Supplementary Table 4. Hybrid-multidimensional MRI (HM-MRI) model fitting based on artificial neural networks.

Signal model	$$s(b,\mathrm{TE}) = s_0 \left(v_l \, e^{-b\, d_l - \frac{\mathrm{TE}}{T_{2l}}} + (1 - v_l) \left(v_e \, e^{-b\, d_e - \frac{\mathrm{TE}}{T_{2e}}} + (1 - v_e) \, e^{-b\, d_s - \frac{\mathrm{TE}}{T_{2s}}} \right) \right)$$ *b*: b-value; TE: echo time
Tissue parameters	s_0 (apparent proton density): fitting range [0.5; 10.0] v_l (luminal water voxel volume fraction): fitting range [0.0; 1.0] v_e (epithelial fraction of non-luminal tissue): fitting range [0.0; 1.0] d_l (luminal water diffusivity): set to 2.5 $\mu m^2 \, ms^{-1}$ d_e (epithelial water diffusivity): set to 0.4 $\mu m^2 \, ms^{-1}$ d_s (stromal water diffusivity): set to 1.4 $\mu m^2 \, ms^{-1}$ T_{2l} (luminal water T2): set to 650 ms T_{2e} (epithelial water T2): set to 45 ms T_{2s} (stromal water T2): set to 95 ms

DNN implementation	• The DNN is built of a cascade of hidden layers (here 7), each consisting of a linear matrix operation followed by ReLU activation, with $\mathrm{ReLU}(x) = \max(0, x)$
	• The DNN features N input neurons (as many as the number of input measurements per voxel) and 9 output neurons, which map the 9 tissue parameters:
	  o for fully-sampled signals we use $\{16, 15, 14, 13, 12, 11, 10, 9\}$ neurons;
	  o for 12-measurement sub-protocols we use $\{12, 12, 11, 11, 10, 10, 9, 9\}$ neurons;
	  o for 9-measurement sub-protocols we use $\{9, 9, 9, 9, 9, 9, 9, 9\}$ neurons
	• The i-th output neuron activation u_i, defined in the range $0 \leq u_i \leq u_{max}$, is mapped to the i-th tissue parameter p_i as
	$$p_i = \frac{2}{1 + e^{-\alpha_i \left(\log(\mathrm{softplus}(u_i)) - \log(\log(2))\right)}} - 1,$$
	where α_i is a learnable scaling factor and $\mathrm{softplus}(x) = \log(1 + e^x)$
	• Output MRI signals are calculated from tissue parameters p_i with the equation above

DNN training	• Input measurements $\{s(b_n, \mathrm{TE}_n) \mid n = 1, \ldots, N\}$ from each voxel are normalised by computing $a(b_n, \mathrm{TE}_n) = s(b_n, \mathrm{TE}_n) \, / \max_n \left(s(b_n, \mathrm{TE}_n)\right)$ • The DNN is optimised by backpropagating the ℓ^2-norm of the error (i.e. mean squared error, MSE) between ground truth MRI measurements and signal prediction • Optimisation is performed with ADAM on synthetic MRI signals for 50 epochs (learning rate of 10^{-4}; one update per mini-batch of 100 voxels) for sub-protocols, repeating the training 10 times with different random DNN initialisations • Synthetic MRI signals are computed for uniformly distributed tissue parameters within the ranges reported above, adding Rician noise with $SNR = \frac{s_0}{\sigma_{\mathrm{noise}}}$ within the range $[35; 250]$. We use 80,000 voxels as training set and 20,000 as validation set • The DNN providing the minimum validation loss is deployed