



# Fitting a Directional Microstructure Model to Diffusion-Relaxation MRI Data with Self-supervised Machine Learning

Jason P. Lim<sup>1</sup>, Stefano B. Blumberg<sup>1,2</sup>, Neil Narayan<sup>1</sup>, Sean C. Epstein<sup>1</sup>, Daniel C. Alexander<sup>1</sup>, Marco Palombo<sup>1,3,4</sup>, and Paddy J. Slator<sup>1</sup>(✉)

<sup>1</sup> Centre for Medical Image Computing, Department of Computer Science, University College London, London, UK

[p.slator@ucl.ac.uk](mailto:p.slator@ucl.ac.uk)

<sup>2</sup> Centre for Artificial Intelligence, Department of Computer Science, University College London, London, UK

<sup>3</sup> Cardiff University Brain Research Imaging Centre (CUBRIC), School of Psychology, Cardiff University, Cardiff, UK

<sup>4</sup> School of Computer Science and Informatics, Cardiff University, Cardiff, UK

**Abstract.** Machine learning is a powerful approach for fitting microstructural models to diffusion MRI data. Early machine learning microstructure imaging implementations trained regressors to estimate model parameters in a supervised way, using synthetic training data with known ground truth. However, a drawback of this approach is that the choice of training data impacts fitted parameter values. Self-supervised learning is emerging as an attractive alternative to supervised learning in this context. Thus far, both supervised and self-supervised learning have typically been applied to isotropic models, such as intravoxel incoherent motion (IVIM), as opposed to models where the directionality of anisotropic structures is also estimated. In this paper, we demonstrate self-supervised machine learning model fitting for a directional microstructural model. In particular, we fit a combined T1-ball-stick model to the multidimensional diffusion (MUDI) challenge diffusion-relaxation dataset. Our self-supervised approach shows clear improvements in parameter estimation and computational time, for both simulated and in-vivo brain data, compared to standard non-linear least squares fitting. Code for the artificial neural net constructed for this study is available for public use from the following GitHub repository: <https://github.com/jplte/deep-T1-ball-stick>.

**Keywords:** Microstructure imaging · Machine learning · Self-supervised learning

## 1 Introduction

Microstructure imaging aims to quantify features of the tissue microstructure from in-vivo MRI [1]. Historically, microstructure imaging utilised diffusion MRI

(dMRI) data. Recently, combined diffusion-relaxation MRI - where relaxation-encoding parameters such as inversion time (TI) and echo time (TE) are varied alongside diffusion-encoding parameters such as b-value and gradient direction - has been emerging as an extension [23]. The typical approach to estimating tissue microstructure from such diffusion or diffusion-relaxation data is multi-compartment modelling, which utilises signal models comprising linear combinations of multiple compartments - such as balls, sticks, zeppelins, and spheres - each representing a distinct tissue geometry [21].

Multi-compartment microstructure models are usually fit to the data with non-linear least squares (NLLS) algorithms. However, these can be computationally expensive and are prone to local minima, necessitating grid searches or parameter cascading [9] to seek global minima. Machine learning is a powerful alternative. Thus far, most machine learning microstructure model fitting approaches have used supervised learning [7, 10, 14, 17–20]. However, a crucial limitation is that the distribution of training data significantly affects fitted parameters [8, 12]. It has also proved difficult to estimate directional parameters, such as fibre direction, with existing machine learning methods instead directly estimating rotationally invariant parameters, such as mean diffusivity, fractional anisotropy, mean kurtosis, and orientation dispersion. This may be due to the difficulty of constructing a training dataset that adequately samples the high-dimensional parameter space, and/or complications due to the periodicity of angular parameters.

Self-supervised (sometimes imprecisely called unsupervised in the microstructure imaging context) learning is an alternative with the potential to address these limitations. Self-supervised algorithms learn feature representations from the input data by inferring supervisory constraints from the data itself. For microstructure imaging, self-supervised learning has been implemented with voxelwise fully connected artificial neural networks (ANNs). However, thus far self-supervised microstructure imaging has been limited to isotropic models [6, 11], including many intravoxel incoherent motion (IVIM) MRI examples [2, 8, 13, 25, 26]. To our knowledge, self-supervised model fitting has not yet been demonstrated for directional microstructural models.

In this paper, we fit an extended T1-ball-stick model to diffusion-relaxation MRI data using self-supervised machine learning and demonstrate several advantages of this approach over classical NLLS, such as higher precision and faster computational time.

## 2 Methods

### 2.1 Microstructure Model

As this is a first attempt at fitting directional multi-compartment models with self-supervised learning, we choose a simple model - the ball-stick model first proposed by Behrens et al. [3]. According to the ball-stick model, the expression for the normalized signal decay is

$$S(b, \mathbf{g}) = f \exp(-b\lambda_{||}(\mathbf{g}, \mathbf{n})) + (1 - f) \exp(-b\lambda_{iso}) \quad (1)$$

where  $b$  is the b-value,  $\mathbf{g}$  is the gradient direction,  $\lambda_{\parallel}$  and  $\lambda_{iso}$  are the parallel and isotropic diffusivities of the stick and ball respectively, and  $\mathbf{n}$  is the stick orientation, which we parameterise using polar coordinates. The relationship between Cartesian and polar coordinates is  $n = [\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta]$  where  $\phi \in [0, \pi]$  and  $\theta \in [-\pi, \pi]$ .

We extend the ball-stick model to account for T1 relaxation time, by assuming the ball and stick compartments have separate T1 times, represented by  $T1_{ball}$  and  $T1_{stick}$  respectively. Note that we assume a single T2 for both compartments, so the volume fraction  $f$  will be affected by the T2 of each compartment. Given a combined T1 inversion recovery [5] and diffusion MRI experiment, where inversion time (TI), b-value and gradient direction are simultaneously varied, we can fit the following T1-ball-stick equation

$$S(b, g, T_I, T_R) = f \exp(-b\lambda_{\parallel}(\mathbf{g}, \mathbf{n})) \left| 1 - 2 \exp\left(-\frac{T_I}{T1_{stick}}\right) \exp\left(-\frac{T_R}{T1_{stick}}\right) \right| + (1 - f) \exp(-b\lambda_{iso}) \left| 1 - 2 \exp\left(-\frac{T_I}{T1_{ball}}\right) \exp\left(-\frac{T_R}{T1_{ball}}\right) \right| \quad (2)$$


In this work, we first fit this model to combined T1-diffusion data with standard NLLS, then demonstrate self-supervised fitting with an ANN. We first describe the data, then the model fitting techniques.

## 2.2 Combined T1-Diffusion in Vivo Data

We utilise in-vivo data from 5 healthy volunteers (3 F, 2 M, age = 19–46 years), acquired from the 2019 multidimensional diffusion (MUDI) challenge [22]. The acquisition sequence comprises simultaneous diffusion, inversion recovery (giving T1 contrast), and multi-echo gradient echo (giving T2\* contrast) measurements. We chose to ignore the subsection of the data that is sensitive to T2\* by only included signals captured with the lowest echo time (80 ms). This is since the two higher TEs have very low signal intensity and the 3 TEs have a small range, and thus there is limited T2\* information in the data. Our subsequent description hence only refers to the subsection of the data with TE = 80 ms.

The datasets were obtained using a clinical 3T Philips Achieva scanner (Best, Netherlands) with a 32-channel adult headcoil. Each scan includes 416 volumes distributed over five b-shells,  $b \in \{0, 500, 1000, 2000, 3000\}$  s/mm<sup>2</sup>, with 16 uniformly spread directions, and 28 inversion times (TI)  $\in [20, 7322]$  ms. For all datasets, the following parameters were fixed: repetition time TR = 7.5 s, resolution = 2.5 mm isotropic, FOV = 220 × 230 × 140 mm, SENSE = 1.9, halfscan = 0.7, multiband factor 2, total acquisition time 52 min (including preparation time).

The MUDI data has already undergone standard pre-processing, see [22] for full details. Upon inspection, we noted that the lowest (20 ms) and highest (7322 ms) inversion times, which comprise 7.14% of the data, were clearly dominated by noise and/or artifacts (see Fig. 1). We therefore removed them from the data prior to model fitting, leaving 416 MRI volumes in total. After removing these TIs, the dataset contains 26 TIs  $\in [176, 4673]$  ms.



**Fig. 1.** Average signal of all  $b=0$  voxels within the brain mask at each inversion time. Note the outlying signals at the smallest (20 ms, blue) and largest (7322 ms, green) inversion times. We removed these from the data before fitting. (Color figure online)

We normalised each voxel’s data independently, by dividing by the signal generated from the  $b=0$  volume with the highest TI, i.e. the volume with the highest expected signal. We then removed all background voxels using the brain mask provided with the MUDI data.

### 2.3 Simulated Data

In-vivo MRI data does not have ground truth tissue-related parameters values, making it hard to quantitatively assess the accuracy of model fitting. We thus simulated 100,000 synthetic signals using the T1-ball-stick model signal equation (Eq. (2)). We used the same 416 acquisition parameters (b-values, gradient directions and inversion times) as our reduced MUDI dataset. The ground truth values of  $\lambda_{||}$ ,  $\lambda_{iso}$ ,  $T1_{ball}$  and  $T1_{stick}$  were sampled randomly from physically-plausible ranges (Table 1). Note that we choose units so that parameter values are close to 1; this prevents having to normalise parameter values before training neural networks. Complex Gaussian noise was added to simulate the Rician distribution of noisy MRI data [16].

### 2.4 Non-linear Least Squares Fitting

The modified T1-ball-stick model (Eq. (2)) was fit with non-linear least squares by modifying the open source diffusion microstructure imaging in python (dmipy) toolbox [9], with parameter constraints as in Table 1. Specifically, we used the “brute2fine” function, which uses a brute force grid search followed by non-linear optimisation.

**Table 1.** Constraints on T1-ball-stick parameters for simulating data and model fitting.

Parameter	Minimum	Maximum
$\lambda_{iso}$	0.1 $\mu\text{m}^2/\text{ms}$	3.0 $\mu\text{m}^2/\text{ms}$
$\lambda_{  }$	0.1 $\mu\text{m}^2/\text{ms}$	3.0 $\mu\text{m}^2/\text{ms}$
$f$	0	1
$T1_{ball}$	0.01 s	5 s
$T1_{stick}$	0.01 s	5 s
$\theta$	0	$\pi$
$\phi$	$-\pi$	$\pi$

## 2.5 Self-supervised Model Fitting

We built an ANN to generate estimates for the T1-ball-stick parameters. The network comprises an input layer, 3 fully connected hidden layers and an output layer, (see Fig. 2). The input layer and hidden layers each have 416 nodes - mirroring the 416 MRI volumes. The final layer has 7 output neurons, one for each parameter of interest. The normalised signal from a single voxel of the MRI data,  $S$ , which comprises 416 measurements, is fed into the input layer and passed through the ANN. The output layer is fed forward into the T1-ball-stick model equation, giving a synthetic signal  $\hat{S}$ . Training loss is the mean squared error between input ( $S$ ) and synthetic ( $\hat{S}$ ) MRI signals across all voxels passed through the ANN.

We implemented the ANN on Python 3.9.5 using PyTorch21 1.10.0. For both simulated and in-vivo data, we used the Adam optimiser [15] with learning rate 0.0001, batch size 128, and dropout [24] with rate 0.5. Parameter constraints (Table 1) were imposed using PyTorch’s clamp feature, which converts any value outside the bounds to the value closest to it within the boundary. Following [2], we trained the network with patience 10, i.e. until there 10 consecutive epochs without loss improvement. As the self-supervised approach estimates T1-ball-stick parameters directly from the data, we didn’t use a train-test split. Instead, the ANN was trained on each dataset separately.

## 3 Results

### 3.1 Simulated Data

The T1-ball-stick model was successfully fit to the simulated signals using both the ANN and NLLS. Scatter plots and Pearson correlation coefficients of parameter estimates against ground truth values are shown in Fig. 3 (note that we don’t report correlations for  $\theta$  and  $\phi$  as the values are confounded by the periodicity of these angular parameters). Correlation coefficients for the ANN fits are higher for all model parameters, with coefficients above 0.9 for all parameters except  $\lambda_{||}$ . NLLS correlation coefficients for T1 relaxation time are particularly low.

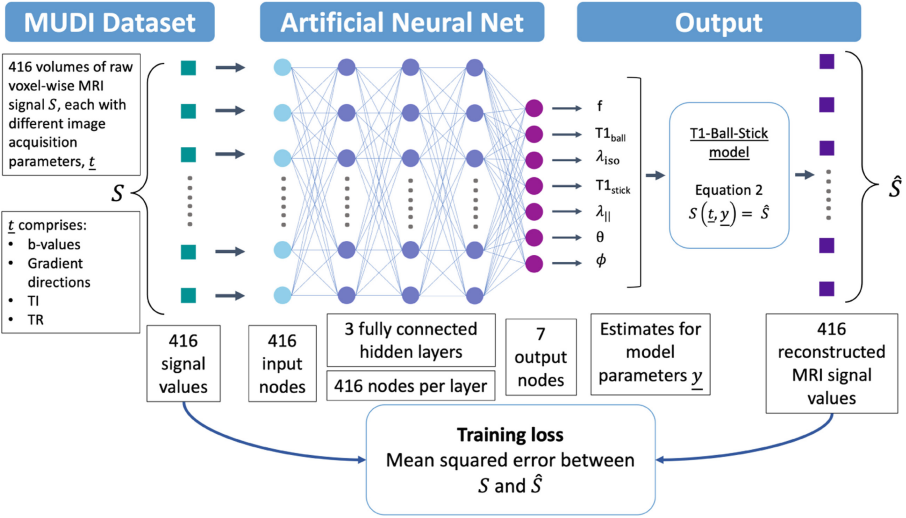


Fig. 2. Our ANN for T1-ball-stick fitting and the flow of data through it.

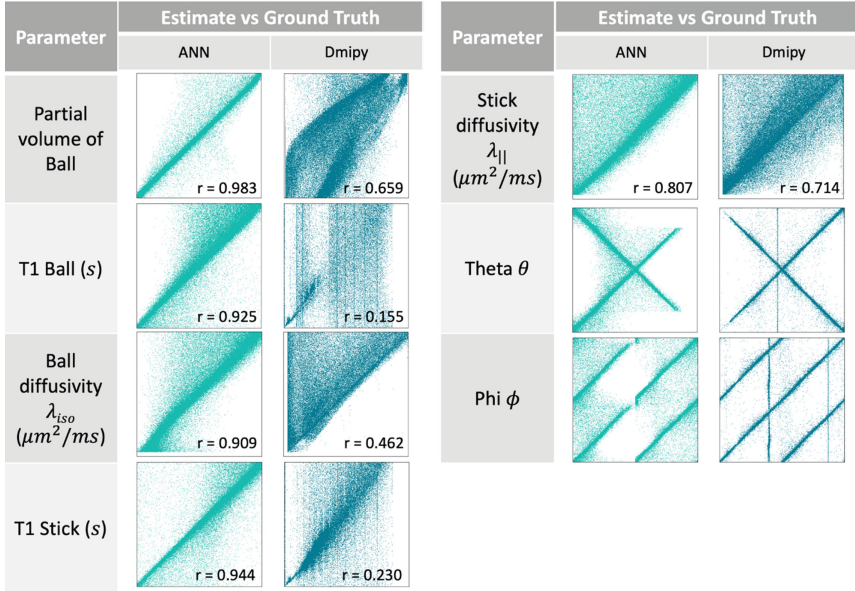
### 3.2 Real Data

Figures 4 and 5 show T1-ball-stick model fits for all MUDI subjects. ANN parameter maps are qualitatively less noisy and show more anatomically plausible contrast than NLLS. This is particularly clear for  $T1_{ball}$ ,  $\lambda_{||}$  and direction encoded colour (DEC) maps. All 5 MUDI subjects showed similar trends. ANN inferred higher and lower values than NLLS for  $T1_{ball}$  and  $T1_{stick}$  respectively. The ANN clearly shows highest  $\lambda_{||}$  values in the corpus callosum, while the Dmipy fit has high  $\lambda_{||}$  values in many places.  $\lambda_{iso}$  maps are generally similar across both methods.

Average time taken for ANN fits on real data was 1966s, compared to 8833s in Dmipy, meaning ANN was 77.25% faster than NLLS on average. All model fits were performed on a 2017 Macbook Pro’s central processing unit (3.1 GHz Dual-Core Intel I5-7267U).

## 4 Discussion

This study demonstrates self-supervised microstructure imaging for a combined T1-ball-stick model. Our ANN approach is faster and more precise than conventional NLLS model fitting. In the ANN model fits, parametric maps show plausible estimates for both diffusivity ( $\lambda_{||}$ ,  $\lambda_{iso}$ ) and relaxation ( $T1_{ball}$  and  $T1_{stick}$ ), whereas some NLLS maps, particularly  $T1_{ball}$ , show dubious contrast.

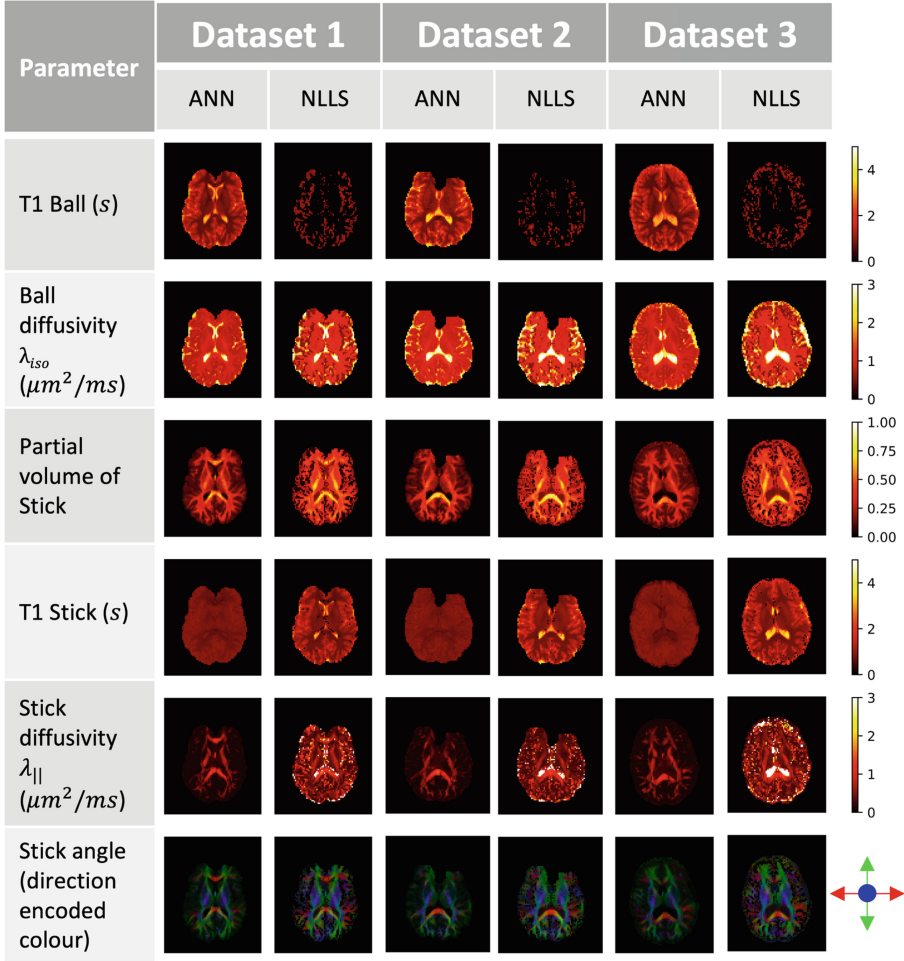


**Fig. 3.** NLLS and ANN T1-ball-stick estimates against ground truth values, after fitting the model to simulated signals.  $r$  denotes the Pearson correlation coefficient.

The ANN model fits potentially reveal more accurate tissue information than NLLS counterparts. Throughout the brain, white matter T1 times are expected to be around 0.7–0.9 s [4]. The ANN estimates fall within this range - the  $T1_{stick}$  voxels displayed in Figs. 4 and 5 average to 0.87 s - while NLLS estimates are higher, with some regions reaching 4 s. T1 times in the CSF are expected to be around 4 s [4], which is reflected in the ANN  $T1_{ball}$  estimate, but not in the NLLS estimate, where it is approximately 0 s (Figs. 4 and 5).

These observations match those in the simulations (Fig. 3), with NLLS correlation being extremely poor in both T1 times. In line with our qualitative analysis of MUDI data fits, Figure 3 shows that the ANN outperforms Dmipy in every parameter, except potentially for  $\theta$  and  $\phi$ , whose correlations are not straightforward to quantify due to their periodicity. Direction encoded colour (DEC) maps are less noisy for ANN fits, with less visible noise, and are hence potentially more useful for tractography.

The ANN approach was faster in all datasets, with a 77% improvement in time on average. This will hopefully improve the feasibility of utilising similar modelling in clinical scenarios. However, the ANN would still require retraining for every new dataset. A possible next step would be to explore the viability of using multiple datasets to train an ANN to be generalisable to unseen data, then fine-tuning the network on each new dataset, effectively combining advantages of supervised and self-supervised learning, as recently demonstrated by Epstein et al. [8].



**Fig. 4.** Parameter maps for all models fit to the first 3 MUDI datasets using both Dmipy and ANN methods. Each parameter map is a cross-sectional view of the brain. The maps presented are generated from the middle Z-axis slice.

If successful, this could significantly reduce time taken to generate parameter estimates for new patients, but possibly at the cost of accuracy.

Whilst the ANN outperforms NLLS in our experiments, we applied NLLS “out of the box” without focusing on improving the fitting. The NLLS fits could be improved - e.g. the vertical lines in Fig. 3 are likely local minima. . In future, we could use a larger grid in the grid search stage, although this can quickly lead to infeasible computational times, or initialise the NLLS fit from “reasonable” parameter values. Whilst these would likely improve the NLLS fits, the fact that our self-supervised ANN bypasses these ad-hoc tuning steps presents a significant advantage. We also only compare parameter estimates with the ground truth



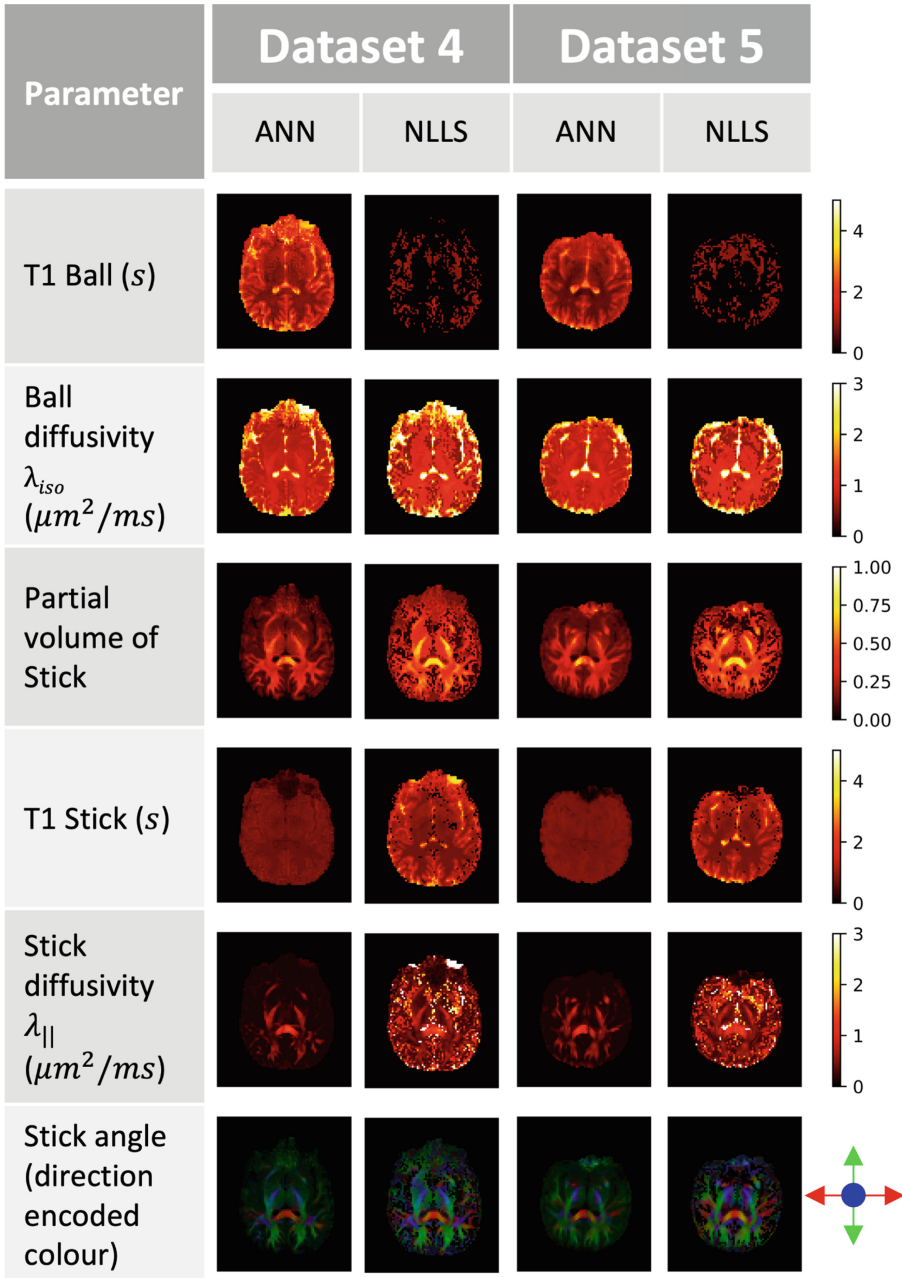


Fig. 5. As Fig. 4 but for the remaining 2 MUDI datasets.

using correlation coefficients, effectively merging bias and variance. In future, we could calculate bias and variance separately, and also explore tuning the training cost function towards accuracy or precision depending on the application. Ball-stick is a very simple single bundle model, in future we will explore multi-fibre models that more accurately reflect brain microstructure.

The small sample size is a limitation of this study. In future, we can adapt our ANN to fit standard microstructural models that only require diffusion MRI data. This would enable us to test self-supervised learning against NLLS, and quantify test-retest repeatability, on large open source datasets. Additionally, the 5 datasets are all from healthy patients with normal physiology. Hence, we are unable to judge the suitability of the T1-ball-stick model parameters as imaging biomarkers. The overall goal is to identify imaging biomarkers that can be used for diagnosis, prognosis, and monitoring of brain conditions such as stroke and dementia, so future studies should involve MRI data from patients with these conditions. This would also help us determine thresholds for model parameters to differentiate between healthy and diseased tissue.

## 5 Conclusion

We demonstrate, for the first time, self-supervised learning fitting of a directional microstructural model, T1-ball-stick. We show vastly improved performance, in terms of speed and accuracy, compared to the current standard fitting technique, NLLS. Self-supervised machine learning model fitting had only been demonstrated in a limited number of simple MC models thus far, such as IVIM [2]. This work can pave the way for self-supervised fitting of a wide range of multi-compartment microstructure models to MRI data.

**Acknowledgments.** This work was supported by EPSRC grants EP/M020533/1, EP/V034537/1, and EP/R014019/1; UKRI Future Leaders Fellowship MR/T020296/2 (MP); an EPSRC and Microsoft scholarship (SBB); and the National Institute for Health Research (NIHR) Biomedical Research Centre at University College London Hospitals NHS Foundation Trust and University College London. The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health.

## References

1. Alexander, D.C., Dyrby, T.B., Nilsson, M., Zhang, H.: Imaging brain microstructure with diffusion MRI: practicality and applications, April 2019. <https://doi.org/10.1002/nbm.3841>
2. Barbieri, S., Gurney-Champion, O.J., Klaassen, R., Thoeny, H.C.: Deep learning how to fit an intravoxel incoherent motion model to diffusion-weighted MRI. *Magn. Reson. Med.* **83**, 312–321 (2020). <https://doi.org/10.1002/mrm.27910>
3. Behrens, T.E., et al.: Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magn. Reson. Med.* **50**, 1077–1088 (2003). <https://doi.org/10.1002/mrm.10609>

4. Bojorquez, J.Z., Bricq, S., Acquitte, C., Brunotte, F., Walker, P.M., Lalande, A.: What are normal relaxation times of tissues at 3t?, January 2017. <https://doi.org/10.1016/j.mri.2016.08.021>
5. Bydder, G.M., Young, I.R.: MR imaging: clinical use of the inversion recovery sequence. *J. Comput. Assist. Tomogr.* **9**, 659–675 (1985)
6. de Almeida Martins, J.P., et al.: Neural networks for parameter estimation in microstructural MRI: application to a diffusion-relaxation model of white matter. *NeuroImage* **244**, 118601 (2021). <https://doi.org/10.1016/j.neuroimage.2021.118601>
7. Diao, Y., Jelescu, I.O.: Parameter estimation for WMTI-Watson model of white matter using encoder-decoder recurrent neural network. *ArXiv Preprint* (2022)
8. Epstein, S.C., Bray, T.J.P., Hall-Craggs, M., Zhang, H.: Choice of training label matters: how to best use deep learning for quantitative MRI parameter estimation. *ArXiv Preprint*, May 2022. <https://doi.org/10.48550/arxiv.2205.05587>
9. Fick, R.H.J., Wassermann, D., Deriche, R.: The Dmipy toolbox: diffusion MRI multi-compartment modeling and microstructure recovery made easy. *Front. Neuroinform.* **13**, 64 (2019). <https://doi.org/10.3389/fninf.2019.00064>
10. Golkov, V., et al.: q-space deep learning: twelve-fold shorter and model-free diffusion MRI scans. *IEEE Trans. Med. Imaging* **35**(5), 1344–1351 (2016). <https://doi.org/10.1109/TMI.2016.2551324>
11. Grussu, F., Battiston, M., Palombo, M., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: Deep learning model fitting for diffusion-relaxometry: a comparative study. In: *Computational Diffusion MRI (CDMRI) 2021*, pp. 159–172. Springer (2021). [https://doi.org/10.1007/978-3-030-73018-5\\_13](https://doi.org/10.1007/978-3-030-73018-5_13)
12. Gyori, N.G., Palombo, M., Clark, C.A., Zhang, H., Alexander, D.C.: Training data distribution significantly impacts the estimation of tissue microstructure with machine learning. *Magn. Reson. Med.* **87**, 932–947 (2022). <https://doi.org/10.1002/mrm.29014>
13. Kaandorp, M.P., et al.: Improved unsupervised physics-informed deep learning for intravoxel incoherent motion modeling and evaluation in pancreatic cancer patients. *Magn. Reson. Med.* **86**(4), 2250–2265 (2021). <https://doi.org/10.1002/mrm.28852>
14. Kerkelä, L., Seunarine, K., Henriques, R.N., Clayden, J.D., Clark, C.A.: Improved reproducibility of diffusion kurtosis imaging using regularized non-linear optimization informed by artificial neural networks (2022). [arxiv:2203.07327](https://arxiv.org/abs/2203.07327)
15. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2014). [arxiv:1412.6980](https://arxiv.org/abs/1412.6980)
16. Koay, C.G., Basser, P.J.: Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. *J. Magn. Reson.* **179**, 317–322 (2006). <https://doi.org/10.1016/j.jmr.2006.01.016>
17. Li, Z., et al.: Fast and robust diffusion kurtosis parametric mapping using a three-dimensional convolutional neural network. *IEEE Access* **7**, 71398–71411 (2019). <https://doi.org/10.1109/ACCESS.2019.2919241>
18. Nedjati-Gilani, G.L., et al.: Machine learning based compartment models with permeability for white matter microstructure imaging. *NeuroImage* **150**, 119–135 (2017). <https://doi.org/10.1016/j.neuroimage.2017.02.013>
19. Palombo, M., et al.: SANDI: a compartment-based model for non-invasive apparent soma and neurite imaging by diffusion MRI. *NeuroImage* **215**, 116835 (2020). <https://doi.org/10.1016/j.neuroimage.2020.116835>

20. Palombo, M., et al.: Joint estimation of relaxation and diffusion tissue parameters for prostate cancer grading with relaxation-VERDICT MRI. medRxiv 1(165) (2021). <https://doi.org/10.1101/2021.06.24.21259440>
21. Panagiotaki, E., Schneider, T., Siow, B., Hall, M.G., Lythgoe, M.F., Alexander, D.C.: Compartment models of the diffusion MR signal in brain white matter: A taxonomy and comparison. *Neuroimage* **59**(3), 2241–2254 (2012). <https://doi.org/10.1016/j.neuroimage.2011.09.081>
22. Pizzolato, M., et al.: Acquiring and predicting multidimensional diffusion (MUDI) data: an open challenge (2020). [https://doi.org/10.1007/978-3-030-52893-5\\_17](https://doi.org/10.1007/978-3-030-52893-5_17)
23. Slator, P.J., et al.: Combined diffusion-relaxometry microstructure imaging: current status and future prospects, December 2021. <https://doi.org/10.1002/mrm.28963>
24. Srivastava, N., Hinton, G., Krizhevsky, A., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting (2014)
25. Vasylechko, S.D., Warfield, S.K., Afacan, O., Kurugol, S.: Self-supervised IVIM DWI parameter estimation with a physics based forward model. *Magn. Reson. Med.* **00**, 1–11 (2021). <https://doi.org/10.1002/mrm.28989>
26. Zhou, X.X., et al.: An unsupervised deep learning approach for dynamic-exponential intravoxel incoherent motion MRI modeling and parameter estimation in the liver. *J. Magn. Reson. Imaging* (2022). <https://doi.org/10.1002/jmri.28074>