.2016151582>. [Epub ahead of print].
- Jenkinson, M., Bannister, P.R., Brady, J.M., Smith, S.M., 2002. Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *NeuroImage* 17 (2), 825–841.
- Jelescu, I.O., Veraart, J., Fieremans, E., Novikov, D.S., 2016. Degeneracy in model parameter estimation for multi-compartmental diffusion in neuronal tissue. *NMR Biomed.* 29, 33–47.
- Jespersen, S.N., Pedersen, M., Stodkilde-Jorgensen, H., 2005. The influence of a cellular size distribution on NMR diffusion measurements. *Eur. Biophys. J.* 34, 890–898.
- Kärger, J., Pfeifer, H., Wilfried, H., 1988. Principles and application of self-diffusion measurements by nuclear magnetic resonance. *Adv. Mag. Res.* 12 (1), 1–89.
- Kaden, E., Kruggel, F., Alexander, D.C., 2016. Quantitative mapping of the per-axon diffusion coefficients in brain white matter. *Magn. Reson. Med.* 75 (4), 1752–1763.
- Lampinen, B., Szczepankiewicz, F., van Westen, D., Ståhlberg, F., Lätt, J., Nilsson, M., 2014. Apparent exchange rate (AXR) mapping in diffusion MRI: An in vivo test-retest study and analysis of statistical power. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*, 2639.
- Lasić, Samo, Jespersen, Sune N., Lundell, Henrik, Nilsson, Markus, Dyrby, Tim B., Topgaard, Daniel, 2016. Apparent exchange rate in multi-compartment anisotropic tissue. In: *Proceedings of the International Society for Magnetic Resonance in Medicine*, 3083.
- Lasić, S., Nilsson, M., Lätt, J., Ståhlberg, F., Topgaard, D., 2011. Apparent exchange rate mapping with diffusion MRI. *Magn. Reson. Med.* 66 (2), 356–365.
- Lätt, J., Nilsson, M., van Westen, D., Wirestam, R., Ståhlberg, F., Brockstedt, S., 2009. Diffusion-weighted MRI measurements on stroke patients reveal water-exchange mechanisms in sub-acute ischaemic lesions. *NMR Biomed.* 22, 619–628.
- Lundell, H., Alexander, D.C., Dyrby, T.B., 2014. High angular resolution diffusion imaging with stimulated echoes: compensation and correction in experiment design and analysis. *NMR Biomed.* 27, 918–925.
- Nilsson, M., Alerstam, E., Wirestam, R., Ståhlberg, F., Brockstedt, S., Lätt, J., 2010. Evaluating the accuracy and precision of a two-compartment Kärger model using Monte Carlo simulations. *J. Mag. Res.* 206, 59–67.
- Nilsson, M., Lätt, J., van Westen, D., Brockstedt, S., Lasić, S., Ståhlberg, F., Topgaard, D., 2013. Noninvasive mapping of water diffusional exchange in the human brain using filter-exchange imaging. *Mag. Res. Med.* 69 (6), 1573–1581.
- Nilsson, M., van Westen, D., Ståhlberg, F., 2013. The role of tissue microstructure and water exchange in biophysical modelling of diffusion in white matter. *Magn. Reson. Mater. Phys.* 26, 345–370.
- Panagiotaki, E., et al., 2012. Compartment models of the diffusion MR signal in brain white matter: a taxonomy and comparison. *NeuroImage* 59 (3), 2241–2254.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. *JMLR* 12, 2825–2830.
- Quirk, J.D., Brethorst, G.L., Duong, T.Q., Snyder, A.Z., Springer, C.S., Ackerman, J.J.H., Neil, J.J., 2003. Equilibrium water exchange between the intra- and extracellular spaces of mammalian brain. *Mag. Res. Med.* 50, 493–499.
- Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E.J., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., Niaz, R.,

- Saunders, J., Vickers, J., Zhang, Y., De Stefano, N., Brady, J.M., Matthews, P.M., 2004. Advances in functional and structural MR image analysis and implementation as FSL. *NeuroImage* 23 (S1), 208–219.
- Stanisz, G.J., Wright, G.A., Henkelman, R.M., Szafer, A., 1997. An analytical model of restricted diffusion in bovine optic nerve. *Mag. Res. Med.* 37 (1), 103–111.
- Stanisz, G.J., Odobina, E.E., Pun, J., Escaravage, M., Graham, S.J., Bronskill, M.J., Henkelman, R.M., TStanisz, G.J., 2005. T1, T2 relaxation and magnetization transfer in tissue at 3T. *Mag. Res. Med.* 54 (3), 507–512.
- Shemesh, N., Jespersen, S.N., Alexander, D.C., Cohen, Y., Drobnjak, I., Dyrby, T.B., Finsterbusch, J., Koch, M.A., Kuder, T., Laun, F., Lawrenz, M., Lundell, H., Mitra, P.P., Nilsson, M., Özarslan, E., Topgaard, D., Westin, C.-F., 2016. Conventions and nomenclature for double diffusion encoding NMR and MRI. *Magn. Reson. Med.* 75m, 82–87.
- Szafer, A., Zhong, J., Gore, J.C., 1995. Theoretical model for water diffusion in tissues. *Mag. Res. Med.* 33, 697–712.
- Tanno, R., Ghosh, A., Grussu, F., Kaden, E., Criminisi, A., Alexander, D.C., 2016. Bayesian Image Quality Transfer. *Proceedings of the MICCAI*, pp. 1–9.
- Tariq, M., Schneider, T., Alexander, D.C., Wheeler-Kingshott, C.A.M., Zhang, H., 2013. Assessing scan-rescan reproducibility of the parameter estimates from NODDI. *ISMRM* 21, (3187–3187).
- Van Gelderen, P., DesPres, D., Van Zijl, P.C.M., Moonen, C.T.W., 1994. Evaluation of restricted diffusion in cylinders. Phosphocreatine in rabbit leg muscle. *J. Magn. Reson. Ser. B* 103, 255–260.
- Vollmar, C., O'Muircheartaigh, J., Barker, G.J., Symms, M.R., Thompson, P., Kumari, V., Duncan, J.S., Richardson, M.P., Koepp, M.J., 2010. Identical, but not the same: intra-site and inter-site reproducibility of fractional anisotropy measures on two 3.0 T scanners. *NeuroImage* 51 (4), 1384–1394.
- Volles, M.J., Lee, S.-J., Rochet, J.-C., Shtilerman, M.D., Ding, T.T., Kessler, J.C., Lansbury, P.T., Jr., 2001. Vesicle permeabilization by protofibrillar α -synuclein: implications for the pathogenesis and treatment of Parkinson's disease. *Biochemistry* 40, 7812–7819.
- Westin, C.-F., Maier, S.E., Khidhir, B., Everett, P., Jolesz, F.A., Kikinis, R., 1999. Image processing for diffusion tensor magnetic resonance imaging. In: *Proceedings of the MICCAI*, pp. 441–452.