

## RESEARCH ARTICLE OPEN ACCESS

# Hierarchical Bayesian Modelling Improves Microstructural Parameter Mapping in Diffusion and Exchange MRI Data

Elizabeth Powell<sup>1</sup>  | Mark Maskery<sup>2,3</sup>  | Hedley C. A. Emsley<sup>2,3</sup>  | Laura M. Parkes<sup>4,5</sup>  | Geoff J. M. Parker<sup>1,6</sup>  | Paddy J. Slator<sup>7,8</sup> 

<sup>1</sup>Department of Medical Physics and Biomedical Engineering, University College London, London, UK | <sup>2</sup>Lancaster Medical School, Lancaster University, Lancaster, UK | <sup>3</sup>Department of Neurology, Lancashire Teaching Hospitals NHS Foundation Trust, Preston, UK | <sup>4</sup>Division of Psychology, Communication and Human Neuroscience, School of Health Sciences, Faculty of Biology, Medicine and Health, University of Manchester, Manchester, UK | <sup>5</sup>Geoffrey Jefferson Brain Research Centre, Faculty of Biology, Medicine and Health, University of Manchester, Manchester, UK | <sup>6</sup>Bioxydyn Limited, Manchester, UK | <sup>7</sup>Cardiff University Brain Research Imaging Centre, School of Psychology, Cardiff University, Cardiff, UK | <sup>8</sup>School of Computer Science and Informatics, Cardiff University, Cardiff, UK

**Correspondence:** Elizabeth Powell ([e-powell@ucl.ac.uk](mailto:e-powell@ucl.ac.uk))

**Received:** 2 June 2025 | **Revised:** 5 February 2026 | **Accepted:** 10 February 2026

**Keywords:** Bayesian modelling | blood–brain barrier | diffusion MRI | filter exchange imaging (FEXI) | kurtosis | microstructure | water exchange

## ABSTRACT

Microstructure modelling quantifies subvoxel tissue features by combining an MRI acquisition with a mathematical model, which is typically fitted voxel-by-voxel with least-squares (LSQ) minimisation to give voxelwise maps of microstructural quantities such as diffusivity and compartmental fractions. Such approaches are susceptible to voxelwise noise, which can lead to erroneous values in parameter maps. Hierarchical Bayesian modelling (HBM) can address this limitation but has only been demonstrated for simple models. We previously derived an HBM approach for an arbitrary microstructure model with flexible parameter constraints, utilising a Markov chain Monte Carlo algorithm for parameter estimation; here, the method is demonstrated and evaluated using simulated and human data for two previously unexplored diffusion MRI techniques, namely, diffusion kurtosis imaging and blood–brain barrier filter exchange imaging. When compared with LSQ minimisation, HBM increased the accuracy, precision, contrast-to-noise ratio and parameter map quality in both simulated and human data. HBM was also able to resolve local parameter variations associated with white matter lesions in a small sample of cerebral small vessel disease subjects, which were obscured by high noise levels in the LSQ-derived parameter maps. Finally, a noise sensitivity assessment in simulations showed that HBM improved the contrast-to-noise ratio and parameter map quality even at low signal-to-noise ratios. This generalised HBM framework can improve parameter estimation for more complex diffusion MRI microstructural models that extend beyond linear combinations of exponentials.

**Abbreviations:** ADC, apparent diffusion coefficient; AXR, apparent exchange rate; BBB, blood–brain barrier; BSP, Bayesian shrinkage prior; CNR, contrast-to-noise ratio; cSVD, cerebral small vessel disease; DKI, diffusion kurtosis imaging; dMRI, diffusion MRI; DTI, diffusion tensor imaging; FA, fractional anisotropy; FEXI, filter exchange imaging; FLAIR, fluid-attenuated inversion recovery; GM, grey matter; HBM, hierarchical Bayesian modelling; HCP, Human Connectome Project; IQR, interquartile range; IVIM, intravoxel incoherent motion; LSQ, least-squares; MCMC, Markov chain Monte Carlo; NSA, number of signal averages; PDF, probability density function; RMSE, root mean squared error; ROI, region of interest; SD, standard deviation; SNR, signal-to-noise ratio; WM, white matter; WMH, white matter hyperintensity.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *NMR in Biomedicine* published by John Wiley & Sons Ltd.

## 1 | Introduction

Diffusion MRI (dMRI) is sensitive to the microscopic motion of water molecules. Microstructure imaging combines specialised dMRI acquisitions with tissue models to estimate quantitative parameters of tissue microstructure and microcirculation. This approach has been widely applied in neuroimaging—demonstrating estimation of neurite density [1], soma density [2], axonal diameter [3] and water exchange [4–6]—as well as being used in body imaging [7–9].

Such tissue models are typically fitted to dMRI data in each voxel separately to estimate the model parameters. This process has traditionally employed least-squares (LSQ) regression, although machine learning is emerging as an attractive alternative that can improve the precision [10] of microstructural maps. However, maps generated through voxel-by-voxel fitting are heavily affected by noise in individual voxels and do not optimally exploit the data, disregarding dependencies between voxels. In supervised machine learning approaches, the distribution of the training data effectively acts as a prior on the parameter estimates, providing some degree of noise robustness; however, this prior is fixed in advance by the choice of training data, and it has been shown that this choice influences the resulting estimates [11]. Methods that explicitly capture dependencies between neighbouring voxels demonstrate fast and robust fits—often using convolutional neural networks [12–15]—but these spatial regularisation techniques only capture patterns across predefined local windows, rather than capturing regional or global patterns.

Hierarchical Bayesian modelling (HBM) offers a promising alternative approach that breaks the assumption of independent voxels, while allowing for the incorporation of dependencies at various scales to enhance fitted maps [16]. HBM assumes a prior, typically Gaussian in this context, on model parameters over a region of interest (ROI). Rather than being user defined, the prior is estimated in a data-driven manner. This prior reduces sensitivity to noise in single voxels, leading to improved parameter maps.

Orton et al. [16] introduced the HBM approach for dMRI, formulating the model and a Markov chain Monte Carlo (MCMC) algorithm for fitting the intravoxel incoherent motion (IVIM) model to dMRI data, and demonstrated improved IVIM parameter maps in the liver. However, HBM has still primarily only been applied to a limited number of simple diffusion microstructure models. Notably, it has been used extensively with the IVIM model [16–26], as well as with a three-compartment model for placental MRI [27].

In our previous preliminary work [28], we introduced a generalised (i.e., not model specific) HBM framework in condensed form and applied it to the ball and stick model. In this paper, we build on our preliminary work, providing a complete derivation and demonstrating our general hierarchical Bayesian framework for microstructure modelling in two advanced neuroimaging applications. Specifically, our contributions in this work constitute: (i) evaluation of the framework in both simulated and in vivo data in two previously unexplored neuroimaging applications, namely, mean signal diffusion kurtosis imaging (DKI)

[29] and blood–brain barrier filter exchange imaging (BBB-FEXI) [5, 6] using the apparent exchange rate (AXR) model [4]; (ii) an assessment of the noise sensitivity of the regional priors; (iii) an analysis of the ability of different regional priors to resolve local parameter variations (e.g., lesions); (iv) a full derivation of the method for completeness, given in Appendix A; and (v) presentation of two open-source toolboxes, one Matlab (<https://github.com/e-powell/hbm-matlab>) and one Python (<https://github.com/e-powell/hbm-dmipy>). The HBM approach is shown to enhance parameter mapping relative to traditional voxel-by-voxel fitting, demonstrating that the advantages of HBM in dMRI extend beyond models that are linear combinations of exponentials.

## 2 | Methods

### 2.1 | Background: A General Hierarchical Bayesian Microstructure Model

The mathematical formulation for a general hierarchical Bayesian microstructural modelling framework is summarised here, building on our preliminary work [28] which extends the methods of Orton et al. [16]; see Appendix A for the full derivation and algorithm pseudocode.

Consider a general microstructure model (Appendix A.1) with  $C$  compartments and a set of microstructure-related parameters given by

$$\boldsymbol{\theta} = \left\{ \{f_c\}_{c=1}^{C-1}, \{x_j\}_{j=1}^J \right\}, \quad (1)$$

where  $f_c$  denotes compartment signal fractions and  $x_j$  all other parameters (e.g., diffusivities, orientations and radii). The signal equation in voxel  $i$  for an acquisition of length  $N$  can be written:

$$\mathbf{S}_i = S_0 g(\boldsymbol{\theta}_i, \mathbf{t}), \quad (2)$$

where  $\mathbf{S}_i = [S_1, \dots, S_N]^T$  is the noise-free model-predicted signal,  $\mathbf{t} = [t_1, \dots, t_N]^T$  is the set of acquisition parameters (e.g., diffusion encodings) and  $\boldsymbol{\theta}_i$  represents the underlying parameter values.

A Bayesian shrinkage prior (BSP) is assumed, following Orton et al. [16], but is here extended to accommodate multiple ROIs (Appendix A.2). The prior on parameters in voxel  $i$ , which is in ROI  $k$ , is therefore

$$p(\boldsymbol{\theta}_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \mathcal{N}(\boldsymbol{\theta}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \quad (3)$$

The posterior distribution and BSP parameters over each ROI,  $\boldsymbol{\mu}_k$  and  $\boldsymbol{\Sigma}_k$ , can be directly formulated [16]. The posterior distribution (Appendix A.3) is

$$p(\boldsymbol{\theta}_{1:I_k}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k | \mathbf{y}_{1:I_k}) \propto p(\mathbf{y}_{1:I_k} | \boldsymbol{\theta}_{1:I_k}) p(\boldsymbol{\theta}_{1:I_k} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) p(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (4)$$

where  $\boldsymbol{\theta}_{1:I_k} = \{\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_{I_k}\}$  are the microstructure parameters and  $\mathbf{y}_{1:I_k} = \{\mathbf{y}_1, \dots, \mathbf{y}_{I_k}\}$  the measured dMRI signal for all  $I_k$  voxels in ROI  $k$ , where each  $\mathbf{y}_i = [y_1, \dots, y_N]^T$  corresponds to  $\mathbf{S}_i$  with added

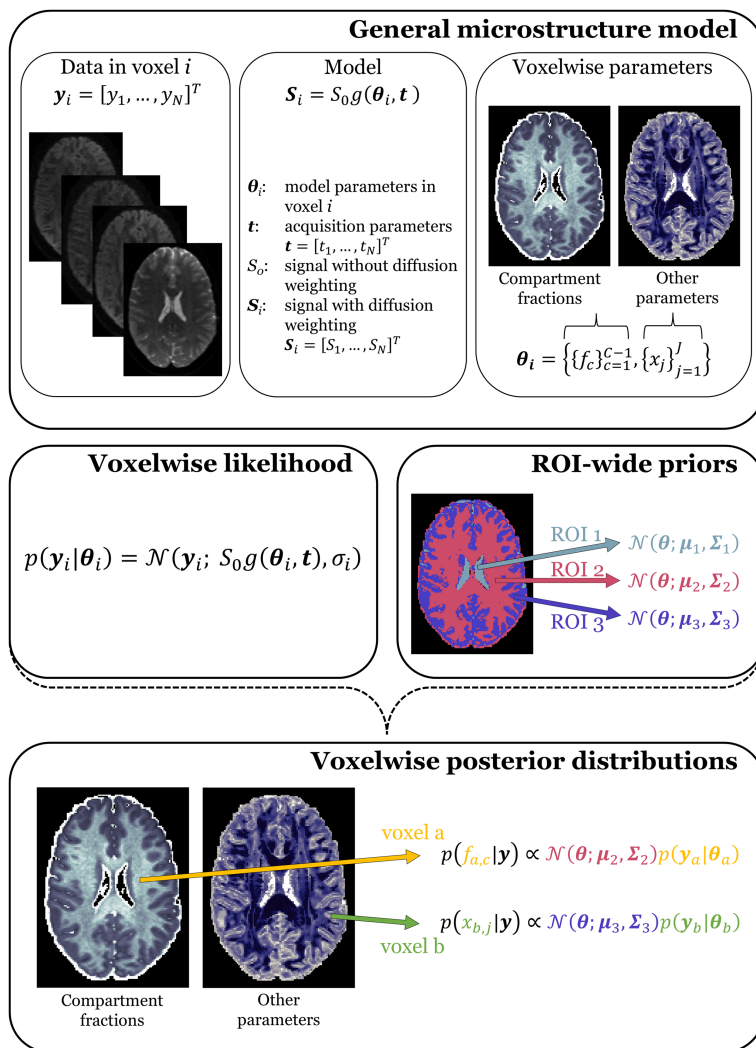
noise. Samples can be drawn from this posterior distribution with an MCMC algorithm. Once more following Orton et al. [16], the updates for  $\mu_k, \Sigma_k$  are Gibbs moves (Appendix A.5.1) and the updates for  $\theta_{1:k}$  are Metropolis–Hastings moves (Appendix A.5.2). Single parameter estimates for each voxel are obtained by taking the mean of the posterior distribution after the MCMC burn-in period (Appendix A.5.4).

A schematic of the full HBM framework is shown in Figure 1. The framework is demonstrated here in two distinct microstructure estimation problems: (i) the well-known mean signal DKI [29] experiment and (ii) the more recent innovation of filter exchange imaging (FEXI) [4] applied to the blood–brain barrier (BBB) [5, 6].

## 2.2 | Exemplar Microstructure Models

### 2.2.1 | DKI

DKI [29], an extension of diffusion tensor imaging (DTI), provides information on the non-Gaussianity of water diffusion in the brain, from which microstructural tissue information can be inferred. DKI modelling uses higher order cumulant expansion terms of the displacement distribution to extract kurtosis information. This requires higher  $b$ -values than needed for DTI, which, in addition to the increased model complexity, increases noise susceptibility and commonly leads to extreme values in the parameter maps when using conventional estimation methods.



**FIGURE 1** | Schematic of the hierarchical Bayesian modelling method. *Top panel.* A general microstructure model,  $g(\boldsymbol{\theta}_i, \mathbf{t})$ , maps the microstructure-related model parameters  $\boldsymbol{\theta}_i$  and the diffusion MRI acquisition parameters  $\mathbf{t}$  to the signal  $\mathbf{S}_i$  in voxel  $i$ , where the acquisition volume in a multivolume acquisition is indexed  $n = 1, \dots, N$ . The model parameters can be grouped according to parameter type as  $\boldsymbol{\theta} = \{ \{f_c\}_{c=1}^{C-1}, \{x_j\}_{j=1}^J \}$ , where  $\{f_c\}_{c=1}^{C-1}$  are the compartmental signal fractions, and  $\{x_j\}_{j=1}^J$  represent all other parameters. *Middle panel.* Definitions of the voxelwise likelihood function and ROI-wide Gaussian priors. *Bottom panel.* Definitions of the voxelwise parameter posteriors and examples of the corresponding parameter maps.

The spherically averaged DKI signal model is given by [30]

$$S = S_0 \exp\left(-bD + \frac{1}{6}b^2D^2K\right), \quad (5)$$

where  $S_0$  is the signal without diffusion encoding,  $D$  is the apparent diffusion coefficient (ADC) and  $K$  is the mean signal kurtosis.

### 2.2.2 | BBB-FEXI

Models parameterising water exchange rates across the BBB are highly noise sensitive, often requiring spatial or cross-subject data averaging [5, 6, 31–33]. Bayesian techniques offer a potential solution to this problem and could enable subject-specific voxel-level exchange rate estimates.

The BBB-FEXI technique [5, 6] utilises a double diffusion encoding sequence that was originally developed to measure the AXR of water across cell membranes [4, 34]. First, a low  $b$ -value diffusion filter selectively suppresses the signal from fast pseudo-diffusing intravascular spins. The intravascular signal recovers during the mixing time via exchange across the BBB, such that the signal measured after the second diffusion encoding block is a weighted combination of intravascular and extravascular signals, dependent on the mixing time and exchange rate. The FEXI AXR model [4] is defined as

$$S = \exp(-b \cdot D'(t_m)), \quad (6)$$

where

$$D'(t_m) = D[1 - \sigma \exp(-t_m \cdot \text{AXR})], \quad (7)$$

and  $D$  and  $D'(t_m)$  are the ADCs at equilibrium (i.e., filter inactive) and at each mixing time  $t_m$  with the filter active, respectively, and  $\sigma$  is the filter efficiency.

## 2.3 | Data

The HBM algorithm is demonstrated in both the DKI and FEXI AXR models first in simulations, where there is a known ground truth, before being applied in human data.

### 2.3.1 | DKI

**2.3.1.1 | Simulations.** Synthetic data were generated for two distinct ROIs representative of healthy human white matter (WM) and grey matter (GM). Parameter values were obtained by fitting the DKI model (Equation 5) using LSQ estimation to a randomly selected subject from the Human Connectome Project (HCP; details below) and calculating the mean and standard deviation (SD) of fitted parameters in WM and GM. Ground truth parameters were then drawn from normal distributions defined over these means and SDs. In WM, they were  $D_{wm} \sim \mathcal{N}(0.87, 0.29)\mu\text{m}^2/\text{ms}$ ,  $K_{wm} \sim \mathcal{N}(1.04, 0.41)$ ; in GM, they were  $D_{gm} \sim \mathcal{N}(1.12, 0.50)\mu\text{m}^2/\text{ms}$ ,  $K_{gm} \sim \mathcal{N}(0.63, 0.22)$ . The total matrix size was  $50 \times 50$ , with 1204 voxels in the ‘WM’ ROI and 1296 in the ‘GM’ ROI. Zero-mean Gaussian noise was added to the complex signal (i.e., to the real and imaginary components independently) to give a signal-to-noise ratio  $\text{SNR} = 20$  in the magnitude  $b = 0$  data. Signals were generated for  $b = 0, 1, 2, 3 \text{ ms}/\mu\text{m}^2$  with a number of signal averages  $\text{NSA} = 2, 9, 9, 9$ , respectively, mimicking multiple gradient directions (Table 1); the magnitude signals were averaged for each  $b$ -value.

**TABLE 1 | DKI protocol.** The DKI parameters shown are those used by the Human Connectome Project (HCP). Note that only 2, 9, 9, 9 gradient directions (29 volumes total) were used for the analysis here. The  $T_1$ -weighted MPRAGE data (resolution =  $0.7 \times 0.7 \times 0.7 \text{ mm}^3$ ; TR = 2400ms; TE = 2.14ms; flip angle =  $8^\circ$ ) were additionally utilised to obtain WM/GM segmentations. *BBB-FEXI protocol.* Cerebral small vessel disease (cSVD) subject data were acquired on a Philips 3T Ingenia Elition X system using an in-house double-diffusion encoding sequence. A  $T_1$ -weighted FFE (resolution =  $1 \times 1 \times 1 \text{ mm}^3$ ; TR = 6.8ms; TE = 3.1ms; flip angle =  $8^\circ$ ) was acquired for WM/GM segmentation, and a FLAIR (resolution =  $0.86 \times 0.86 \times 1 \text{ mm}^3$ ; TR = 4800ms; TE = 340ms; inversion time TI = 1650ms) for WM hyperintensity segmentation.

	DKI (HCP data)		BBB-FEXI (cSVD data)		
Resolution ( $\text{mm}^3$ )	1.25 × 1.25 × 1.25		3 × 3 × 5		
No. slices	111		18		
Repetition time, TR (ms)	5520		5000		
Echo time, TE (ms)	90		63		
Filter block echo time, $\text{TE}_f$ (ms)	—		42		
$b$ -values, $b$ ( $\text{ms}/\mu\text{m}^2$ )	0, 1, 2, 3	0, 0.25	0, 0.25	0, 0.25	0, 0.25
Gradient directions	18, 90, 90, 90	6, 6	6, 6	6, 6	6, 6
No. repetitions, $N_{\text{reps}}$	1	1, 1	1, 1	1, 1	1, 1
Mixing time, $t_m$ (ms)	—	16	16	200	400
Filter block $b$ values, $b_f$ ( $\text{ms}/\mu\text{m}^2$ )	—	0	0.25	0.25	0.25
Total no. volumes	288		48		

Abbreviations: BBB-FEXI, blood–brain barrier filter exchange imaging; DKI, diffusion kurtosis imaging; FFE, fast field echo; FLAIR, fluid-attenuated inversion recovery; MPRAGE, magnetisation-prepared rapid acquisition gradient echo.

**2.3.1.2 | Human Connectom Project Data.** Images for two subjects (in native space) were obtained from the publicly available data sets provided by the HCP [35] (WU-Minn Consortium; 1U54MH091657; funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research and by the McDonnell Center for Systems Neuroscience at Washington University) (Test Retest Data Release, release date: 1 March 2017, available online at [humanconnectome.org](http://humanconnectome.org)). To better demonstrate the efficacy of the HBM approach, which is most beneficial when applied to lower SNR data, the number of gradient directions was reduced to 2, 9, 9, 9 for  $b = 0, 1, 2, 3 \text{ ms}/\mu\text{m}^2$  respectively, in line with the NSA of the simulated data. WM/GM segmentations were obtained by first fitting the diffusion tensor then calculating ADC and fractional anisotropy (FA) maps with the MRtrix3 [36] functions *dwi2tensor* and *tensor2metric*, respectively. GM masks were defined by selecting voxels with  $\text{ADC} < 2 \mu\text{m}^2/\text{ms}$  and  $\text{FA} < 0.15$  and WM masks by selecting voxels with  $\text{ADC} < 2 \mu\text{m}^2/\text{ms}$  and  $\text{FA} > 0.15$ .

Average ROI sizes across the two subjects for WM/GM (one slice) were 2757/7626 voxels, respectively.

### 2.3.2 | BBB-FEXI

**2.3.2.1 | Simulations.** Model parameters were generated for two ROIs, again representing healthy WM/GM with a total matrix size of  $50 \times 50$ . Parameter values were drawn from normal distributions with means and SDs taken from literature values for healthy subjects [6]. In the ‘WM’ ROI, these were  $D_{wm} \sim \mathcal{N}(0.90, 0.04) \mu\text{m}^2/\text{ms}$ ,  $\sigma_{wm} \sim \mathcal{N}(0.14, 0.02)$ ,  $\text{AXR}_{wm} \sim \mathcal{N}(2.00, 0.30) \text{s}^{-1}$ ; in the ‘GM’ ROI they were  $D_{gm} \sim \mathcal{N}(1.20, 0.08) \mu\text{m}^2/\text{ms}$ ,  $\sigma_{gm} \sim \mathcal{N}(0.18, 0.03)$ ,  $\text{AXR}_{gm} \sim \mathcal{N}(1.40, 0.35) \text{s}^{-1}$ . Six signals were generated for each parameter combination using the acquisition parameters in Table 1; zero-mean Gaussian noise was added to each complex signal to give  $\text{SNR} = 30$  in the equilibrium magnitude data (i.e., with the filter inactive,  $b_f = 0 \text{ ms}/\mu\text{m}^2$ , the minimum mixing time,  $t_m = 16 \text{ ms}$ , and without diffusion weighting,  $b = 0 \text{ ms}/\mu\text{m}^2$ ), comparable with the human data (see below and Table 1). Each set of six noisy magnitude signals were then averaged (i.e.,  $\text{NSA} = 6$ ) to mimic the six gradient directions used in the human data.

**2.3.2.2 | Cerebral Small Vessel Disease (cSVD) Subjects.** Two exemplar subjects from an ongoing cSVD study were analysed to provide proof of principle in the presence of pathology. Patients were recruited from Lancashire Teaching Hospitals NHS Foundation Trust and Salford Royal NHS Foundation Trust. The study was approved by Yorkshire and The Humber (Leeds West) Research Ethics Committee (21/YH/0119), the Health Research Authority and local research governance panels.

Subjects were scanned on a Philips 3T Ingenia Elition X system using a 32-channel head coil. BBB-FEXI,  $T_1$ -weighted and fluid-attenuated inversion recovery (FLAIR) data were acquired; see Table 1 for protocols. Susceptibility and eddy current distortions in the BBB-FEXI data were corrected using FSL *topup*

and *eddy* [37, 38]. The  $T_1$ -weighted data (in native space) were segmented using FreeSurfer [39] to produce WM/GM ROIs. The  $T_1$ -weighted and FLAIR data were used to produce WM hyperintensity (WMH) masks using the Lesion Segmentation Toolbox [40] in SPM12 with a locally established threshold of 0.3. The  $T_1$ -weighted and FLAIR images were registered to the BBB-FEXI data using FSL FLIRT [41] and the respective warps used to propagate the WM/GM/WMH masks to the BBB-FEXI data. Average ROI sizes across the two subjects for WM/GM/WMH (18 slices) were 8373/10,202/1825 voxels, respectively.

## 2.4 | Model Fitting

### 2.4.1 | LSQ Estimation

LSQ fitting was performed on all data sets using in-house code to provide priors for the HBM method and to provide conventional fitting results against which the HBM results could be compared. The LSQ algorithm (*fminsearchbnd* in Matlab R2024b) was initialised using 25 and 27 starting values for the DKI and FEXI AXR models, respectively, uniformly distributed between parameter constraints to avoid local minima. The DKI model (Equation 5; estimating  $D$  and  $K$ ) fits were constrained such that  $D \in (0.1, 3.5) \mu\text{m}^2/\text{ms}$  and  $K \in (0, 3)$ ; parameter constraints for the FEXI AXR model (Equations 6 and 7; estimating  $D$ ,  $\sigma$  and AXR, with  $D'(t_m)$  calculated from the data) were  $D \in (0.1, 3.5) \mu\text{m}^2/\text{ms}$ ,  $\sigma \in (0, 1)$ ,  $\text{AXR} \in (0, 50) \text{s}^{-1}$ . The AXR upper bound was chosen as this is approximately the fastest exchange rate that is likely to be detectable given the shortest mixing time.

### 2.4.2 | Hierarchical Bayesian Estimation

The HBM framework implemented in Matlab (R2024b) was used for all experiments. Parameter values were initialised with the voxelwise LSQ fit and constrained with the same bounds (Section 2.4.1); the ROI priors were initialised using the mean and standard deviation of the LSQ fits in each ROI. For the simulated and human DKI data, the MCMC algorithm was run for  $N_s = 1 \times 10^5$  steps. For the human BBB-FEXI data the MCMC algorithm was similarly run with  $N_s = 1 \times 10^5$  steps; for the simulated BBB-FEXI data, the MCMC chain length was increased to  $N_s = 4 \times 10^5$  steps to improve the conditioning of the regional priors given the reduced number of voxels (2500 voxels) compared with the human data ( $\sim 20,000$  voxels). In all cases, the burn-in period was defined as the first  $N_{burn} = N_s/2$  steps. Weights were updated every 100 steps during the first half of the burn-in period (Appendix A.5.3), which was heuristically found to be sufficient to sample the posterior distributions. Two ROIs were used for the regional priors—one each for WM/GM—for the DKI data (simulated and human) and the simulated BBB-FEXI data; a third ROI was utilised in cSVD patient BBB-FEXI data for WMH.

Parameter posterior distributions and representative statistics were calculated from the MCMC samples after the burn-in. Voxelwise parameter maps were generated using the mean of the posterior distributions in each voxel  $i$ :

$$\hat{\theta}_i \simeq \frac{1}{N_s - N_{burn}} \sum_{s=N_{burn}+1}^{N_s} \theta_i^{(s)}, \quad (8)$$

where  $\hat{\theta}_i$  is the expected value of a single parameter (Equation 1) and  $\theta_i^{(s)}$  the value sampled at each MCMC step.

Summary statistics for the regional priors were calculated as

$$\hat{\mu}_{\theta,k} \simeq \frac{1}{N_s - N_{burn}} \sum_{s=N_{burn}+1}^{N_s} \mu_{\theta,k}^{(s)}, \quad (9)$$

where  $\hat{\mu}_{\theta,k}$  is the prior mean of parameter  $\theta$  in the  $k$ th ROI and  $\mu_{\theta,k}^{(s)}$  the prior mean estimated at each MCMC step.

## 2.5 | Evaluation

A series of metrics were used to test the ability of the methods to (i) accurately and precisely recover ground truth parameter values under variable SNR conditions (simulations) and (ii) obtain contrast between different ROIs (simulations and human data).

### 2.5.1 | Accuracy and Precision

In simulations, the precision of fitted parameters was evaluated using the root mean squared error (RMSE):

$$\text{RMSE} = \left( \frac{1}{n} \sum_{i=1}^n (\theta_i - \hat{\theta}_i)^2 \right)^{\frac{1}{2}}, \quad (10)$$

where  $\theta_i$  is the ground truth of a single model parameter in voxel  $i$  and  $\hat{\theta}_i$  the estimated value. Bias (accuracy) was defined as

$$\text{bias}(\hat{\theta}, \theta) = \frac{1}{n} \sum_{i=1}^n (\hat{\theta}_i - \theta_i). \quad (11)$$

Maps of the percent relative error between ground truth and fitted values were also generated, with the error for voxel  $i$  given by

$$\epsilon_i = 100 \times \frac{\hat{\theta}_i - \theta_i}{\theta_i}. \quad (12)$$

The data were also smoothed prior to LSQ fitting to evaluate the performance of more simple filtering and non-local means methods against the HBM approach; full details are provided in Supporting Information S1 and S2.

### 2.5.2 | Contrast-to-Noise Ratio (CNR)

The CNR between two ROIs in parameter maps was quantified according to [42]:

$$\text{CNR} = \frac{|\bar{\mu}_{\theta,k_1} - \bar{\mu}_{\theta,k_2}|}{\left( \text{IQR}_{\theta,k_1}^2 + \text{IQR}_{\theta,k_2}^2 \right)^{\frac{1}{2}}}, \quad (13)$$

where  $\bar{\mu}$  and IQR are the median and interquartile range of voxelwise parameter estimates in each ROI.

### 2.5.3 | Noise Sensitivity of Regional Priors

The noise sensitivity of the regional priors was investigated in the DKI and BBB-FEXI simulated data. Data were simulated with additional noise characteristics to give three data sets per model: for the DKI model, data were simulated with  $\text{SNR}_{dki} = 10, 20, 30$ ; for the FEXI AXR model, data were simulated with  $\text{SNR}_{axr} = 10, 30, 50$ .

### 2.5.4 | Resolving Local Parameter Variations


The ability of the HBM algorithm, and its dependence on the defined regional priors, to resolve small local variations in AXR—as is hypothesised to occur in WMH regions in cSVD subjects—was investigated in the human data. The HBM model was run three times using different prior configurations: (i) one whole brain global prior ( $k = 1$ ); (ii) two regional priors, one each for WM/GM ( $k = 2$ ); and (iii) three regional priors, with an additional prior for WMH ( $k = 3$ ). The CNR between WM and WMH was evaluated as a surrogate for lesion discrimination quality for the LSQ output and for each HBM output. Simulations were performed post hoc to support the in vivo experiments; details are provided in Supporting Information S3.

## 3 | Results

### 3.1 | Accuracy and Precision

Figure 2 shows example parameter maps for the simulated data (with  $\text{SNR}_{dki} = 20$  and  $\text{SNR}_{axr} = 30$ ), estimated using the LSQ and HBM approaches. Qualitatively, HBM parameter maps were less affected by noise and showed improved parameter estimation over LSQ maps; the HBM approach also outperformed more simple filtering (Supporting Information S1) and non-local means (Supporting Information S2) methods. The effect of this noise on parameter accuracy and precision is visualised in Figure 3, which shows percent error maps and correlations with ground truth values. Improvements in the estimation of kurtosis and AXR—the most challenging parameters to fit in the DKI and FEXI AXR models, respectively—were most apparent: the number of extreme values in the LSQ kurtosis fit was reduced in the HBM fit (Figure 3A; see also Table 2), while the large and numerous errors in the LSQ FEXI AXR outputs were substantially reduced using the HBM approach (Figure 3B). Correlations with ground truth values were also improved using the HBM method: the correlation coefficient increased from  $R_{lsq} = 0.87$  to  $R_{hbm} = 0.93$  for kurtosis (Figure 3C) and from  $R_{lsq} = 0.14$  to  $R_{hbm} = 0.41$  for AXR (Figure 3D).

The RMSE (Table 2) was also improved using the HBM approach, decreasing by 25% for kurtosis (from  $\text{RMSE}_{lsq} = 0.20$  to  $\text{RMSE}_{hbm} = 0.15$ ) and by 89% for AXR (from  $\text{RMSE}_{lsq} = 19.2 \text{ s}^{-1}$  to  $\text{RMSE}_{hbm} = 2.1 \text{ s}^{-1}$ ). Bias was minimal in the DKI model for the LSQ and HBM outputs of both parameters; for the FEXI AXR model, the HBM fit reduced bias in the AXR estimation by 88% (from  $\text{bias}_{lsq} = 8.5 \text{ s}^{-1}$  to  $\text{bias}_{hbm} = 1.0 \text{ s}^{-1}$ ).



**FIGURE 2** | Parameter maps for the simulated data, showing ground truth (GT) values (top row) along with the least-squares (LSQ) outputs (middle row) and hierarchical Bayesian modelling (HBM) outputs (bottom row). LSQ output maps were generally noisier than the HBM output maps. (A) DKI model, showing the apparent diffusion coefficient ( $D$ ) and kurtosis ( $K$ ) maps. (B) FEXI AXR model, showing the apparent diffusion coefficient ( $D$ ), filter efficiency ( $\sigma$ ) and apparent exchange rate (AXR) maps.

Minimal or no improvements were seen in the HBM diffusivity estimates of either model as it is typically a stable parameter to fit, even when using LSQ methods.

Parameter estimation using the LSQ and HBM approaches took approximately 6 and 64 min, respectively, for the DKI model, and 6 min and 5.5 h, respectively, for the FEXI AXR model.

### 3.2 | CNR

With the HBM method, CNR improved in the simulated data by a factor of 1.2 for kurtosis (from  $\text{CNR}_{lsq} = 0.6$  to  $\text{CNR}_{hbm} = 0.7$ ) and by a factor of 6 for AXR (from  $\text{CNR}_{lsq} = 0.1$  to  $\text{CNR}_{hbm} = 0.6$ ) (Table 2). This improvement was most visually striking in the AXR parameter maps (Figure 2B), where ROI boundaries were obscured by high noise in the LSQ output maps but were well resolved in the HBM output maps.

### 3.3 | Noise Sensitivity of Regional Priors

Figure 4 shows the ground truth distribution of parameter values, along with the parameter distributions obtained from the LSQ and HBM outputs and the HBM prior distributions (averaged over the MCMC chain after the burn-in period). For the DKI model (Figure 4A), the LSQ and HBM outputs and

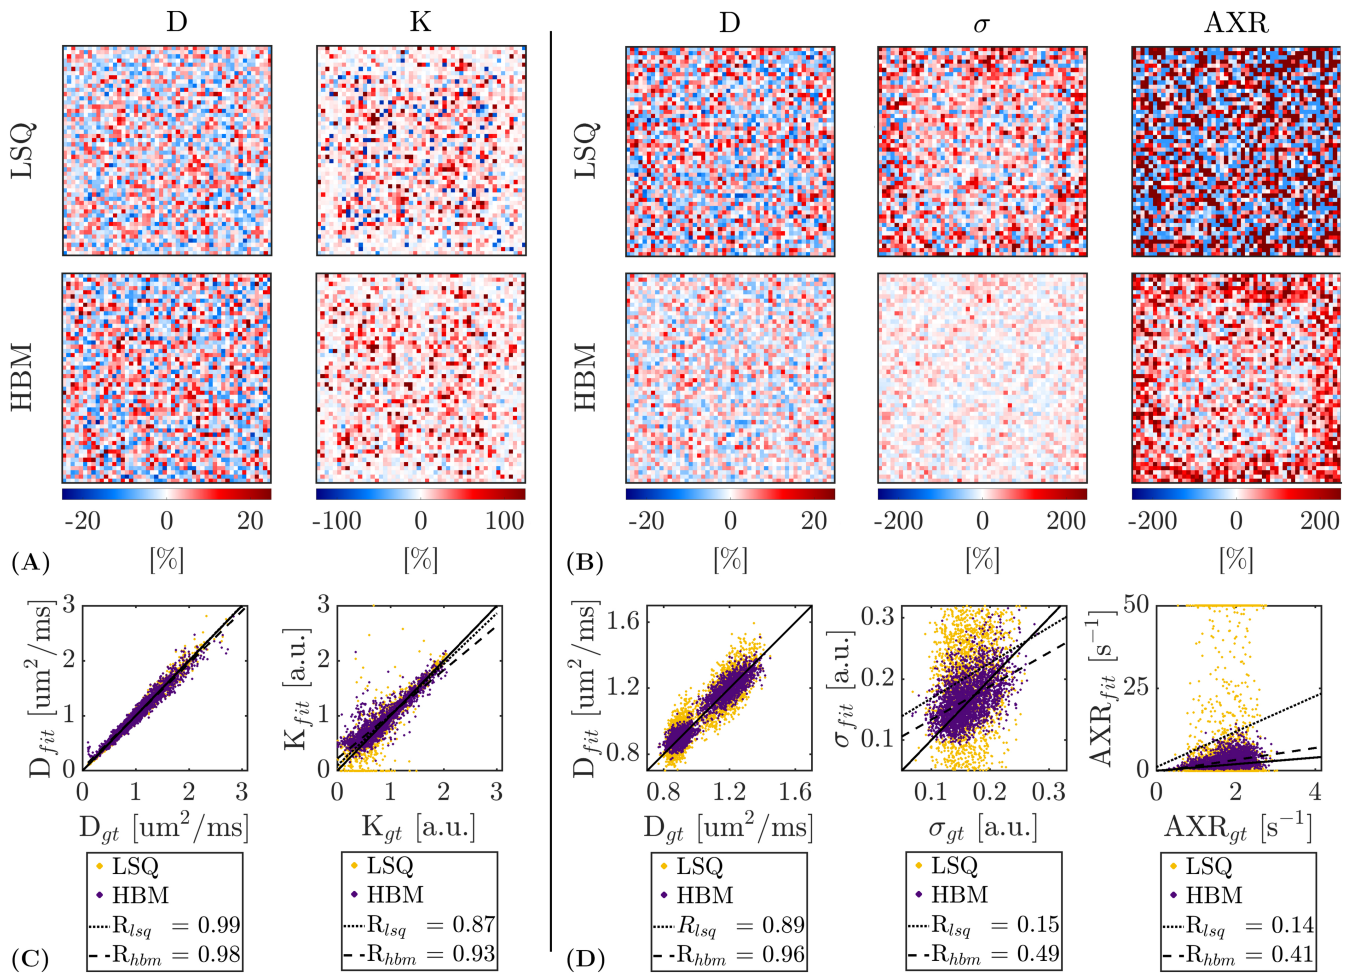
the HBM priors were well matched to the ground truth across SNR levels. For the FEXI AXR model (Figure 4B), the HBM priors became less informative at lower SNRs, displaying wide, flat profiles; however, this weak influence still helped resolve the WM/GM ROIs, although biases were introduced. LSQ AXR outputs provided no contrast between the separate ROIs at low SNRs, with parameter estimates clustered at the bounds for both ROIs. The information content of the HBM priors improved at higher SNRs, more closely resembling the ground truth distributions.

### 3.4 | Healthy Volunteer Data

Figure 5 shows parameter maps for the healthy volunteer DKI data, estimated using the LSQ method and the HBM method with  $k = 2$  regional priors. A reduction in the number of extreme fits in the kurtosis estimation was shown using the HBM approach (Figure 5, inset), in accordance with the simulated data (Figure 2A). A reduction in diffusivity was observed near the ventricles in the HBM-derived maps.

### 3.5 | Exemplar cSVD Subject Data

AXR parameter maps of the two exemplar cSVD subjects are shown in Figure 6 for the LSQ approach and the HBM approach

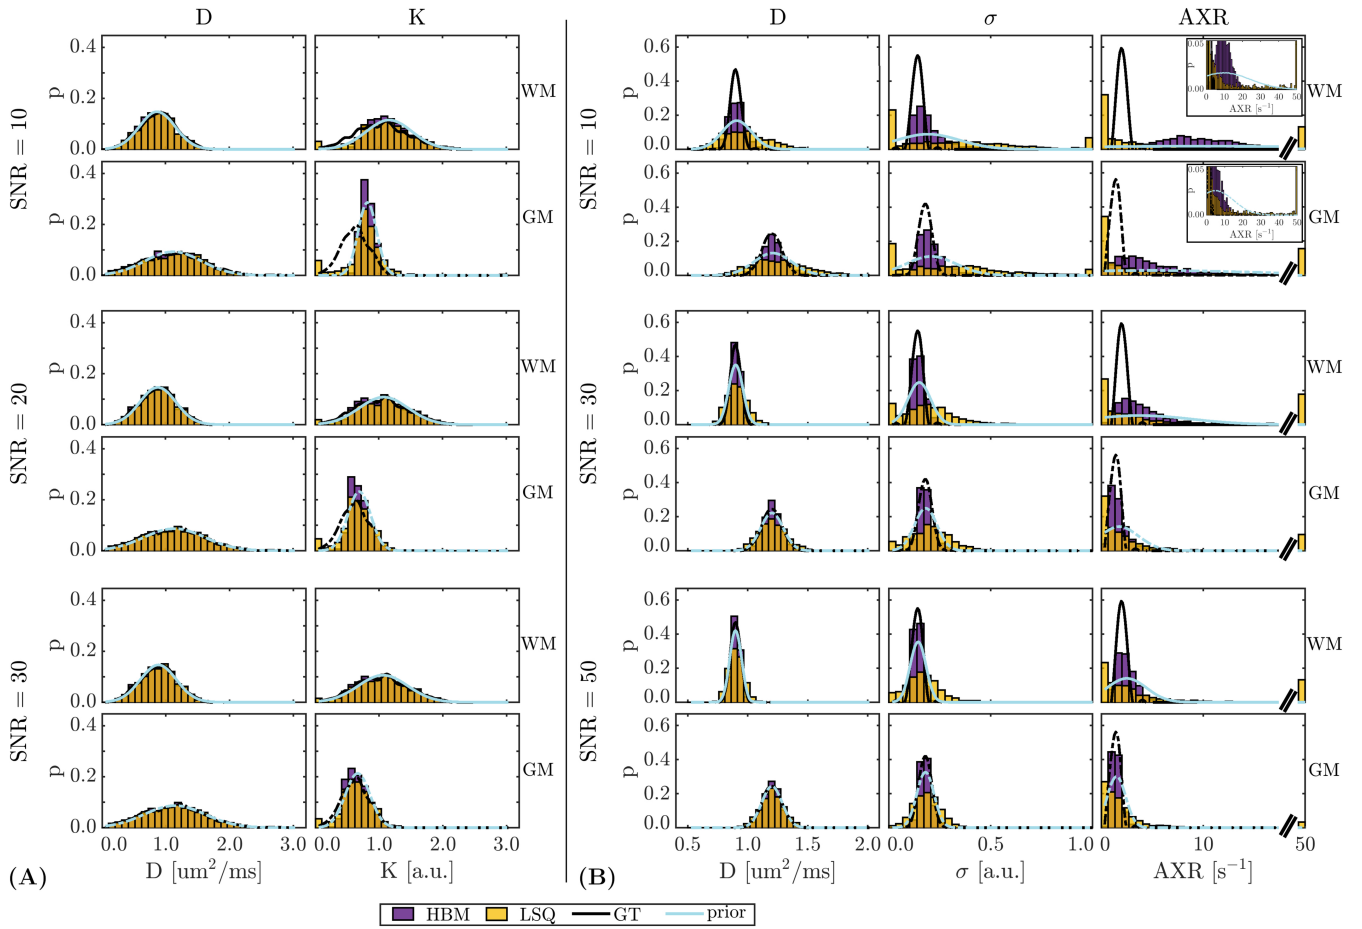


**FIGURE 3** | Error metrics for the least-squares (LSQ) and hierarchical Bayesian modelling (HBM) outputs in simulated data. The HBM outputs typically showed reduced errors and improved correlations with ground truth values relative to the LSQ outputs. (A) Percent error maps for the apparent diffusion coefficient ( $D$ ) and kurtosis ( $K$ ) parameters of the DKI model. (B) Percent error maps for the apparent diffusion coefficient ( $D$ ), filter efficiency ( $\sigma$ ) and apparent exchange rate (AXR) parameters of the FEXI AXR model. (C) Correlations with ground truth (GT) values for the LSQ- and HBM-derived DKI model parameters. (D) Correlations with ground truth values for the LSQ- and HBM-derived FEXI AXR model parameters.

**TABLE 2** | Error metrics for the simulated DKI (SNR = 20) and BBB-FEXI (SNR = 30) data. Extreme fits were defined as parameter values within 1% of the fitting bounds; values in brackets indicate the total proportion of voxels containing one or more parameters with extreme fits.

		DKI model			AXR model		
		$D$ ( $\mu\text{m}^2/\text{ms}$ )	$K$ (a.u.)		$D$ ( $\mu\text{m}^2/\text{ms}$ )	$\sigma$ (a.u.)	AXR ( $\text{s}^{-1}$ )
RMSE	LSQ	$6.5 \times 10^{-2}$	0.20		$8.1 \times 10^{-2}$	$13.2 \times 10^{-2}$	19.2
	HBM	$7.9 \times 10^{-2}$	0.15		$4.6 \times 10^{-2}$	$3.5 \times 10^{-2}$	2.1
Bias	LSQ	$0.1 \times 10^{-2}$	0.03		$0.6 \times 10^{-2}$	$4.3 \times 10^{-2}$	8.5
	HBM	$0.2 \times 10^{-2}$	0.05		$0.1 \times 10^{-2}$	$0.6 \times 10^{-2}$	1.0
CNR	GT	0.3	0.7		2.6	0.8	1.0
	LSQ	0.3	0.6		1.5	0.1	0.1
	HBM	0.3	0.7		2.8	0.6	0.6
Extreme fits (%)	LSQ	0.0	2.6	(2.6)	0.0	7.4	32.9 (39.6)
	HBM	0.0	0.0	(0.0)	0.0	0.0	0.0 (0.0)

Abbreviations: AXR, apparent exchange rate; CNR, contrast-to-noise ratio; DKI, diffusion kurtosis imaging; FEXI, filter exchange imaging; RMSE, root mean squared error.



**FIGURE 4** | Impact of SNR on parameters fitted using least-squares (LSQ) (yellow) and hierarchical Bayesian modelling (HBM) (purple) approaches. The histograms show voxelwise parameter estimates from the simulated ‘white matter’ (WM) and ‘grey matter’ (GM) regions. Ground truth distributions are indicated in black; HBM priors are shown in blue. (A) DK1 model, showing the apparent diffusion coefficient ( $D$ ) and kurtosis ( $K$ ) parameter distributions. (B) FEXI AXR model, showing the apparent diffusion coefficient ( $D$ ), filter efficiency ( $\sigma$ ) and apparent exchange rate (AXR) parameter distributions. Inset figures for SNR = 10 are re-scaled to better show the prior distributions.

with  $k = 3$  regional priors; other parameters are shown for each subject, respectively, in Figures S4.1 and S4.2. Noise was substantially reduced in the HBM parameter maps, revealing an increase in AXR in WMH relative to both WM and GM; by comparison, high noise levels in the LSQ maps obscured any differences in the WMH AXR.

Box plots of parameter estimates in WM, GM and WMH are shown in Figure 7A,B for each of the different prior configurations ( $k = 1, 2, 3$ ). Without a separate prior for WMH ( $k = 2$ ), median WM and WMH values were similar, with  $\tilde{A}\tilde{X}R_{k=2}^{wm} = 5.9, 4.5 \text{ s}^{-1}$  and  $\tilde{A}\tilde{X}R_{k=2}^{wmh} = 5.5, 8.0 \text{ s}^{-1}$  for subjects S1, S2, respectively. When using a separate prior for WMH ( $k = 3$ ), a substantial increase in the median WMH AXR compared with the median WM AXR was observed:  $\tilde{A}\tilde{X}R_{k=3}^{wm} = 4.9, 3.8 \text{ s}^{-1}$  and  $\tilde{A}\tilde{X}R_{k=3}^{wmh} = 9.1, 15.7 \text{ s}^{-1}$  for subjects S1, S2, respectively. Median WMH AXR values estimated using the LSQ approach were considerably lower at  $\tilde{A}\tilde{X}R_{lsq}^{wmh} = 2.4, 3.6 \text{ s}^{-1}$  for subjects S1, S2 (with  $\tilde{A}\tilde{X}R_{lsq}^{wm} = 1.4, 1.3 \text{ s}^{-1}$ ) but showed a wider distribution across voxels than all of the HBM-derived estimates.

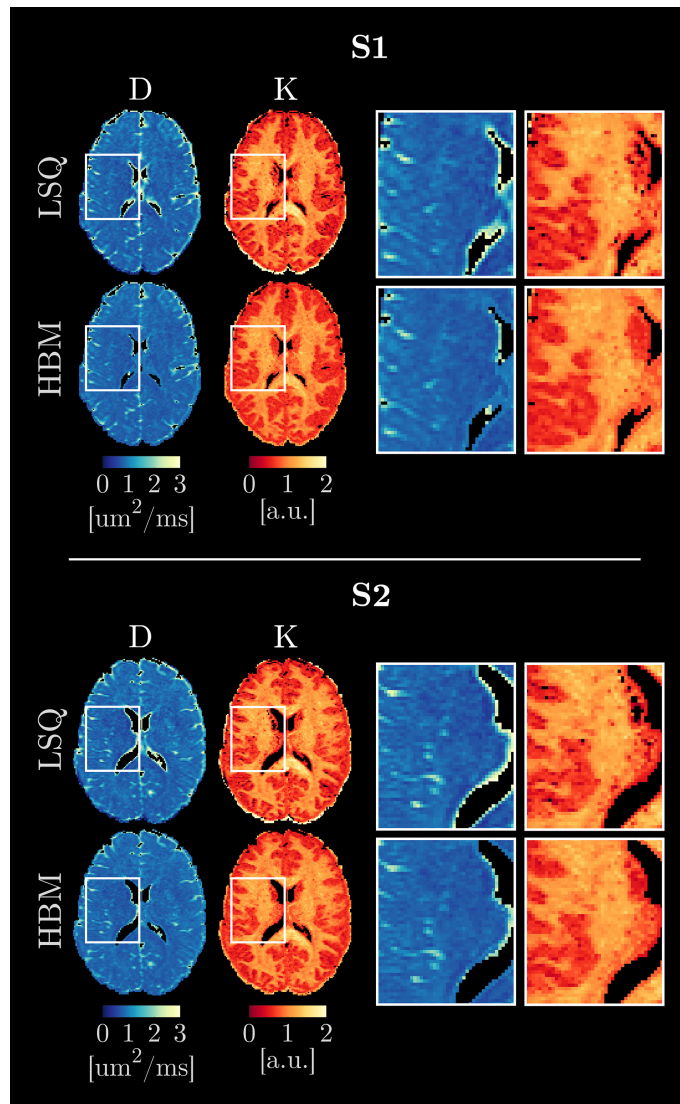
Figure 7C shows the CNR between WM and WMH for each prior configuration. The CNR was higher for  $k = 3$

( $\text{CNR}_{k=3} = 0.2, 0.5$  for subjects S1, S2, respectively) than for  $k = 1, 2$  (e.g.,  $\text{CNR}_{k=2} = 0.02, 0.2$  for subjects S1, S2), demonstrating the potential for differentiating between WM and WMH AXR values when using a separate WMH prior. The CNR between WM and WMH was minimal when using the LSQ approach ( $\text{CNR}_{lsq} = 0.04, 0.08$  for subjects S1, S2). Post hoc simulations showed similar behaviour (Supporting Information S3).


The average SNR in the brain, estimated using the mean and SD of the six repetitions with  $b_f = 0 \text{ ms}/\mu\text{m}^2$ ,  $t_m = 16 \text{ ms}$ ,  $b = 0 \text{ ms}/\mu\text{m}^2$ , was SNR = 26, 32 in subjects S1, S2, respectively.

## 4 | Discussion

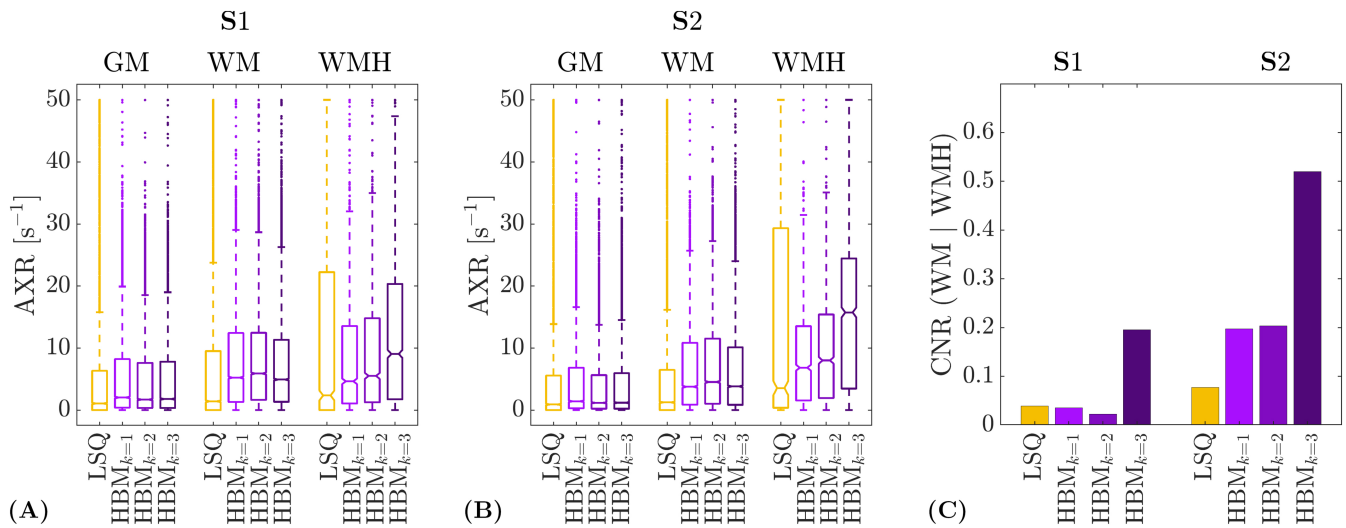
Our generalised algorithm, which enables HBM fitting for any dMRI microstructural model and incorporates arbitrary parameter constraints and regional priors, is demonstrated here in two previously unexplored neuroimaging applications. The algorithm is implemented in both Matlab (R2024b) and Python (v3.9.20) and is publicly available at <https://github.com/e-powell/hbm-matlab> and <https://github.com/e-powell/hbm-dmipy>, respectively; in Python, the Dmipy [43] software package is utilised and extended.



**FIGURE 5** | Diffusion kurtosis imaging (DKI) parameter maps for the two test cases (top panel, subject S1; bottom panel, subject S2), showing the least-squares (LSQ) and hierarchical Bayesian modelling (HBM) (with  $k = 2$  regional priors: white matter and grey matter) outputs. Inset images highlight regions of reduced diffusivity ( $D$ ) around the ventricles in the HBM maps (indicating a potential reduction in CSF contamination) and reduced noise in the HBM kurtosis ( $K$ ) maps.



**FIGURE 6** | Apparent exchange rate (AXR) parameter maps for cerebral small vessel disease (cSVD) subject S1 (top row) and subject S2 (bottom row), derived using least-squares (LSQ) and hierarchical Bayesian modelling (HBM) (with  $k = 3$  regional priors: white matter, grey matter and white matter hyperintensities) approaches. Noise was substantially reduced in the HBM parameter maps, revealing regions of increased AXR within cSVD white matter hyperintensities (outlined in red) that were obscured in the LSQ maps.



**FIGURE 7** | Impact of different regional priors in cerebral small vessel disease (cSVD) subjects. (A) Boxplots of voxelwise apparent exchange rate (AXR) estimates are shown for cSVD subject S1 in grey matter (GM), white matter (WM) and white matter hyperintensities (WMH) for the least-squares (LSQ) and hierarchical Bayesian modelling (HBM) approaches; the number of regional priors used for the HBM approach is denoted by the subscript, where  $k = 1$  indicates one global prior,  $k = 2$  indicates two priors (WM and GM) and  $k = 3$  indicates three priors (WM, GM and WMH). Using a separate prior for WMH (i.e.,  $k = 3$ ) lead to a substantial increase in the median AXR in WMH, while median WMH AXR values for  $k = 1, 2$  were closer to WM AXR values. (B) As for (A), but for cSVD subject S2. (C) Contrast-to-noise ratio (CNR) between WM and WMH for subjects S1 and S2. The CNR was higher for  $k = 3$  than for  $k = 1, 2$ , indicating superior discrimination of WMH AXR alterations when using a separate WMH prior.

Note that all results reported in this work used the Matlab implementation; examples are provided alongside the implementation on GitHub to aid users in generalising the code for other microstructural models. The HBM algorithm was shown to outperform conventional LSQ approaches in both simulated and human data for the diffusion-based DKI and FEXI AXR models.

The greatest improvements were observed for the least stable parameters in the LSQ fit (kurtosis in the DKI model; AXR in the FEXI AXR model), particularly for the more complex FEXI AXR model. Complex models often suffer from parameter degeneracy—where multiple combinations of model parameters give rise to similar MRI signals—and high sensitivity to noise; consequently, voxelwise fitting techniques often return values associated with erroneous local minima in the fitting objective function. Exploiting information across a region—as is the role of the prior distributions in the HBM method—can reduce the sensitivity of a model to these effects, as the prior information increases the likelihood of selecting parameter combinations more similar to the rest of the region.

This is exemplified in the DKI model (Figures 2A and 5), where the HBM method demonstrated potential reductions in partial volume effects compared with the LSQ approach. For example, areas of high diffusivity around the ventricles in the LSQ maps were reduced in the HBM maps, potentially indicating a reduction in partial volume effects arising from CSF contamination by the HBM method, while small areas of noise in the LSQ-derived kurtosis maps—which are demonstrably fitting artefacts in the simulated data—were also removed in the HBM maps.

The benefits of the HBM approach were most apparent for the FEXI AXR model (Figures 2B and 6), likely reflecting the higher noise susceptibility (and subsequently noisier LSQ fits)

of FEXI AXR model parameters relative to DKI model parameters. Experiments with simulated data in this work provide confidence that, for the noise levels expected in human data [6], HBM can still infer the true FEXI AXR model parameter values, as demonstrated by the error maps in Figure 3. The cSVD case studies (Figure 6) best exemplify the HBM advantages: A clear contrast between AXR values in WMHs and surrounding brain tissue was revealed in HBM maps, while noise in the LSQ maps obscured any regional variation (reflecting the simulation results in Figure 2B). The AXR was generally higher in WMH (although this was variable) than in surrounding tissues, suggesting increased BBB damage; this has been reported in studies using dynamic contrast-enhanced MRI [44]. As this work was designed to introduce and demonstrate proof of principle of the HBM algorithm, a detailed interpretation of the observed WMH AXR values in relation to underlying pathophysiology is out of scope here; however, it is noted that cSVD is a diffuse and heterogeneous disease process with multiple factors contributing to small vessel and BBB damage [45, 46], potentially explaining the AXR variability within WMH. Previous studies have demonstrated altered BBB water exchange in cSVD subjects using a diffusion-prepared pseudo-continuous arterial spin labelling sequence [47, 48], although a decrease in water exchange rate was observed while an increase was detected here. It may be important to note that these previous studies included genetic causes of cSVD (representing a minority of cSVD cases) in contrast to the two cSVD subjects included in this work who are older adults with vascular risk factors; careful consideration of the differences between MRI acquisition techniques and factors such as relaxation time differences (known to influence the FEXI AXR model [4, 6]) in WMH, along with increased sample sizes, are also required before drawing conclusions. However, without loss of generalisability, these data

provide evidence that localised pathology can be resolved and detected using the HBM method, even in high noise conditions.

A limitation of the HBM method is that, in high noise conditions, it is possible that the prior (and subsequently parameter estimates derived using HBM) may tend towards the mean of the parameter range if individual voxels do not contain sufficient information for inference. Noise sensitivity analysis on the regional priors demonstrated that, for the models chosen here, some priors became heavily biased at low SNR (e.g., AXR at SNR = 10; Figure 4B) but maintained enough information content to resolve the separate ROIs and did not reduce parameter distributions to the mean of the parameter bounds.

A related issue under high noise conditions—and one of the primary concerns of any hierarchical method—is that regions with underlying parameter values distinct from the regional prior (e.g., lesions) may become obscured. For example, in the cSVD patient data, using a separate regional prior for WMH substantially improved discrimination between WM and WMH AXR estimates (Figure 7). The unbalanced ROI sample sizes between WM (~10,000 voxels) and WMH (~1200 voxels) in combination with the large difference between underlying WM/WMH AXR values (as observed for  $k = 3$  regional priors) could explain this finding: a prior combining two distinct ROIs may be dominated by the larger ROI (i.e., WM here) as voxels in the smaller ROI (i.e., WMH here) may be treated as outliers by the prior, such that greater biases will be incurred in the smaller ROI for larger differences between underlying regional parameter means. For example, AXR prior means (Equation 9) were substantially different for WM and WMH when using  $k = 3$  regional priors ( $\hat{\mu}_{axr,wm} = 8.0, 6.8 s^{-1}$  and  $\hat{\mu}_{axr,wmh} = 13.8, 16.6 s^{-1}$  for subjects S1, S2, respectively), while the WM AXR prior mean for  $k = 2$  (i.e., no WMH regional prior) was skewed towards the WM prior mean observed for  $k = 3$  ( $\hat{\mu}_{axr,wm} = 9.3, 8.6 s^{-1}$ ). This supports the suggestion that, when WM and WMH are combined under one regional prior, the likelihood of selecting parameter combinations that reflect the true WMH distribution may be reduced. Post hoc simulations (Supporting Information S3) supported the use of a separate regional prior for pathological tissues.

The above considerations suggest a requirement for identification of localised pathology before applying the HBM method, particularly in the context of noisier data with more complex models. This could be addressed by developing an adaptive hierarchical model that selects the number of ROIs during inference to determine the appropriate number of regional priors. Identifying additional independent ROIs—that is, spatially localised voxel clusters with similar parameter values—and updating voxel ROI membership accordingly at each MCMC step may improve parameter estimation in small distinct regions. However, small ROIs are more susceptible to noise, and hence, the trade-off between regional prior size and noise bias needs to be carefully managed. Simulations could help elucidate any such sample size dependencies. An adaptive model could additionally remove the need for obtaining tissue segmentations from separate MRI acquisitions, as was required here for the cSVD data.

Another avenue of future work is to develop non-Gaussian priors. It was assumed here that a Gaussian prior was suitable for all model parameters; however, particularly for models that include fibre orientation parameters (e.g., NODDI [1]), a uniform prior may instead be more appropriate. A further assumption of the HBM method is that signal noise is Gaussian, but this may not be suitable for very low signal voxels (SNR  $\lesssim 2$ ) [49], where the noise distribution becomes Rician. Parallel imaging further alters the noise characteristics [50], which may show non-central chi or Rayleigh distribution properties at low-to-moderate SNRs (SNR  $\leq 30$  [51]); more advanced noise reduction methods, for example, denoising complex channel data prior to image reconstruction [52], could be included to reduce these effects. Future work could also extend HBM to approaches that calculate spectra rather than single parameter values by reformulating the algorithm to place priors over multidimensional spectra, although this may necessitate methods such as variational Bayes to manage the associated computational complexity. It would also be prudent in future work to perform comparisons between different MCMC model fitting methods (e.g., Harms and Roebroek [53]), which may be able to reduce the computational burden of MCMC sampling and minimise the adjustments and tuning traditionally required when applying MCMC sampling to new microstructure models.

## 5 | Conclusions

Our generalised HBM and MCMC algorithm, which enables parameter inference for dMRI signal models with arbitrary parameter constraints, is shown to improve the accuracy, precision, CNR and detection of localised pathology over the conventional voxelwise LSQ approach, providing a new option for reducing erroneous fits in complex dMRI models. The algorithm is open source and implemented in Matlab (<https://github.com/e-powell/hbm-matlab>) and Python (<https://github.com/e-powell/hbm-dmipy>).

---

### Author Contributions

**Elizabeth Powell:** conceptualization; methodology; investigation; formal analysis; writing – original draft. **Mark Maskery:** investigation; resources; writing – review and editing. **Hedley C. A. Emsley:** investigation; resources; writing – review and editing. **Laura M. Parkes:** investigation; resources; writing – review and editing; funding acquisition. **Geoff J. M. Parker:** conceptualization; writing – review and editing; funding acquisition. **Paddy J. Slator:** conceptualization; methodology; investigation; writing – original draft.

### Acknowledgements

Thanks to the radiographers at Manchester University imaging facilities for facilitating the cSVD data acquisition. Thanks to Dr David Higgins of Philips Healthcare MR Clinical Science for supporting this work.

### Funding

This study was supported by EPSRC (Grants EP/S031510/1 and EP/M020533/1) and the Alzheimer's Society Heather Corrie Impact Fund (Grant/Award number: 577 [AS-PG-21-045]).

## Conflicts of Interest

G.J.M.P. receives salary from, is a director of and a shareholder in Bioxydyn Limited, a company with an interest in imaging biomarkers. G.J.M.P. is also a director of and a shareholder in Queen Square Analytics Limited and is a director of and a shareholder in Quantitative Imaging Limited, companies with an interest in imaging biomarkers. G.J.M.P. has current grant funding from Eli Lilly and GSK and received other support from Philips and Siemens Healthineers. E.P. has received compensation for consultation services to Bioxydyn Limited.

## Data Availability Statement

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

## References

1. H. Zhang, T. Schneider, C. A. Wheeler-Kingshott, and D. C. Alexander, "NODDI: Practical In Vivo Neurite Orientation Dispersion and Density Imaging of the Human Brain," *NeuroImage* 61, no. 4 (2012): 1000–1016, <https://doi.org/10.1016/j.neuroimage.2012.03.072>.
2. M. Palombo, A. Ianus, M. Guerrerri, et al., "SANDI: A Compartment-Based Model for Non-Invasive Apparent Soma and Neurite Imaging by Diffusion MRI," *NeuroImage* 215 (2020): 116835, <https://doi.org/10.1016/j.neuroimage.2020.116835>.
3. J. Veraart, D. Nunes, U. Rudrapatna, et al., "Noninvasive Quantification of Axon Radii Using Diffusion MRI," *eLife* 9 (2020): e49855, <https://doi.org/10.7554/eLife.49855>.
4. S. Lasic, M. Nilsson, J. Lätt, F. Ståhlberg, and D. Topgaard, "Apparent Exchange Rate Mapping With Diffusion MRI," *Magnetic Resonance in Medicine* 66, no. 2 (2011): 356–365, <https://doi.org/10.1002/mrm.22782>.
5. R. Bai, Z. Li, C. Sun, Y. C. Hsu, H. Liang, and P. Basser, "Feasibility of Filter-Exchange Imaging (FEXI) in Measuring Different Exchange Processes in Human Brain," *NeuroImage* 219, no. June (2020): 117039, <https://doi.org/10.1016/j.neuroimage.2020.117039>.
6. E. Powell, Y. Ohene, M. Battiston, B. R. Dickie, L. M. Parkes, and G. J. M. Parker, "Blood-Brain Barrier Water Exchange Measurements Using FEXI: Impact of Modeling Paradigm and Relaxation Time Effects," *Magnetic Resonance in Medicine* 90, no. 1 (2023): 34–50, <https://doi.org/10.1002/mrm.29616>.
7. E. Panagiotaki, S. Walker-Samuel, B. Siow, et al., "Noninvasive Quantification of Solid Tumor Microstructure Using VERDICT MRI," *Cancer Research* 74, no. 7 (2014): 1902–1912, <https://doi.org/10.1158/0008-5472.CAN-13-2511>.
8. E. Panagiotaki, R. W. Chan, N. Dikaios, et al., "Microstructural Characterization of Normal and Malignant Human Prostate Tissue With Vascular, Extracellular, and Restricted Diffusion for Cytometry in Tumours Magnetic Resonance Imaging," *Investigative Radiology* 50, no. 4 (2015): 218–227, <https://doi.org/10.1097/RLL.0000000000000115>.
9. Y. T. Li, J. P. Cercueil, J. Yuan, W. Chen, R. Loffroy, and Y. X. J. Wng, "Liver Intravoxel Incoherent Motion (IVIM) Magnetic Resonance Imaging: A Comprehensive Review of Published Data on Normal Values and Applications for Fibrosis and Tumor Evaluation," *Quantitative Imaging in Medicine and Surgery* 7, no. 1 (2017): 59–78, <https://doi.org/10.21037/qims.2017.02.03>.
10. E. R. Sabidussi, S. Klein, B. Jeurissen, and D. H. J. Poot, "dtiRIM: A Generalisable Deep Learning Method for Diffusion Tensor Imaging," *NeuroImage* 269 (2023): 119900, <https://doi.org/10.1016/j.neuroimage.2023.119900>.
11. N. G. Gyori, M. Palombo, C. A. Clark, H. Zhang, and D. C. Alexander, "Training Data Distribution Significantly Impacts the Estimation of Tissue Microstructure With Machine Learning," *Magnetic Resonance in Medicine* 87, no. 2 (2022): 932–947, <https://doi.org/10.1002/mrm.29014>.
12. Z. Li, T. Gong, Z. Lin, et al., "Fast and Robust Diffusion Kurtosis Parametric Mapping Using a Three-Dimensional Convolutional Neural Network," *IEEE Access* 7 (2019): 71398–71411, <https://doi.org/10.1109/ACCESS.2019.2919241>.
13. Q. Tian, B. Bilgic, Q. Fan, et al., "DeepDTI: High-Fidelity Six-Direction Diffusion Tensor Imaging Using Deep Learning," *NeuroImage* 219 (2020): 117017, <https://doi.org/10.1016/j.neuroimage.2020.117017>.
14. S. D. Vasylechko, S. K. Warfield, O. Afacan, and S. Kurugol, "Self-Supervised IVIM DWI Parameter Estimation With a Physics Based Forward Model," *Magnetic Resonance in Medicine* 87, no. 2 (2022): 904–914, <https://doi.org/10.1002/mrm.28989>.
15. S. HashemizadehKolowri, R. R. Chen, G. Adluru, and E. V. R. DiBella, "Jointly Estimating Parametric Maps of Multiple Diffusion Models From Undersampled q-Space Data: A Comparison of Three Deep Learning Approaches," *Magnetic Resonance in Medicine* 87, no. 6 (2022): 2957–2971, <https://doi.org/10.1002/mrm.29162>.
16. M. R. Orton, D. J. Collins, D. M. M. Koh, and M. O. Leach, "Improved Intravoxel Incoherent Motion Analysis of Diffusion Weighted Imaging by Data Driven Bayesian Modeling," *Magnetic Resonance in Medicine* 71, no. 1 (2013): 411–420, <https://doi.org/10.1002/mrm.24649>.
17. S. Barbieri, O. F. Donati, J. M. Froehlich, and H. C. Thoeny, "Impact of the Calculation Algorithm on Biexponential Fitting of Diffusion-Weighted MRI in Upper Abdominal Organs," *Magnetic Resonance in Medicine* 75, no. 5 (2016): 2175–2184, <https://doi.org/10.1002/mrm.25765>.
18. P. T. While, "A Comparative Simulation Study of Bayesian Fitting Approaches to Intravoxel Incoherent Motion Modeling in Diffusion-Weighted MRI," *Magnetic Resonance in Medicine* 78, no. 6 (2017): 2373–2387, <https://doi.org/10.1002/mrm.26598>.
19. G. R. Spinner, C. von Deuster, K. C. Tezcan, C. T. Stoeck, and S. Kozerke, "Bayesian Intravoxel Incoherent Motion Parameter Mapping in the Human Heart," *Journal of Cardiovascular Magnetic Resonance* 19, no. 1 (2017): 85, <https://doi.org/10.1186/s12968-017-0391-1>.
20. O. Gustafsson, M. Montelius, G. Starck, and M. Ljungberg, "Impact of Prior Distributions and Central Tendency Measures on Bayesian Intravoxel Incoherent Motion Model Fitting," *Magnetic Resonance in Medicine* 79, no. 3 (2018): 1674–1683, <https://doi.org/10.1002/mrm.26783>.
21. O. J. Gurney-Champion, R. Klaassen, M. Froeling, et al., "Comparison of Six Fit Algorithms for the Intra-Voxel Incoherent Motion Model of Diffusion-Weighted Magnetic Resonance Imaging Data of Pancreatic Cancer Patients," *PLoS ONE* 13, no. 4 (2018): e0194590, <https://doi.org/10.1371/journal.pone.0194590>.
22. O. Jalnefjord, M. Andersson, M. Montelius, et al., "Comparison of Methods for Estimation of the Intravoxel Incoherent Motion (IVIM) Diffusion Coefficient ( $D$ ) and Perfusion Fraction ( $f$ )," *Magnetic Resonance Materials in Physics, Biology and Medicine* 31, no. 6 (2018): 715–723, <https://doi.org/10.1007/s10334-018-0697-5>.
23. T. Koopman, R. Martens, O. J. Gurney-Champion, et al., "Repeatability of IVIM Biomarkers From Diffusion-Weighted MRI in Head and Neck: Bayesian Probability Versus Neural Network," *Magnetic Resonance in Medicine* 85, no. 6 (2021): 3394–3402, <https://doi.org/10.1002/mrm.28671>.
24. G. R. Spinner, C. Federau, and S. Kozerke, "Bayesian Inference Using Hierarchical and Spatial Priors for Intravoxel Incoherent Motion MR Imaging in the Brain: Analysis of Cancer and Acute Stroke," *Medical Image Analysis* 73 (2021): 102144, <https://doi.org/10.1016/j.media.2021.102144>.
25. E. Scalco, G. Rizzo, and A. Mastropietro, "The Quantification of IntraVoxel Incoherent Motion MRI Maps Cannot Preserve Texture Information: An Evaluation Based on Simulated and *In-Vivo* Images," *Computers in Biology and Medicine* 154 (2023): 106495, <https://doi.org/10.1016/j.compbiomed.2022.106495>.

26. O. Jalnefjord and I. M. Bjrkman-Burtscher, "Comparison of Methods for Intravoxel Incoherent Motion Parameter Estimation in the Brain From Flow-Compensated and Non-Flow-Compensated Diffusion-Encoded Data," *Magnetic Resonance in Medicine* 92, no. 1 (2024): 303–318, <https://doi.org/10.1002/mrm.30042>.
27. D. Flouri, D. Owen, R. Aughwane, et al., "Improved Fetal Blood Oxygenation and Placental Estimated Measurements of Diffusion-Weighted MRI Using Data-Driven Bayesian Modeling," *Magnetic Resonance in Medicine* 83, no. 6 (2020): 2160–2172, <https://doi.org/10.1002/mrm.28075>.
28. E. Powell, M. Battocchio, C. S. Parker, and P. J. Sclator, "Generalised Hierarchical Bayesian Microstructure Modelling for Diffusion MRI," in *Computational Diffusion MRI*. Lecture Notes in Computer Science, vol. 13006, eds. S. Cetin-Karayumak, D. Christiaens, M. Figini, et al. (Springer International Publishing, 2021), 36–47, <https://doi.org/10.1007/978-3-030-87615-9-4>.
29. J. H. Jensen, J. A. Helpert, A. Ramani, H. Lu, and K. Kaczynski, "Diffusional Kurtosis Imaging: The Quantification of Non-Gaussian Water Diffusion by Means of Magnetic Resonance Imaging," *Magnetic Resonance in Medicine* 53, no. 6 (2005): 1432–1440, <https://doi.org/10.1002/mrm.20508>.
30. R. N. Henriques, S. N. Jespersen, and N. Shemesh, "Microscopic Anisotropy Misestimation in Spherical-Mean Single Diffusion Encoding MRI," *Magnetic Resonance in Medicine* 81, no. 5 (2019): 3245–3261, <https://doi.org/10.1002/mrm.27606>.
31. J. A. Wells, B. Siow, M. F. Lythgoe, and D. L. Thomas, "Measuring Biexponential Transverse Relaxation of the ASL Signal at 9.4 T to Estimate Arterial Oxygen Saturation and the Time of Exchange of Labeled Blood Water Into Cortical Brain Tissue," *Journal of Cerebral Blood Flow and Metabolism* 33, no. 2 (2013): 215–224, <https://doi.org/10.1038/jcbfm.2012.156>.
32. X. Shao, S. J. Ma, M. Casey, L. D'Orazio, J. M. Ringman, and D. J. J. Wang, "Mapping Water Exchange Across the Bloodbrain Barrier Using 3D Diffusion-Prepared Arterial Spin Labeled Perfusion MRI," *Magnetic Resonance in Medicine* 81, no. 5 (2019): 3065–3079, <https://doi.org/10.1002/mrm.27632>.
33. E. Powell, B. R. Dickie, Y. Ohene, M. Maskery, G. J. M. Parker, and L. M. Parkes, "Bloodbrain Barrier Water Exchange Measurements Using Contrast-Enhanced ASL," *NMR in Biomedicine* 36, no. 11 (2023): e5009, <https://doi.org/10.1002/nbm.5009>.
34. M. Nilsson, J. Lätt, D. Van Westen, et al., "Noninvasive Mapping of Water Diffusional Exchange in the Human Brain Using Filter-Exchange Imaging," *Magnetic Resonance in Medicine* 69, no. 6 (2013): 1573–1581, <https://doi.org/10.1002/mrm.24395>.
35. D. C. Van Essen, K. Ugurbil, E. Auerbach, et al., "The Human Connectome Project: A Data Acquisition Perspective," *NeuroImage* 62, no. 4 (2012): 2222–2231, <https://doi.org/10.1016/j.neuroimage.2012.02.018>.
36. J. D. Tournier, R. Smith, D. Raffelt, et al., "MRtrix3: A Fast, Flexible and Open Software Framework for Medical Image Processing and Visualisation," *NeuroImage* 202, no. August (2019): 116137, <https://doi.org/10.1016/j.neuroimage.2019.116137>.
37. J. L. R. Andersson, S. Skare, and J. Ashburner, "How to Correct Susceptibility Distortions in Spin-Echo Echo-Planar Images: Application to Diffusion Tensor Imaging," *NeuroImage* 20, no. 2 (2003): 870–888, [https://doi.org/10.1016/S1053-8119\(03\)00336-7](https://doi.org/10.1016/S1053-8119(03)00336-7).
38. S. M. Smith, M. Jenkinson, M. W. Woolrich, et al., "Advances in Functional and Structural MR Image Analysis and Implementation as FSL," *NeuroImage* 23 (2004): 208–219, <https://doi.org/10.1016/j.neuroimage.2004.07.051>.
39. B. Fischl, "FreeSurfer," *NeuroImage* 62, no. 2 (2012): 774, <https://doi.org/10.1016/j.neuroimage.2012.01.021>.
40. P. Schmidt, C. Gaser, M. Arsic, et al., "An Automated Tool for Detection of FLAIR-Hyperintense White-Matter Lesions in Multiple Sclerosis," *NeuroImage* 59, no. 4 (2012): 3774–3783, <https://doi.org/10.1016/j.neuroimage.2011.11.032>.
41. M. Jenkinson and S. Smith, "A Global Optimisation Method for Robust Affine Registration of Brain Images," *Medical Image Analysis* 5, no. 2 (2001): 143–156, [https://doi.org/10.1016/S1361-8415\(01\)00036-6](https://doi.org/10.1016/S1361-8415(01)00036-6).
42. F. Grussu, T. Schneider, H. Zhang, D. C. Alexander, and C. A. M. WheelerKingshott, "Neurite Orientation Dispersion and Density Imaging of the Healthy Cervical Spinal Cord *in Vivo*," *NeuroImage* 111 (2015): 590–601, <https://doi.org/10.1016/j.neuroimage.2015.01.045>.
43. R. H. Fick, D. Wassermann, and R. Deriche, "The Dmipy Toolbox: Diffusion MRI Multi-Compartment Modeling and Microstructure Recovery Made Easy," *Frontiers in Neuroinformatics* 13, no. October (2019): 1–26, <https://doi.org/10.3389/fninf.2019.00064>.
44. M. J. Thrippleton, W. H. Backes, S. Sourbron, et al., "Quantifying Blood-Brain Barrier Leakage in Small Vessel Disease: Review and Consensus Recommendations," *Alzheimer's and Dementia* 15, no. 6 (2019): 840–858, <https://doi.org/10.1016/j.jalz.2019.01.013>.
45. J. M. Wardlaw, "Blood-Brain Barrier and Cerebral Small Vessel Disease," *Journal of the Neurological Sciences* 299, no. 1 (2010): 66–71, <https://doi.org/10.1016/j.jns.2010.08.042>.
46. J. M. Wardlaw, C. Smith, and M. Dichgans, "Small Vessel Disease: Mechanisms and Clinical Implications," *Lancet Neurology* 18, no. 7 (2019): 684–696, [https://doi.org/10.1016/S1474-4422\(19\)30079-1](https://doi.org/10.1016/S1474-4422(19)30079-1).
47. Y. Li, Y. Ying, T. Yao, et al., "Decreased Water Exchange Rate Across Bloodbrain Barrier in Hereditary Cerebral Small Vessel Disease," *Brain* 146, no. 7 (2023): 3079–3087, <https://doi.org/10.1093/brain/awac500>.
48. Y. Ying, Y. Li, T. Yao, et al., "Heterogeneous Blood-Brain Barrier Dysfunction in Cerebral Small Vessel Diseases," *Alzheimer's & Dementia* 20, no. 7 (2024): 4527–4539, <https://doi.org/10.1002/alz.13874>.
49. H. Gudbjartsson and S. Patz, "The Rician Distribution of Noisy MRI Data," *Magnetic Resonance in Medicine* 34, no. 6 (1995): 910–914, <https://doi.org/10.1002/mrm.1910340618>.
50. O. Dietrich, J. G. Raya, S. B. Reeder, M. Ingrisch, M. F. Reiser, and S. O. Schoenberg, "Influence of Multichannel Combination, Parallel Imaging and Other Reconstruction Techniques on MRI Noise Characteristics," *Magnetic Resonance Imaging* 26, no. 6 (2008): 754–762, <https://doi.org/10.1016/J.MRI.2008.02.001>.
51. M. Bouhrara and R. G. Spencer, "Fisher Information and Cramér-Rao Lower Bound for Experimental Design in Parallel Imaging," *Magnetic Resonance in Medicine* 79, no. 6 (2018): 3249–3255, <https://doi.org/10.1002/mrm.26984>.
52. E. Powell, T. Schneider, M. Battiston, et al., "SENSE EPI Reconstruction With 2D Phase Error Correction and Channel-Wise Noise Removal," *Magnetic Resonance in Medicine* 88, no. 5 (2022): 2157–2166, <https://doi.org/10.1002/mrm.29349>.
53. R. L. Harms and A. Roebroek, "Robust and Fast Markov Chain Monte Carlo Sampling of Diffusion MRI Microstructure Models," *Frontiers in Neuroinformatics* 12 (2018): 97, <https://doi.org/10.3389/fninf.2018.00097>.
54. P. Neal and G. Roberts, "Optimal Scaling for Partially Updating MCMC Algorithms," *Annals of Applied Probability* 16, no. 2 (2006): 475–515, <https://doi.org/10.1214/105051605000000791>.

### Supporting Information

Additional supporting information can be found online in the Supporting Information section. Supporting Information

## Appendix A

### Full Derivation of the General Hierarchical Bayesian Microstructure Model

In this section, our generalised HBM algorithm, which is extended from Orton et al. [16] to include the formulation for arbitrary upper/lower parameter bounds and the extension to multiple ROIs, together with the MCMC algorithm used for inference are motivated and introduced in full.

#### A.1 | General Microstructure Model

Consider a general multicompartment model of  $C$  compartments, with an underlying microstructure-related parameter set  $\theta$ :

$$\theta = \left\{ \{f_c\}_{c=1}^{C-1}, \{x_j\}_{j=1}^J \right\}, \quad (\text{A1})$$

where  $f_c$  denotes compartment signal fractions and  $x_j$  all other parameters (e.g., diffusivities, orientations and radii). Assuming relaxation times are fixed across compartments, the signal fractions sum to one ( $\sum_{c=1}^C f_c = 1$ ), meaning that  $f_C$  is not a free parameter but fixed as  $f_C = 1 - \sum_{c=1}^{C-1} f_c$ . Grouping model parameters  $\theta$  by type, as in Equation (A1), allows for notational convenience in enforcing this additional constraint on the signal fraction priors (e.g., in Equation A8).

The general multi-compartment model comprises a function, denoted  $g$ , that relates the underlying tissue-related parameters  $\theta$  and dMRI acquisition parameters (e.g.,  $b$ -value and gradient orientations) to the measured signal intensity. For a dMRI acquisition with  $N$  diffusion weightings, the set of acquisition parameters is defined as  $\mathbf{t} = [t_1, \dots, t_N]^T$  and the measured signal in voxel  $i$  as  $\mathbf{S}_i = [S_1, \dots, S_N]^T$ . The signal equation across a whole dMRI acquisition can thus be written:

$$\mathbf{S}_i = S_0 g(\theta_i, \mathbf{t}), \quad (\text{A2})$$

where  $S_0$  is the signal intensity without diffusion weighting and  $\theta_i$  are the model parameters in voxel  $i$ . The RHS of Equation (A2) is a column vector where the  $n$ th element,  $S_0 g(\theta_i, t_n)$ , is the model-predicted signal for diffusion weighting  $t_n$  given model parameters  $\theta_i$ .

The experimentally measured signal for voxel  $i$  in the presence of Gaussian noise is hence modelled as

$$\mathbf{y}_i = \mathbf{S}_i + \boldsymbol{\epsilon}_i, \quad (\text{A3})$$

where  $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$  is the measured signal and  $\boldsymbol{\epsilon}_i \sim \mathcal{N}(0, \sigma_i^2 \mathcal{I})$  is the noise and where  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  is a multivariate normal distribution with mean  $\boldsymbol{\mu}$  and covariance matrix  $\boldsymbol{\Sigma}$ , with  $\sigma_i$  a scalar giving the noise standard deviation in voxel  $i$  and  $\mathcal{I}$  the  $N \times N$  identity matrix.

The likelihood for the measurement in voxel  $i$  is

$$p(\mathbf{y}_i | \theta_i, S_0, \sigma_i^2) = (2\pi\sigma_i^2)^{-\frac{N}{2}} \times \exp\left(-\frac{1}{2\sigma_i^2} \sum_{n=1}^N (y_n - S_0 g(\theta_i, t_n))^2\right) \quad (\text{A4})$$

Orton et al. [16] demonstrated that the ‘nuisance parameters’  $S_0$  and  $\sigma_i^2$  can be marginalised out from Equation (A4) to give the following marginalised likelihood:

$$p(\mathbf{y}_i | \theta_i) \propto [\mathbf{y}_i^T \mathbf{y}_i - (\mathbf{y}_i^T \mathbf{g}_i)^2 / \mathbf{g}_i^T \mathbf{g}_i]^{-N/2}, \quad (\text{A5})$$

where  $\mathbf{g}_i = [g(\theta_i, t_1), \dots, g(\theta_i, t_N)]$  are the model-predicted signals for voxel  $i$ .

#### A.2 | BSPs

Orton et al. [16] used a multivariate Gaussian BSP on the IVIM model parameters over a single user-defined ROI, defined as

$$p(\theta_i | \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \mathcal{N}(\theta_i; \boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (\text{A6})$$

where  $\boldsymbol{\mu}$  and  $\boldsymbol{\Sigma}$  are estimated from the data and effectively encode the mean and covariance, respectively, of all model parameters across the ROI.

The method of Orton et al. [16] is extended here to account for multiple ROIs. Given  $K$  ROIs, denoted  $k = 1, \dots, K$ , each voxel is a member of exactly one ROI. Each ROI  $k$  has its own Gaussian BSP distribution. For voxel  $i$  in ROI  $k$ , the prior on the parameters  $\theta_i$  is then

$$p(\theta_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \mathcal{N}(\theta_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (\text{A7})$$

where  $\boldsymbol{\mu}_k$  is a vector whose elements encode the prior means of the parameters in ROI  $k$ ,  $\boldsymbol{\Sigma}_k$  is their covariance and  $\mathcal{N}(\theta; \boldsymbol{\mu}, \boldsymbol{\Sigma})$  denotes the multivariate normal probability density function (PDF). Again, it is emphasised that  $\boldsymbol{\mu}_k$  and  $\boldsymbol{\Sigma}_k$  are estimated from the data.

To generalise from Orton et al.’s [16] two-compartment model to an arbitrary multi-compartment model, all signal fractions must sum to one. This is enforced here (following Harms and Roebroeck [53]) by modifying the priors to

$$p(\theta_i | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \begin{cases} \mathcal{N}(\theta_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) & \text{if } \sum_{c=1}^{C-1} f_c \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A8})$$

To complete the model, hyperpriors are defined on all  $\boldsymbol{\mu}_k$ ’s and  $\boldsymbol{\Sigma}_k$ ’s as non-informative Jeffrey’s priors:

$$p(\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K, \boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K) = \prod_{k=1}^K |\boldsymbol{\Sigma}_k|^{-1/2} \quad (\text{A9})$$

#### A.3 | Posterior Distributions

Each ROI has an independent posterior distribution. If ROI  $k$  comprises  $I_k$  voxels, its posterior distribution can be written as [16]:

$$p(\theta_{1:I_k}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k | \mathbf{y}_{1:I_k}) \propto p(\mathbf{y}_{1:I_k} | \theta_{1:I_k}) p(\theta_{1:I_k} | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) p(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (\text{A10})$$

where  $\theta_{1:I_k} = \{\theta_1, \dots, \theta_{I_k}\}$  are the parameters and  $\mathbf{y}_{1:I_k} = \{\mathbf{y}_1, \dots, \mathbf{y}_{I_k}\}$  the dMRI data for all  $I_k$  voxels in the ROI. Substituting in Equations (A5), (A7) and (A9) gives

$$p(\theta_{1:I_k}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k | \mathbf{y}_{1:I_k}) \propto \left( \prod_{i=1}^{I_k} [\mathbf{y}_i^T \mathbf{y}_i - (\mathbf{y}_i^T \mathbf{g}_i)^2 / \mathbf{g}_i^T \mathbf{g}_i]^{-N/2} \right) \times \left( \prod_{i=1}^{I_k} \mathcal{N}(\theta_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right) |\boldsymbol{\Sigma}_k|^{-1/2}, \quad (\text{A11})$$

from which samples can be drawn with an MCMC algorithm (see Appendix A.5).

#### A.4 | Parameter Transforms

Microstructure model fitting needs to enforce physically reasonable minimum and maximum values of parameters; for example, diffusivities need to be positive. Here, the transforms used by Orton et al. [16] are generalised to enable arbitrary minimum and maximum constraints. For a single parameter  $\theta$ , its transform is defined as

$$\theta' = \log(\theta - \theta_{\min}) - \log(\theta_{\max} - \theta), \quad (\text{A12})$$

which maps the interval  $(\theta_{\min}, \theta_{\max})$  to  $\mathbb{R}$ , allowing updates from unconstrained distributions in both the Gibbs and Metropolis–Hastings algorithms. By defining the Bayesian prior on the transformed parameter

$\theta'$ ,  $\theta$  is therefore constrained between  $\theta_{min}$  and  $\theta_{max}$ . The associated inverse transform is

$$\theta = \frac{\exp(\theta') \theta_{max} + \theta_{min}}{1 + \exp(\theta')}, \quad (\text{A13})$$

For notational convenience, from here on,  $\theta$  denotes the transformed parameter.

## A.5 | MCMC Algorithm

The MCMC algorithm is derived here and given as pseudocode in Algorithm A1.

### A.5.1 | ROI-Wide Parameters

Following Orton et al. [16], the MCMC updates for the ROI-wide prior parameters  $\boldsymbol{\mu}_{1:k}$  and  $\boldsymbol{\Sigma}_{1:k}$  are Gibbs moves. The conditional distributions are (up to proportionality)

$$p(\boldsymbol{\mu}_k | \boldsymbol{\theta}_{1:I_k}, \boldsymbol{\Sigma}_k, \mathcal{Y}_{1:I_k}) \propto \prod_{i=1}^{I_k} \mathcal{N}(\boldsymbol{\theta}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (\text{A14})$$

where  $\prod_{i=1}^{I_k}$  denotes the product over all voxels in ROI  $k$ . By rearranging the multivariate normal PDF so that  $\boldsymbol{\mu}_k$  is the variable, it is given that

$$\prod_{i=1}^{I_k} \mathcal{N}(\boldsymbol{\theta}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \mathcal{N}(\boldsymbol{\mu}_k; \mathbf{m}, V), \quad (\text{A15})$$

where

$$\mathbf{m} = I_k^{-1} \sum_{i=1}^{I_k} \boldsymbol{\theta}_i, \text{ and } V = I_k^{-1} \boldsymbol{\Sigma}_k. \quad (\text{A16})$$

The MCMC updates are therefore sampled as

$$\boldsymbol{\mu}_k \sim \mathcal{N}(\mathbf{m}, V), \quad (\text{A17})$$

where  $\mathcal{N}(\mathbf{m}, V)$  is a multivariate normal distribution with mean  $\mathbf{m}$  and covariance  $V$ . Following the same steps for  $\boldsymbol{\Sigma}_k$  (full details in Orton et al. [16]) gives the MCMC update:

$$\boldsymbol{\Sigma}_k \sim W^{-1}(\Phi, I_k - N_\theta - 1), \quad (\text{A18})$$

where  $W^{-1}$  is the inverse-Wishart distribution,  $N_\theta$  is the number of free model parameters to be fitted and

$$\Phi = \sum_{i=1}^{I_k} (\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)(\boldsymbol{\theta}_i - \boldsymbol{\mu}_k)^T. \quad (\text{A19})$$

### A.5.2 | Voxelwise Parameters

For the non-signal fraction voxelwise parameters the posterior distribution for voxel  $i$  in ROI  $k$ , up to proportionality, is

$$p(x_{i,j} | x_{i,-j}, \{f_{i,c}\}_{c=1}^{C-1}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \propto [\mathbf{y}_i^T \mathbf{y}_i - (\mathbf{y}_i^T \mathbf{g}_i)^2 / \mathbf{g}_i^T \mathbf{g}_i]^{-N/2} \mathcal{N}(\boldsymbol{\theta}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k), \quad (\text{A20})$$

where  $x_{i,j}$  is the  $j$ th parameter value,  $x_{i,-j} = \{x_{i,1}, \dots, x_{i,j-1}, x_{i,j+1}, \dots, x_{i,J}\}$  denotes the set of all non-signal fraction parameters except  $x_{i,j}$ , and  $\{f_{i,c}\}_{c=1}^{C-1}$  are the signal fractions.

As in Orton et al. [16], this is sampled with a Metropolis–Hastings algorithm. Proposed parameters,  $x_{i,j}^*$ , are first sampled from Gaussian distributions as

$$x_{i,j}^* \sim \mathcal{N}(x_{i,j}, w_{x_{i,j}}), \quad (\text{A21})$$

where  $x_{i,j}$  is the current parameter value and  $w_{x_{i,j}}$  the proposal distribution variance, which should reflect the parameter scale and can be tuned for optimal algorithm performance.

The acceptance probability utilises the ratio of the posterior distributions for  $x_{i,j}$  and  $x_{i,j}^*$ :

$$\alpha(x_{i,j} \rightarrow x_{i,j}^*) = \min \left\{ 1, \frac{p(x_{i,j}^* | x_{i,-j}, \{f_{i,c}\}_{c=1}^{C-1}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{p(x_{i,j} | x_{i,-j}, \{f_{i,c}\}_{c=1}^{C-1}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)} \right\}, \quad (\text{A22})$$

where the numerator and denominator are calculated using Equation (A20).

MCMC moves for the signal fractions are the same, except that the posterior distributions now contain the terms enforcing  $\sum_{c=1}^C f_c = 1$ :

$$p(f_{i,c} | x_{i,1}, \dots, x_{i,J}, f_{i,-c}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \propto \begin{cases} [\mathbf{y}_i^T \mathbf{y}_i - (\mathbf{y}_i^T \mathbf{g}_i)^2 / \mathbf{g}_i^T \mathbf{g}_i]^{-N/2} & \text{if } \sum_{c=1}^{C-1} f_{i,c} \leq 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{A23})$$

where  $f_{i,-c} = \{f_{i,1}, \dots, f_{i,c-1}, f_{i,c+1}, \dots, f_{i,C}\}$  are the other signal fractions apart from  $f_{i,c}$ . Again, proposed values are sampled as

$$f_{i,c}^* \sim \mathcal{N}(f_{i,c}, w_{f_{i,c}}), \quad (\text{A24})$$

where  $f_{i,c}$  is the current signal fraction. The acceptance probabilities are

$$\alpha(f_{i,c} \rightarrow f_{i,c}^*) = \min \left\{ 1, \frac{p(f_{i,c}^* | x_{i,1}, \dots, x_{i,J}, f_{i,-c}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{p(f_{i,c} | x_{i,1}, \dots, x_{i,J}, f_{i,-c}, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)} \right\} \quad (\text{A25})$$

### A.5.3 | Metropolis–Hastings Acceptance Ratio

Metropolis–Hastings jumping variances are tuned during the first half of the burn-in period (see Appendix A.5.4) to achieve an acceptance ratio that efficiently samples the posterior distribution. Variances are adjusted voxelwise for all compartment signal fractions,  $\{w_{f_{i,c}}\}_{c=1}^{C-1}$ , and other parameters,  $\{w_{x_{i,j}}\}_{j=1}^J$ . The tuning strategy of Orton et al. [16] is extended here by applying the following update rule at every  $n_s$  MCMC steps:

$$w_{x_{i,j}} = \frac{w_{x_{i,j}}(n_s + 1)}{\left( \frac{1}{1 - w_i} \left( (n_s + 1) - R_{x_{i,j}} \right) \right)}, \quad (\text{A26})$$

where  $R_{x_{i,j}}$  is the number of accepted updates for parameter  $x_{i,j}$  during the previous  $n_s$  steps and  $w_i$  is the target acceptance rate. Compartment signal fractions are updated similarly:

$$w_{f_{i,c}} = \frac{w_{f_{i,c}}(n_s + 1)}{\left( \frac{1}{1 - w_i} \left( (n_s + 1) - R_{f_{i,c}} \right) \right)}. \quad (\text{A27})$$

The target acceptance rate is set to  $w_i = 0.25$  [54], adjusting the jumping variances to aim for an acceptance rate of 25%.

### A.5.4 | Burn-In

The burn-in period defines the first  $N_{burn}$  samples that are discarded from the calculation of the posterior distribution. Removing these early samples reduces biases introduced by low-probability regions and subsequent large parameter fluctuations owing to poor initialisation, as the posterior is estimated from the latter part of the chain that has converged to a more stationary distribution.

---

**ALGORITHM A1.** Pseudocode for hierarchical Bayesian model fitting of a general microstructural model and the MCMC algorithm used for inference.

---

```

for ROIs  $k = 1$  to  $k = K$  do
  for voxels  $i = 1$  to  $i = I_k$  do
     $\{f_c^{(0)}\}_{c=1}^{C-1}, \{x_j^{(0)}\}_{j=1}^J \leftarrow$  LSQ estimates
  end for
   $\{f_c^{(0)}\}_{c=1}^{C-1}, \{x_j^{(0)}\}_{j=1}^J \leftarrow$  parameter transform (Equation A12)
   $N_s \leftarrow$  number of MCMC steps
  for steps  $s = 1$  to  $s = N_s$  do
    for ROIs  $k = 1$  to  $k = K$  do
       $\mu_k^{(s)} \leftarrow$  sample using Equation (A17)
       $\Sigma_k^{(s)} \leftarrow$  sample using Equation (A18)
      for voxels  $i = 1$  to  $i = I_k$  do
        for parameters  $j = 1$  to  $j = J$  do
           $x_{ij}^* \leftarrow$  sample using Equation (A21)
           $\alpha(x_{i,j} \rightarrow x_{ij}^*) \leftarrow$  Equation (A22)
           $u \leftarrow$  sample from Uniform (0, 1)
          if  $u < \alpha$  then
             $x_{ij}^{(s)} \leftarrow x_{ij}^*$ 
          else
             $x_{ij}^{(s)} \leftarrow x_{ij}^{(s-1)}$ 
          end if
        end for
        for compartments  $c = 1$  to  $c = C - 1$  do
           $f_{i,c}^* \leftarrow$  sample using Equation (A24)
           $\alpha(f_{i,c} \rightarrow f_{i,c}^*) \leftarrow$  Equation (A25)
           $u \leftarrow$  sample from Uniform (0, 1)
          if  $u < \alpha$  then
             $f_{i,c}^{(s)} \leftarrow f_{i,c}^*$ 
          else
             $f_{i,c}^{(s)} \leftarrow f_{i,c}^{(s-1)}$ 
          end if
        end for
         $f_{i,c}^{(s)} \leftarrow 1 - \sum_{c=1}^{C-1} f_{i,c}^{(s)}$ 
        if  $s \bmod n_s = 0$  and  $s \leq N_{burn}/2$  then
          for voxels  $i = 1$  to  $i = I_k$  do
            for parameters  $j = 1$  to  $j = J$  do
               $w_{x_{ij}} \leftarrow$  Equation (A26)
            end for
            for compartments  $c = 1$  to  $C - 1$  do
               $w_{f_{i,c}} \leftarrow$  Equation (A27)
            end for
          end for
        end if
      end for
    end for
  end for
   $\{f_c^{(s)}\}_{c=1}^{C-1}, \{x_j^{(s)}\}_{j=1}^J \leftarrow$  inverse parameter transform (Equation A13)

```

---