"*Select and retrieve via direct upsampling*" network (SARDU-Net): a data-driven, model-free, deep learning approach for quantitative MRI protocol design

Francesco Grussu[1,2]*, Stefano B. Blumberg[2], Marco Battiston[1], Lebina S. Kakkar[3], Hongxiang Lin[2], Andrada Ianuş[4], Torben Schneider[5], Saurabh Singh[3], Roger Bourne[6], Shonit Punwani[3], David Atkinson[3], Claudia A. M. Gandini Wheeler-Kingshott[1,7,8], Eleftheria Panagiotaki[2], Thomy Mertzanidou[2] and Daniel C. Alexander[2]

[1]Queen Square MS Centre, Queen Square Institute of Neurology, Faculty of Brain Sciences, University College London, London (UK)
[2]Centre for Medical Image Computing, Department of Computer Science, University College London, London (UK)
[3]Centre for Medical Imaging, University College London, London (UK)
[4]Champalimaud Research, Champalimaud Centre for the Unknown, Lisbon (Portugal)
[5]Philips United Kingdom, Surrey (UK)
[6]Discipline of Medical Imaging Science, Faculty of Medicine and Health, The University of Sydney, Sydney (Australia)
[7]Department of Brain and Behavioural Sciences, University of Pavia, Pavia (Italy)
[8]Brain MRI 3T Center, IRCCS Mondino Foundation, Pavia (Italy)

*Corresponding author.
Correspondence to: Russell Square House, 10-12 Russell Square, London WC1B 5EH (UK). Tel.: +442031087499, email: f.grussu@ucl.ac.uk

An expanded version of this preprint will be submitted for consideration to Magnetic Resonance in Medicine

Abstract

Purpose

We introduce "*Select and retrieve via direct upsampling*" network (SARDU-Net), a data-driven deep learning framework for model-free quantitative MRI (qMRI) experiment design. Here we provide a practical demonstration of its utility on *in vivo* joint diffusion-relaxation imaging (DRI) of the prostate.

Methods

SARDU-Net selects subsets of informative qMRI measurements within lengthy pilot scans. The algorithm consists of two deep neural networks (DNNs) that are trained jointly end-to-end: a *selector*, identifying a subset of input qMRI measurements, and a *predictor*, estimating fully-sampled signals from such a subset. We studied 3T prostate DRI scans performed on 3 healthy volunteers with 16 unique (b,TE) values (diffusion-/T2-weighting), and used SARDU-Net to identify sub-protocols of 12 and 9 measurements. The reproducibility of the sub-protocol selection was evaluated, and sub-protocols were assessed for their potential of informing multi-contrast analysis, as for example Hybrid Multi-dimensional MRI (HM-MRI).

Results

SARDU-Net identifies informative sub-protocols of specified size from a small number of pilot scans. The procedure is reproducible across training folds and random initialisations. Moreover, SARDU-Net sub-protocols corresponding to up to ~50% scan time reduction support downstream HM-MRI modelling for which they were not optimised explicitly, providing quality of fit in the top 5% of all tested sub-protocols.

Conclusions

SARDU-Net gives new opportunity to identify economical but informative qMRI data sets for clinical applications under high time pressure in a fully data-driven way from a few pilot scans. The simple SARDU-Net architecture makes the algorithm easy to train and appealing when extensive sub-protocol searches are intractable.

Keywords

Quantitative MRI; Protocol design; Deep learning; Diffusion-relaxation

1. Introduction

Quantitative MRI (qMRI) techniques enable the estimation of biophysical properties of imaged tissues from multi-contrast image sets (1), which hold promise as novel, system-independent biomarkers in several clinical contexts. Notable examples include: relaxation time constants, useful to assess myelination in the brain (2) or luminal structures in the prostate (3); diffusion characteristics, linked to cytoarchitecture in various anatomical districts (4–9); blood flow (10); mechanical stiffness of organs such as liver (11); tissue temperature (12).

qMRI methods can potentially overcome the main limitation of routine clinical MRI, i.e. its limited sensitivity and specificity to early and diffuse alterations that precede the appearance of focal lesions (13,14). Importantly, recent advances in acquisition have increased dramatically the number of images that can be acquired per unit time (15–17), enabling rich multi-modal qMRI sampling schemes, as for example joint diffusion-relaxation imaging (DRI) (18–20). Such novel approaches exploit the complementary information from multiple MRI contrasts, and may enable a better estimation of microstructural properties as compared to single-contrast methods (21,22). Nonetheless, the increased acquisition complexity makes the design of clinically viable protocols challenging. Clinical acquisitions should capture salient signal features within vast sampling spaces in acceptable times, thereby trading off between scan duration and information content.

Previous literature has dealt extensively with qMRI protocol optimisation, i.e. with the design of maximally informative sampling schemes given a specified budget in terms of number of images/scan time. Important examples include the design of optimal diffusion-weighting gradient strength, direction or waveforms in diffusion imaging (4,23–27); number and spacing of temporal sampling in relaxometry (28,29); DRI sampling (18). Such previous studies adopt different optimisation strategies, as for example Cramér–Rao lower bound (CRLB) minimisation based on Fisher information (30), Monte Carlo (MC) samplings (31), mutual information computation (32), empirical approaches (33) or discrete searches (34). Importantly, all such previous optimisation approaches rely on fixed, *a priori* representation of the measured MRI signals, such as biophysical models (e.g. multi-compartment models (4,7,9,35)) or phenomenological signal representations (e.g. cumulant expansions (36) or continuous distributions of MRI signal sources (18)). Such representations are useful to capture salient patterns of measured signals in response to changes in the prescribed MRI pulse sequence. However, adopting *a priori* signal representations inevitably leads to strong assumptions. Models are often formulated with the healthy tissue in mind (4) under the hypothesis that certain microstructural degrees of freedom are more relevant than others to

explain observed MRI contrasts (e.g. number and intrinsic properties of water compartments (4,9)). Moreover, such previous optimisation procedures typically focus on signal variability induced by thermal noise (37) under some expected signal-to-noise ratio (SNR) level (30), and do not account for additional variation induced by instrument-dependent factors (38) and/or physiological noise (39). Ultimately this may limit the generalisability of optimised protocols in real clinical settings and in presence of complex pathophysiological processes.

In this work we introduce the *"Select and retrieve via direct upsampling"* network (SARDU-Net), a new framework for qMRI protocol design that extends previous statistical optimisation approaches in a fully data-driven way, without relying on any *a priori* explicit signal model or representation. SARDU-Net is based on two deep neural networks (DNNs), which are trained on actual *in vivo* qMRI measurements with the aim of identifying informative sub-samplings within lengthy acquisitions. Following previous preliminary investigation (40), here we present the implementation of SARDU-Net and demonstrate its capability of finding economical but informative sub-protocols within lengthy state-of-the-art qMRI approaches, such as joint DRI of the prostate (41–43), whose deployment in real-world clinical settings is subject to high time pressure.

2. Methods

In this section we introduce the SARDU-Net algorithm and describe computational experiments performed on *in vivo* prostate qMRI data to demonstrate the method. SARDU-Net has been implemented using PyTorch (`http://pytorch.org`). A command-line toolbox is made freely available online (`http://github.com/fragrussu/sardunet`). All post-processing/analysis scripts written for this study are also made freely available (`http://github.com/fragrussu/PaperScripts/tree/master/sardunet`).

2.1. SARDU-Net algorithm

SARDU-Net is based on two fully-connected DNNs and is inspired by the deep autoencoder architecture (44,45). Autoencoders are powerful unsupervised learning tools useful to find low-dimensional representations of high dimensional data, and have shown promise to guide MRI sampling within vast acquisition spaces (46). In more detail, SARDU-Net is designed to analyse *in vivo* voxel signals, each made of a unidimensional array of M qMRI measurements, with the aim of identifying $D < M$ informative measurements within the input set.

Figure 1 shows the architecture of SARDU-Net. The first DNN (i.e. the encoder) acts as a *selector*. It takes a voxel signal made of M measurements $\mathbf{s} = [s_1 \quad \cdots \quad s_M]^{\mathrm{T}}$ as input, and it outputs a set of M corresponding scores $\mathbf{w} = [w_1 \quad \cdots \quad w_M]^{\mathrm{T}}$ of which $M - D$ are zero. Afterwards, the Hadamard product

$$\mathbf{w} \odot \mathbf{s} = [w_1 \, s_1 \quad \cdots \quad w_M \, s_M]^{\mathrm{T}} \quad (1)$$

is evaluated to select a subset of D out of M measurements, and passed to the second DNN (i.e. the decoder), which acts as a *predictor*. The predictor outputs $\mathbf{u} = [u_1 \quad \cdots \quad u_M]^{\mathrm{T}}$, an estimate of the fully-sampled signal $\mathbf{s}$ derived directly from the subset $\mathbf{w} \odot \mathbf{s}$.

Both selector and predictor are constructed as multi-layer, fully-connected feedforward DNNs, where each layer is obtained as a linear matrix operation followed by element-wise Rectified Linear Units (ReLU). Additionally, the activations of the selector output neurons are normalised to add up to 1 via *softmin* normalisation, and then thresholded so that only the D top firing neurons are kept and the remaining $M - D$ are zeroed, providing output scores $\mathbf{w}$ that effectively select D out of M measurements. The selector and predictor DNNs are optimised jointly end-to-end with aim of finding a subset of measurements that carries the most information regarding the fully-sampled MRI signal. For this, a loss function

$$L = \| \mathbf{u}(\mathbf{w} \odot \mathbf{s}) - \mathbf{s} \|_2^2 \qquad (2)$$

measuring the ℓ^2-norm (i.e. mean squared error, MSE) of the reconstruction error is minimised via error back-propagation (47) with ADAM stochastic optimisation (48) and optional dropout regularisation (49). In practice, the differentiable product $\mathbf{w} \odot \mathbf{s}$ enables the propagation of error gradients from the predictor back to the selector, and hence their joint optimisation.

For practical implementation, input voxels are normalised between 0 and 1 and then split into training and validation sets. Training voxels are grouped and passed through the network in mini-batches. During one such forward network pass, the selector's output neuron activations are averaged over the mini-batch before computation of $\mathbf{w}$, and network parameters then updated. The whole set of training voxels is inputted to the system for a fixed number of epochs, during which the value of the loss function is evaluated on the validation set as well. In doing so, the selector DNN is temporarily deactivated and the scores $\mathbf{w}$ are fixed to the latest values provided by training. Ultimately, a trained SARDU-Net is obtained by deploying the parameters at the epoch with minimum validation loss.

2.2. Prostate MRI: acquisition

We analysed *in vivo* DRI scans acquired on 3 healthy volunteers as part of an ongoing study (50). Scans were performed with a 3T Philips Achieva scanner in ethically approved imaging sessions after obtaining informed written consent. DRI scans featured different diffusion-/T2-weightings, which were achieved by varying the *b*-value and the echo time TE. A multi-slice diffusion-weighted (DW) 2D echo planar imaging (EPI) sequence was used, with the vendor's 32-channel torso coil for reception. Salient sequence parameters were: 14 axial slices, 5 mm-thick; field-of-view of 220×220 mm^2; in-plane resolution of 1.75×1.75 mm^2; repetition time TR = 2800 ms; SENSE factor of 1.6; half-scan factor of 0.62; number of averages of 2; 1 coronal REST slab for spatial saturation; readout bandwidth of 2.39 KHz/pixel. Scans were performed with 16 different (*b,*TE) values among *b* = {0, 500, 1000, 1500} s/mm^2×TE = {55, 87, 121, 150} ms. For each (*b*,TE) pair, 3 images were acquired, using 3 orthogonal diffusion-weighting gradients when *b* was not zero, for a total of 48 EPI images and a total DRI scan time of 6 min: 15 sec. A detailed list of sequence parameters is available in the appendix.

2.3. Prostate MRI: post-processing

DRI scans were denoised slice-by-slice with the Marchenko-Pastur Principal Component Analysis technique (51) (sliding window size: 7×7 voxel×voxel) in Matlab (The MathWorks, Inc., Natick, Massachusetts, USA; routine: `http://github.com/NYU-DiffusionMRI/mppca_denoise/blob/master/MPdenoising.m`). Afterwards, noise floor was mitigated on the denoising output with a custom Matlab implementation of the method of moments (52).

Subsequently, motion and eddy current distortions were corrected slice-by-slice by co-registering each 2D image to a reference with affine registration. For this purpose, NiftyReg `reg_aladin` (`http://cmictig.cs.ucl.ac.uk/wiki/index.php/NiftyReg`) was used, and the reference was obtained by averaging all EPI volumes. Finally, the three images at any fixed (*b*,TE) were averaged and the entire set of 16 volumes was normalised by dividing all voxel signals by the median signal across all volumes and voxels within the prostate. For this, a prostate mask was manually segmented on the mean EPI image calculated after co-registration using `FSLView` (53). For quality control, routine Apparent Diffusion Coefficient (ADC) and quantitative T2 were also calculated via linear fitting with MyRelax (`http://github.com/fragrussu/MyRelax`) (54): ADC was computed on the 4 images obtained at shortest TE, while T2 and on the 4 *b* = 0 images obtained at varying TE.

2.4. Prostate MRI: experiments

We ran computational experiments to study the ability of SARDU-Net to select informative sub-protocols within the full DRI scan. Measurement subsets selected by SARDU-Net were tested for their potential of informing downstream tissue parameter estimation, as for example via HM-MRI (41,42). We followed a leave-one-out approach and used two out of three subjects to train a SARDU-Net in turn. The remaining subject was then used to test whether SARDU-Net selected an informative sub-protocol.

2.4.1. SARDU-Net training

For each leave-one-out fold, the DRI signals from prostate voxels of the two training subjects were extracted and assigned at random to actual training set (80% of voxels) and validation set (20%). Sub-protocols of size $D = \{12, 9\}$ out of $M = 16$ measurements were searched by training a SARDU-Net for 300 epochs. Four hidden layers were used for selector/predictor DNNs (selector: $\{16, 15, 14, 13, 16\}$ neurons for $D = 12$ and $\{16, 14, 13, 11, 16\}$ for $D = 9$; predictor architecture mirrors the selector), and 18 different sets of learning options (mini-batch size $= \{100, 800, 1500\}$ voxels $\times$ learning rate $= \{10^{-4}, 10^{-3}\}$ $\times$ dropout regularisation $= \{0.0, 0.2, 0.4\}$). Moreover, for each of the 18 configurations of the learning options, the training was performed 8 independent times using different random network initialisation seeds to assess reproducibility.

2.4.2. Multi-contrast evaluation

Measurement subsets selected by SARDU-Net were evaluated for their potential to inform multi-contrast signal analyses, as for example HM-MRI (42). HM-MRI is a simple multi-exponential approach that models the total prostate signal as the sum of luminal, epithelial and stromal components:

$$s(b, \mathrm{TE}) = s_0 \left(v_l \, e^{-b\, d_l - \frac{\mathrm{TE}}{T_{2l}}} + (1 - v_l) \left(v_e \, e^{-b\, d_e - \frac{\mathrm{TE}}{T_{2e}}} + (1 - v_e)\, e^{-b\, d_s - \frac{\mathrm{TE}}{T_{2s}}} \right) \right). \quad (3)$$

In Equation 3, $s(b, \mathrm{TE})$ is the prostate signal at a fixed b-value and echo time TE. The 9 tissue parameters of the model are: s_0 (apparent proton density); v_l (luminal volume fraction); v_e (epithelial tissue fraction); d_l, d_e, d_s and T_{2l}, T_{2e}, T_{2s} (ADC and T2 of luminal/epithelial/stromal water).

Firstly, reference HM-MRI metrics were obtained on the fully-sampled DRI scans for the 3 subjects with deep learning-based model fitting (55,56) for quality control. A DNN was implemented with the qMRI-Net toolbox (`http://github.com/fragrussu/qMRINet`, class `qmrinetsig`), and consisted of 7 hidden layers ($\{16, 15, 14, 13, 12, 11, 10, 9\}$ neurons), each characterised by a linear matrix operation followed by element-wise ReLU, with the 9 output neurons mapping tissue parameters. The DNN was optimised by converting predicted tissue parameters to MRI signals with Equation 3, and then backpropagating the ℓ^2-norm of the error (i.e. MSE) between ground truth MRI measurements and signal prediction. Optimisation was performed with ADAM for 1500 epochs (learning rate of 10^{-3}; one update per mini-batch of 100 voxels) on synthetic MRI signals (80,000 voxels for training; 20,000 for validation. The synthetic MRI signal was computed for uniformly distributed tissue parameters within previously reported ranges (42), adding Gaussian noise with $SNR = \frac{s_0}{\sigma_{\text{noise}}}$ within the range $[35; 250]$. The DNN at the epoch with minimum validation loss was used for final HM-MRI parameter computation.

Afterwards, we tested whether SARDU-Net sub-protocols are indeed valid choices to inform downstream multi-contrast modelling among all possible sub-protocols of fixed size that could be extracted from the fully-sampled signals. To this end, we performed HM-MRI fitting also for each $\binom{M}{D} = \frac{M!}{(M-D)!\ D!}$ subset of $D = \{12, 9\}$ out of $M = 16$ measurements (1,820 for $D = 12$; 11,440 for $D = 9$) following a dictionary-based procedure in MATLAB (The MathWorks, Inc., Natick, Massachusetts, USA). Sub-protocols corresponding to SARDU-Net measurement selections were identified within the exhaustive list of all sub-protocols, as well as naïve sub-protocols obtained by uniform downsampling of *b,* TE or both. For each sub-protocol, voxel-wise measurements were compared to a data base of ~500,000 reference signals synthesised by varying HM-MRI uniformly on a grid of previously reported (42) compartment-specific average values. Tissue parameters providing the best fit were used to synthesise back the fully-sampled signals, which were compared to the actual full set of $M = 16$ measurements to compute average mean squared error (MSE) within the central slice of the prostate.

3. Results

3.1. SARDU-Net training

Table 1 shows the results of selecting sub-protocols of $D = 9$ and $D = 12$ out of $M = 16$ (*b*,TE) measurements on the *in vivo* DRI prostate data. The table reports sub-protocols for

each of the three independent training folds (i.e. training performed on MRI signals from two out of three subjects in turn). For each fold, results obtained using the training options providing the lowest validation loss are reported. For all folds and all values of D, these options were: mini-batch size of 100 voxels; learning rate of 10^{-3}; dropout regularisation probability of both selector and predictor DNN of 0.0 (i.e. no dropout). Table 1 demonstrates that in all cases sub-protocols sampled the full range of diffusion and relaxation weightings. However, measurements with lowest SNR (i.e. maximum *b* and longest TE) were selected less often. Moreover, the table also demonstrates the measurement selection was consistent across fold, in spite of each fold sharing only ~50% of the voxel data with another fold (for instance, the same measurements were selected in fold 1 and fold 2 for sub-protocols of size $D = 12$).

Figure 2 shows an example of the evolution of sub-protocol selection during SARDU-Net training on the *in vivo* prostate data for $D = 9$. In each panel of Figure 2, (*b*,TE) measurements are reported along columns while training epochs along rows; white indicates that a particular measurement was selected as part of the sub-protocol, while black indicates that it was not. In all leave-one-out folds, training converges in roughly 100 epochs. Some measurements are selected in the initial epochs and then dropped, while other measurements are kept until convergence once selected. Overall the selection process is consistent across folds, especially when fold 1 and 2 are compared.

Figure 3 shows the reproducibility of sub-protocol selection in each leave-one-out fold when SARDU-Net parameters are initialised with 8 different random seeds. Each panel reports as a 2D histogram the normalised count of each (*b*,TE) measurement being selected over the 8 different random seeds (top row: $D = 12$; bottom row: $D = 9$; different folds along different columns). The illustrations demonstrate that SARDU-Net sub-protocol selection is consistent across leave-one-out folds and random initialisations. For instance, the same 8 measurements out of 12 were consistently selected in all 8 random training repetitions for $D = 12$ in fold 2 and 3, while in no cases SARDU-Net selected the measurement corresponding to sequence parameters (*b*,TE) = (1500 s mm^{-2},150 ms).

3.2. Multi-contrast evaluation

Figure 4 illustrates examples of mean EPI images, ADC and T2 maps as well as luminal, epithelial and stromal HM-MRI indices obtained via DNN-based fitting *in vivo*. Quantitative values of all maps vary within the expected range for the healthy prostate (41,42), and regional metric variation is observed. HM-MRI disentangles a number of factors contributing to the patterns of macroscopic diffusion/relaxation decay, as highlighted for

instance in the prostate peripheral zone, where both hyperintense luminal (i.e. free) and epithelial (i.e. restricted) water fractions are observed as compared to the surrounding tissue.

Figure 5 shows results from assessment of the HM-MRI quality of fit. Each panel reports the distribution of average MSE within the prostate obtained for all possible sub-protocols of a fixed size on the test subject of the leave-one-out folds. Values of MSE obtained for SARDU-Net and naïve sub-protocols are reported explicitly with a vertical bar. The panels demonstrate that SARDU-Net identifies sub-protocol that enable robust HM-MRI fitting, which in almost all cases are in the top 5% sub-protocols with lowest average MSE.

4. Discussion

4.1. Summary and key findings

This paper presents SARDU-Net, a data-driven, model-free, deep-learning method for qMRI experiment design. SARDU-Net is designed to identify informative sub-protocols within lengthy *in vivo* scans, in order to facilitate the deployment of the latest qMRI techniques in contexts where scan time is limited. In this proof-of-concept we studied the performance of SARDU-Net on *in vivo* DRI scans of the prostate. Our main finding is that SARDU-Net identifies informative sub-protocols in a reproducible manner, and that the selected sub-protocols support multi-contrast signal analyses for which the measurement selection was not optimised explicitly, as demonstrated by quality-of-fit (i.e. MSE) metrics.

4.2. Prostate MRI analysis

We studied prospective DRI EPI scans of the prostate acquired at 3T on three healthy volunteers. The scans were characterised by 16 unique diffusion-/relaxation-weightings (i.e. by 16 unique (b,TE) measurements), and were analysed in leave-one-out computational experiments. SARDU-Net was used to identify shorter sub-protocols of 12 and 9 measurements out of 16, and the reproducibility of the measurement selection procedure across leave-one-out folds and random initialisations was assessed. Moreover, the selected sub-protocols were tested for their potential to inform independent multi-contrast signal analyses.

Our results demonstrate that sub-protocols identified by SARDU-Net do not necessarily feature uniform downsampling of the (b,TE) space. While the selected measurements span the whole range of diffusion-weightings and relaxation-weightings, measurements with extremely low SNR levels (i.e. maximum b and longest TE) are

consistently avoided. This suggests that noise is an important factor to consider when designing qMRI samplings. While conventional model-based approaches typically rely on some *a priori* hypotheses on the level and statistics of thermal noise, our fully-data driven approach enables the design of qMRI samplings from real-world noisy measurements.

Importantly, our investigation characterised the reproducibility of SARDU-Net. Results show that the sub-protocol selection procedure presented here shows enough stability to enable practical protocol design choices from a limited number of pilot scans. Nonetheless, our results also highlight that some variation in the measurement selection process can be observed across training folds (i.e. with different input data), as well as across random algorithm initialisations for a given training fold (i.e. with fixed input data). The former is likely to originate at least in part from intrinsic between-subject variability, and could therefore be minimised by ensuring that the pilot training cohort is large enough to capture biological variability. The latter suggests instead the presence of different local minima in the network loss function, a known issue in complex optimisation problems. This variability appears to be small, since a number of key information-carrying measurements are selected consistently. Nonetheless, in future we aim to further improve the training procedure by reducing its sensitivity to the initial conditions.

Finally, in this work we have tested sub-protocols selected by SARDU-Net for their ability of informing downstream multi-contrast analyses, for which they were not optimised explicitly. To this end, we considered the HM-MRI technique, a simple multi-exponential model that aims to capture jointly the (b,TE) dependence of the MRI signal. The HM-MRI model was fitted for all possible sub-protocols of fixed size in each test subject, and the quality of fit with respect to fully-sampled signals was measured via MSE. SARDU-Net sub-protocols provided MSE in the top 5%, and also enabled the robust computation of HM-MRI metrics as compared to the extensive list of all possible sub-protocols. Overall, these results demonstrate that SARDU-Net is successful in capturing the salient characteristics of qMRI signal. For this particular piece of analysis, we chose to adopt the HM-MRI model as it is a simple and convenient choice for the demonstration. However, different approaches within a wider landscape of alternative models could have been equally adopted, each with its own advantages and disadvantages. In particular VERDICT (7) and Relaxed-VERDICT (57), which SARDU-Net could adapt to naturally, would enable modelling of diffusion time dependence, neglected in this first demonstration. Future work on this aspect is warranted in the future.

4.3. Methodological considerations

The envisioned deployment of SARDU-Net would require the acquisition of a small number of rich, pilot qMRI scans when a new clinical study is being set up, from which informative sub-protocols could be identified given a scan time budget. Pilot test scans are typically performed any way for quality control when developing new MRI standard operating procedures. Importantly, such pilot scans could be included in subsequent group-level analyses, since the final qMRI protocol would be contained within the pilot measurement set.

In this study we demonstrate our new framework on healthy volunteers. However, the inclusion of patients in the pilot training set would provide important information on signal features typical of pathological tissues, as for example focal lesions. Data augmentation techniques could be used to increase the number of examples of such under-represented signals, as well as from any other tissue whose accurate characterisation may be of interest (e.g. grey matter as compared to white matter in the brain).

Furthermore, our proposed approach necessarily relies on the hypothesis that the training scans suffice to capture the salient features of the MRI signal, i.e. that they densely sample the qMRI sequence parameter space. The latest acquisition technologies (15) make it achievable to sample richly two to four sequence parameters (e.g. echo time, inversion time and diffusion encoding) in under one hour of scan time (17). However, our data-driven approach would become impractical if larger acquisition spaces were of interest, implying pilot protocols that exceed substantially the hour of scan time. In practice, in this work the prostate DRI acquisition was well under the hour limit, as the nominal scan time was about six minutes. This was due to the fact that the scan was performed as part of an ongoing multi-parametric MRI study (50). In future, we plan to sample richer (b,TE) samplings and include the dependence on the diffusion time (7,57), which is not considered in this first demonstration.

Notably, SARDU-Net is designed to identify measurements that capture most of the information conveyed by the fully-sampled signal without relying on fixed signal models, i.e. in a purely data-driven way. We remark that our approach is not in contrast with previous model-based qMRI optimisation (30), which remains a valid option when a particular microstructural model is the main interest of a study. Here we propose an alternative approach, which could be adopted when multiple downstream analysis methods are of interest or when the qMRI model of interest is not necessarily known *a priori* at the time of the acquisition.

Regarding the algorithmic details SARDU-Net, our implementation is based on two fully-connected DNNs, which are trained by analysing signals on a voxel-by-voxel basis. This simple two-network structure suffices to demonstrate the potential and flexibility of data-driven qMRI protocol design. Moreover, the light architectural design implies that the algorithm can be trained easily with limited computational resources (e.g. with a CPU on a laptop), which is particularly advantageous when extensive sub-protocols searches become unfeasible (i.e. a situation quickly reached when M is of the order of 50). Nonetheless, more sophisticated architectures as well as training strategies will be explored in the future to further improve the performance and stability of the technique. These refinements may focus for instance on the measurement selection stage, which could relax the current hard thresholding procedure performed on the selector's output to further stabilise the training. Also, state-of-the-art neural architecture searches (58) could also be adopted to more effectively design the two networks and their learning settings, as compared to the simpler grid search approach used here as a proof-of-concept.

4.4. Conclusions

Our new SARDU-Net framework identifies economical but informative qMRI protocols for clinical application under high time pressure. Using a small number of rich pilot acquisitions with long acquisition times, our framework identifies informative sub-protocols given a specific time constraint that capture the salient characteristics of fully-sampled MRI signals.

Appendix

This appendix reports a detailed list of the DRI sequence parameters.

Nucleus =	"H1";
SmartSelect =	"yes";
Coil 1 (exclude) =	"None";
Uniformity =	"CLEAR";
FOV RL (mm) =	220;
AP (mm) =	220;
FH (mm) =	70;
ACQ voxel size RL (mm) =	1.72;
AP (mm) =	1.76;
Slice thickness (mm) =	5;
Recon voxel size RL (mm) =	1.375;
AP (mm) =	1.375;
ACQ matrix M x P =	"128 x 125";
ACQ voxel MPS (mm) =	"1.72 / 1.76 / 5.00";
REC voxel MPS (mm) =	"1.375 / 1.375 / 5.00";
Fold-over suppression =	"rest";
slab thickness =	"half FOV";
Reconstruction matrix =	160;
SENSE =	"yes";
P reduction (AP) =	1.60000002;

```
MB SENSE =                          "no";
CS-SENSE =                          "no";
k-t BLAST =                         "no";
Stacks =                    1;
   type =                   "parallel";
   slices =                         14;
   slice gap =                      "user defined";
       gap (mm) =           0;
   slice orientation =              "transverse";
   fold-over direction =    "AP";
   fat shift direction =    "P";
Stack Offc. AP (P=+mm) = -17.4494019;
         RL (L=+mm) =      -8.37893963;
         FH (H=+mm) =      15.134407;
     Ang.  AP (deg) =                0;
         RL (deg) =                  -0;
         FH (deg) =                  2.9550209;
     Free rotatable =                "no";
Minimum number of packages =         1;
Slice scan order =                  "interleaved";
Large table movement =              "no";
PlanAlign =                         "no";
REST slabs =                        1;
   type =                   "free";
   orientation =            "coronal";
   thickness (mm) =                 40;
Rest  Offc. AP (P=+mm) = -109.625175;
         RL (L=+mm) =      -1.66951835;
         FH (H=+mm) =      43.0401611;
     Ang.  AP (deg) =                0.837242603;
         RL (deg) =                  0.313441545;
         FH (deg) =                  2.44953847;
   power =                           "1";
Shim  Size  AP (mm) =               181.93161;
         RL (mm) =                  213.879593;
         FH (mm) =                  90;
     Offc. AP (P=+mm) =    -17.4500008;
         RL (L=+mm) =      -8.38000011;
         FH (H=+mm) =      15.134407;
     Ang.  AP (deg) =                0;
         RL (deg) =                  -0;
         FH (deg) =                  2.9550209;
Catheter tracking =                 "no";
Interactive positioning =   "no";
Allow table movement =              "no";
Patient position =          "feet first";
Patient body position =             "feet first";
Patient orientation =               "supine";
Patient body orientation =  "supine";
Scan type =                         "Imaging";
Scan mode =                         "MS";
   technique =                      "SE";
Modified SE =                       "no";
Acquisition mode =                  "cartesian";
Fast Imaging mode =                 "EPI";
   shot mode =                      "single-shot";
Echoes =                            1;
   partial echo =                   "no";
Flip angle (deg) =                   90;
Halfscan =                          "yes";
   factor =                         0.623931646;
Water-fat shift =                   "minimum";
RF Shims =                          "adaptive";
Shim =                              "volume";
ShimAlign =                         "yes";
mDIXON =                            "no";
Fat suppression =                    "SPIR";
```

strength = "strong";
frequency offset = "default";
Grad Rev Fat suppr = "yes";
Water suppression = "no";
BB pulse = "no";
MTC = "no";
APT = "no";
Research prepulse = "no";
MDME = "no";
Act. TR (ms) = "2800";
Act. TE (ms) = "55, 55, 55, 87, 87, 87, 121, 121, 121, 150, 150, 150"
Diffusion mode = "DWI";
sequence = "SE";
gradient duration = "maximum";
gradient overplus = "yes";
b-factors =
"500, 1000, 1500, 500, 1000, 1500, 500, 1000, 1500, 500, 1000, 1500"
Diffusion gradient timing DELTA (ms) =
"26.3, 26.3, 26.3, 42.3, 42.3, 42.3, 59.3, 59.3, 59.3, 73.8, 73.8, 73.8;
Diffusion gradient timing delta (ms) =
"11.3, 11.3, 11.3, 27.3, 27.3, 27.3, 44.3, 44.3, 44.3, 58.8, 58.8, 58.8"
average high b = "no";
Multi-transmit = "yes";
Transmit channels = "both";
SAR mode = "high";
B1 mode = "default";
SAR allow first level = "yes";
PNS mode = "moderate";
Gradient mode = "enhanced";
SofTone mode = "no";
Cardiac synchronization = "no";
Heart rate > 250 bpm = "no";
Respiratory compensation = "no";
Navigator respiratory comp = "no";
Flow compensation = "no";
Temporal slice spacing = "default";
NSA = 2;
SMART = "no";
Manual start = "no";
Dynamic study = "no";
Arterial Spin labeling = "no";
Preparation phases = "auto";
Interactive F0 = "no";
SmartPlan survey = "no";
MIP/MPR = "no";
Images = "M", (3) "no";
Reference tissue = "Liver";
Recon compression = "No";
Preset window contrast = "soft";
Uniformity correction = "no";
Geometry correction = "default";
IF_info_seperator = 1634755923;
Min. slice gap (mm) = -0;
EPI factor = 95;
WFS (pix) / BW (Hz) = "31.624 / 13.7";
BW in EPI freq. dir. (Hz) = "2394.4";
PNS / level = "80 % / normal";
dB/dt = "53.8 T/s";

Data and code availability statement

A command-line implementation of SARDU-Net is made freely available online (`http://github.com/fragrussu/sardunet`). Post-processing and analysis scripts written to perform the experiments for this paper are also made openly available (`http://github.com/fragrussu/PaperScripts/tree/master/sardunet`).
The following open-access toolboxes have been used:

- qMRI-Net: `http://github.com/fragrussu/qMRINet`;
- MyRelax: `http://github.com/fragrussu/MyRelax`;
- MRITools: `http://github.com/fragrussu/MRItools`;
- FSL: `http://fsl.fmrib.ox.ac.uk/fsl/fslwiki`;
- NiftyReg: `http://cmictig.cs.ucl.ac.uk/wiki/index.php/NiftyReg`;
- MP-PCA: `http://github.com/NYU-DiffusionMRI/mppca_denoise/blob/master/MPdenoising.m`.

The *in vivo* MRI scans cannot be made openly available online due to privacy issues of clinical data according to GDPR regulations. Researchers interested in accessing the scans can contact Prof Daniel C. Alexander (`d.alexander@ucl.ac.uk`). A data sharing agreement enabling non-commercial research use will be stipulated.

Acknowledgements

The authors would like to thank the volunteers for their time; the radiographers at UCH and F. Gong for their assistance with prostate MRI protocol set up and data acquisition; Philips Healthcare for assistance in protocol development and for access to research protocols. FG is grateful to R. Tanno, E. Kaden, M. Palombo, P. Slator and C. Jin for useful discussion.

Funding

This project was funded by the Engineering and Physical Sciences Research Council (EPSRC EP/R006032/1, M020533/1, G007748, I027084, N018702). This project has received funding under the European Union's Horizon 2020 research and innovation programme under grant agreement No. 634541 and 666992, and from: Rosetrees Trust (UK, funding FG); Prostate Cancer UK Targeted Call 2014 (Translational Research St.2, project reference PG14-018-TR2); Spinal Research (UK), Wings for Life (Austria), Craig H. Neilsen Foundation (USA) for jointly funding the INSPIRED study; Wings for Life (#169111); UK Multiple Sclerosis Society (grants 892/08 and 77/2017); the Department of Health's National Institute for Health

Research (NIHR) Biomedical Research Centres and UCLH NIHR Biomedical Research Centre; Champalimaud Centre for the Unknown, Lisbon (Portugal); European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101003390.

Conflicts of interest disclosures

TS is an employee of Philips Healthcare UK.

Figures

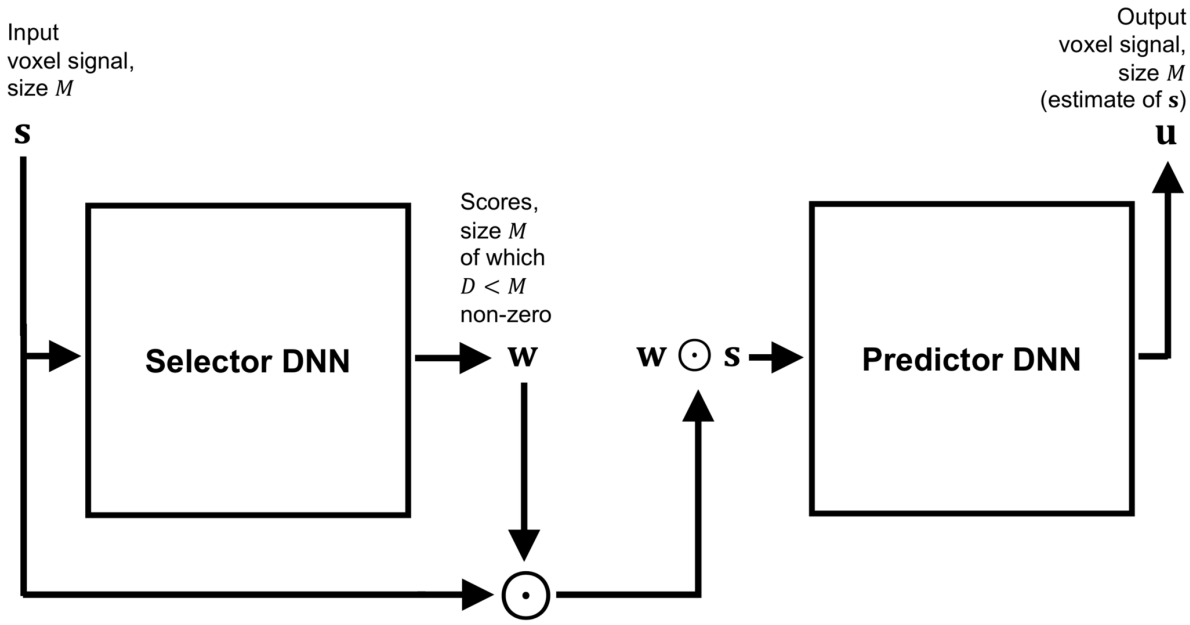

FIGURE 1 Block diagram summarising the architecture of SARDU-Net, made of two fully-connected DNN that are optimised jointly end-to-end to find optimal subsets of measurements within lengthy quantitative MRI acquisitions. The first DNN is a *selector*: it takes as input a signal made of multiple MRI measurements from a voxel, and outputs a corresponding set of scores, which nullify non-salient measurements effectively extracting a subset of the fully-sampled signal. The second DNN is a *predictor*: it aims to retrieve the fully-sampled signal from the subset received from the selector.

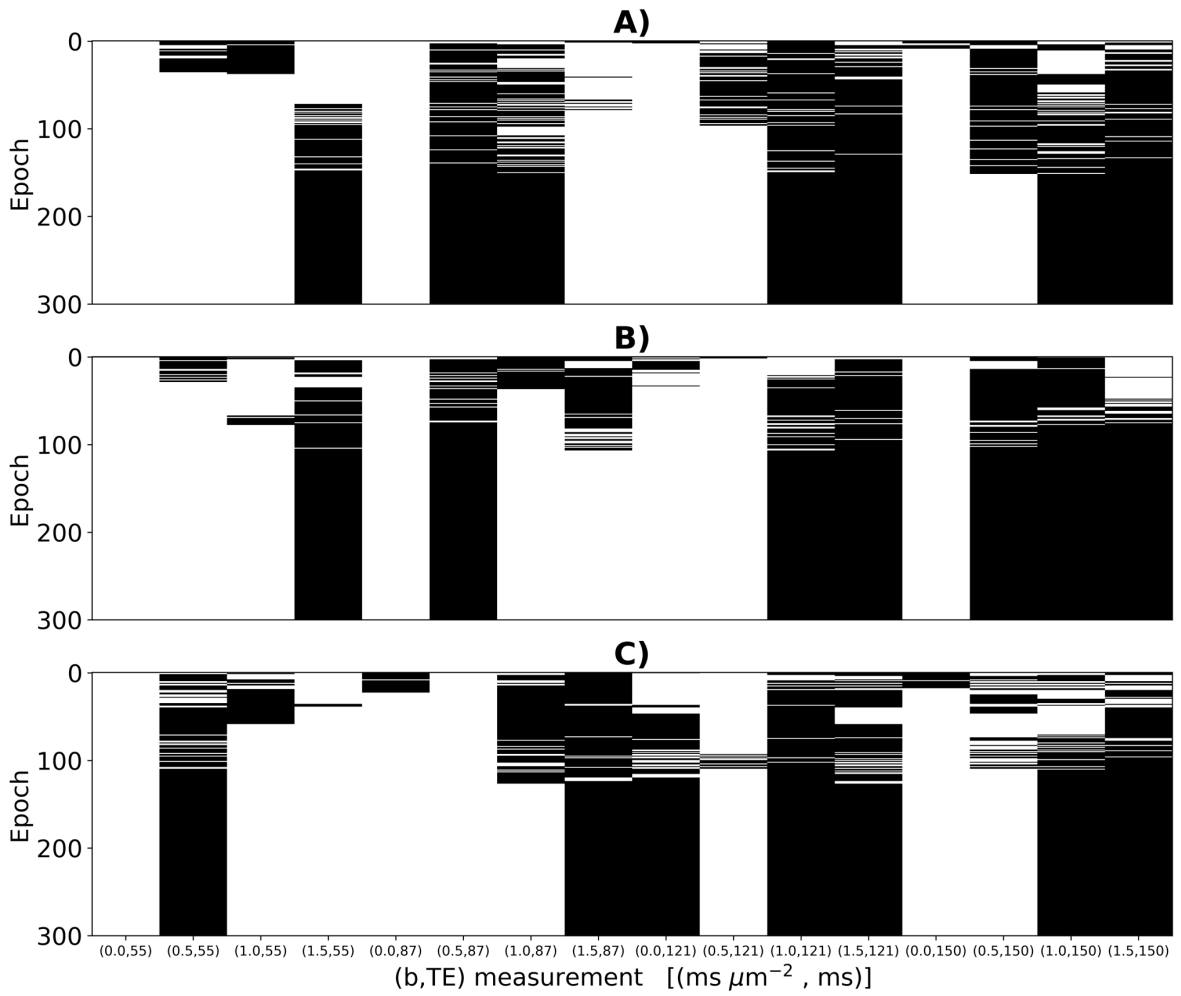

FIGURE 2 Illustration of measurement selection during SARDU-Net training for the *in vivo* DRI prostate data. In each panel, an image illustrates in white measurements that are selected by SARDU-Net at each training epoch, while black is used for measurements that are not selected. Different columns refer to different (*b*,TE) measurements, while training epochs are reported along rows. Panel A (top), B (middle) and C (bottom) respectively refer to the case when training is performed leaving subject 1, 2 and 3 (panels A, B and C) out as test set.

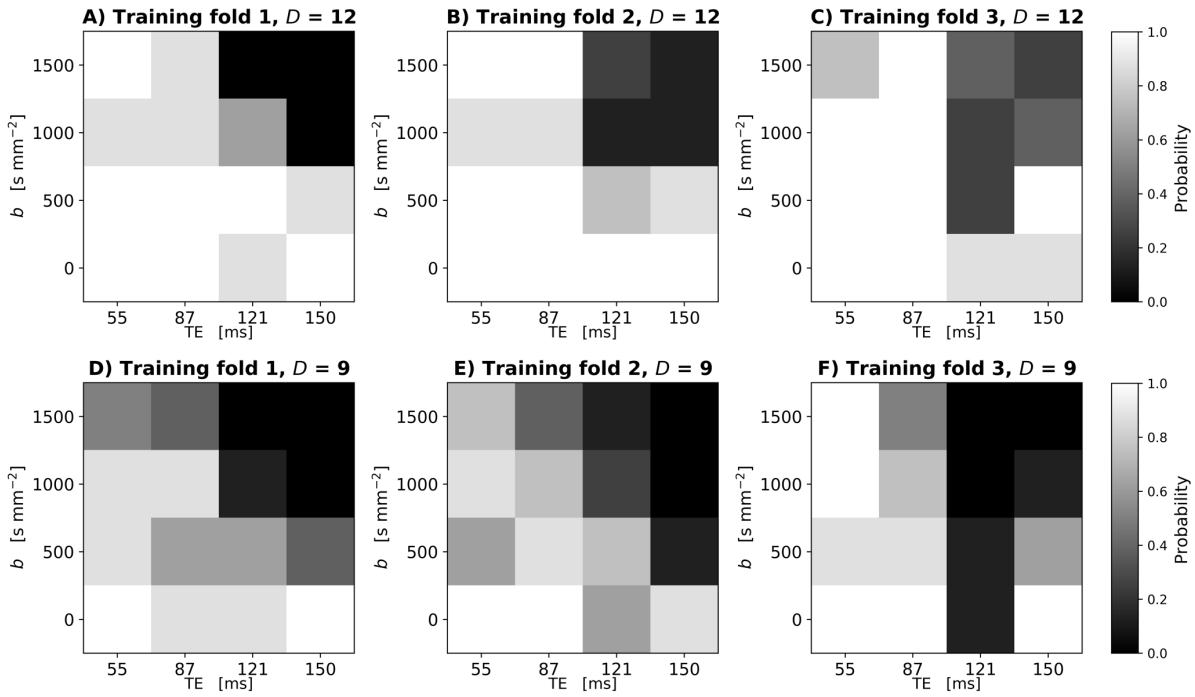

FIGURE 3 Reproducibility of SARDU-Net measurement selection for *in vivo* prostate DRI over different leave-one-out training folds and random initialisations. The image in each panel shows a normalised 2D histogram describing the probability of each (*b*,TE) measurement being selected as part of an optimal sub-protocol over 8 different repetition of the SARDU-Net training. Each repetition featured a unique random initialisation of the SARDU-Net parameters, with all other training options (i.e. mini-batch size, dropout regularisation, learning rate) kept fixed to the configuration providing the lowest validation loss. Panels A to C (top): selection of $D = 12$ out of $M = 16$ measurements (A, left: subject 1 left out for training and used for testing; B, middle: subject 2 left out; C, right: subject 3 left out). Panels D to F (bottom): selection of $D = 9$ out of $M = 16$ measurements (D, left: subject 1 left out of training and reserved for testing; E, middle: subject 2 left out; F, right: subject 3 left out).

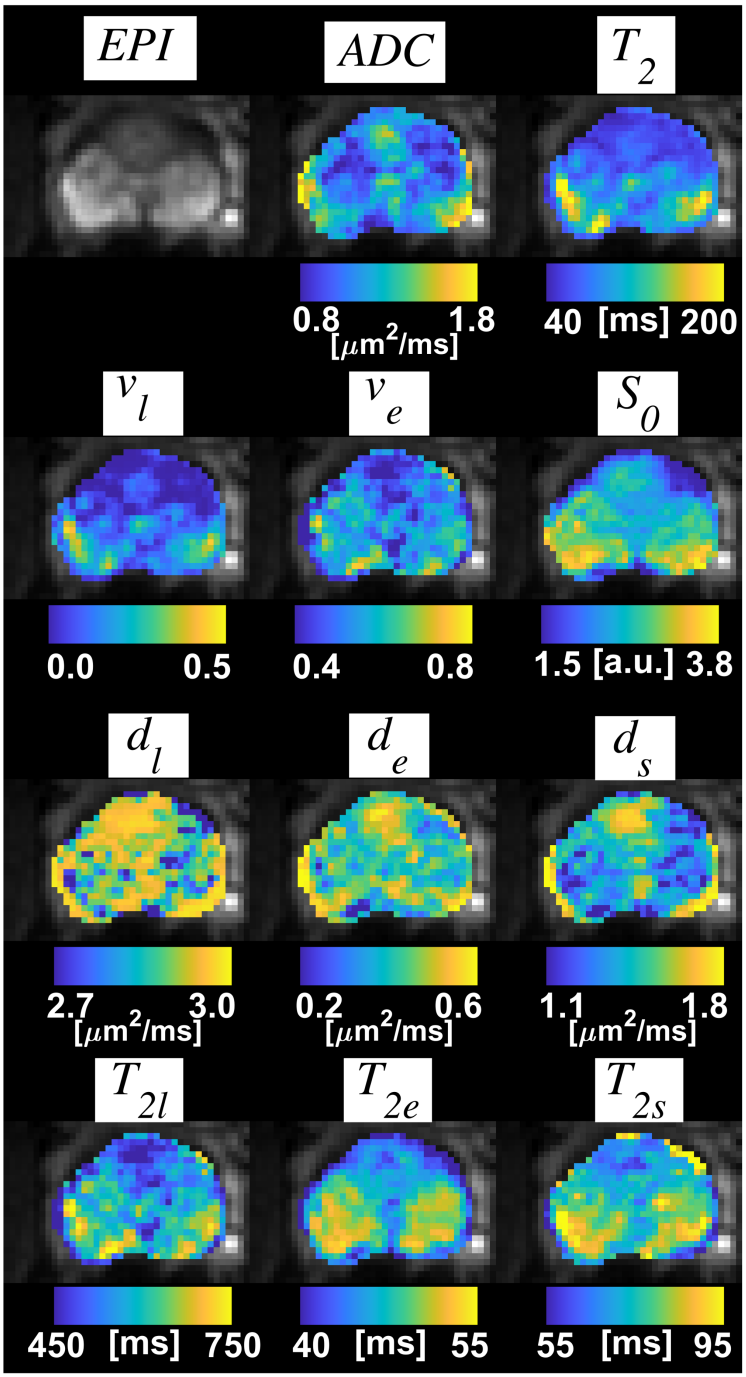

FIGURE 4 Examples of parametric tissue maps obtained in the prostate of one of the subjects. Top row: mean EPI image after post-processing (left); routine ADC (middle) and T2 (right) from mono-exponential decay fits. Second to fourth row from top: reference HM-MRI metrics obtained via DNN-based fitting. Second row: lumen voxel fraction v_l (left), epithelial tissue fraction v_e (middle), apparent proton density s_0 (right). Third row: diffusion coefficients of luminal (d_l, left), epithelial (d_e, middle) and stromal (d_s, right) water pools. Fourth row: T2 relaxation times of luminal (T_{2l}, left), epithelial (T_{2e}, middle) and stromal (T_{2s}, right) water pools.

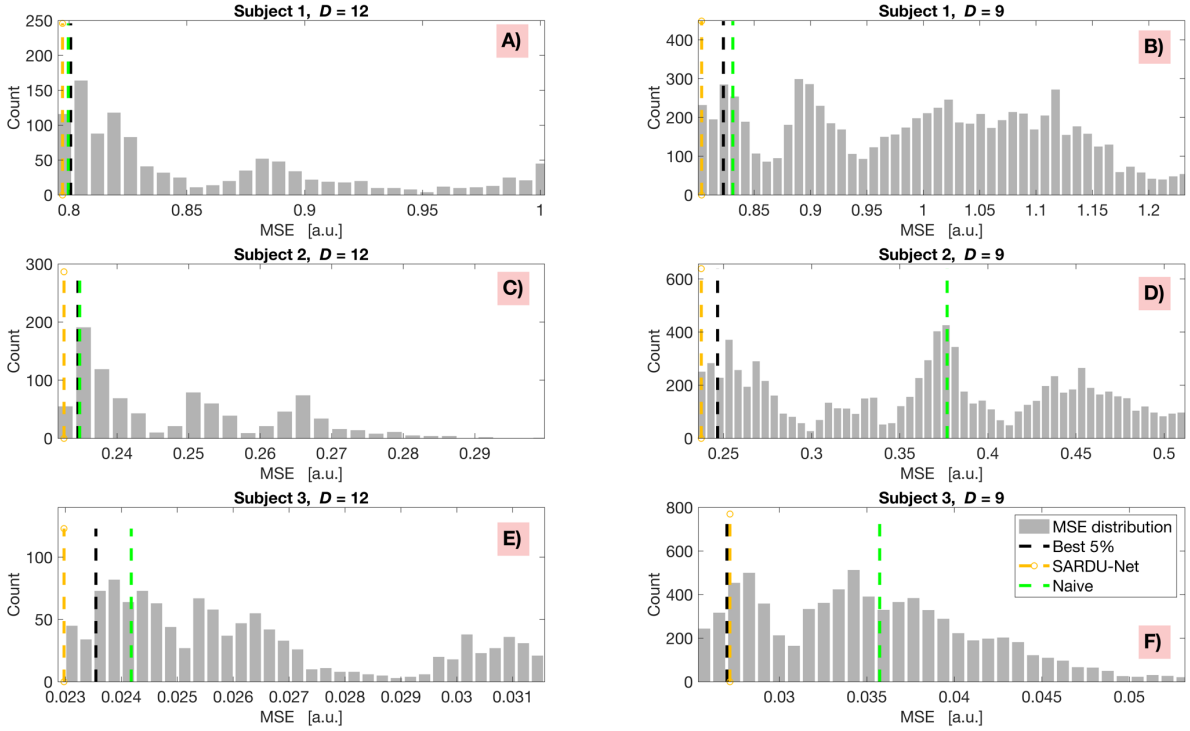

FIGURE 5 Bar plots visualising the distribution of MSE from HM-MRI dictionary fitting performed in the test subject of each leave-one-out fold, for all possible sub-protocols of fixed size (the right tail of the distribution has been cut for better visualisation). The MSE for sub-protocols selected by SARDU-Net and by naïve uniform downsampling are respectively indicated by a vertical yellow and green line. The top 5% MSE values are also flagged by a vertical black bar. Results obtained when subject 1, 2 and 3 are used as test set are reported respectively on top (panels A and B); in the middle (panels C and D); in the bottom (panels E and F). Panels A, C, E (left) refer to MSE computations performed for sub-protocols of size $D = 12$ out of $M = 16$ measurements, while panels B, D, F (right) for sub-protocols of size $D = 9$ out of $M = 16$.

Tables

TABLE 1 SARDU-Net measurement selection in each of the 3 leave-one-out folds obtained with the learning options providing the lowest validation loss (in all cases mini-batch size 100 of voxels, learning rate of 10^{-3}, dropout regularisation probability of both selector/predictor of 0.0). The table shows results for the selection of both $D = 12$ and $D = 9$ measurements out of $M = 16$ input measurements. In training fold 1, 2 and 3, subject 1, 2 and 3 was respectively left out as test set and not used to train a SARDU-Net via back-propagation.

	Fold 1	**Fold 2**	**Fold 3**
(*b*, TE) measurements selected for $D = 12$ [(s mm^{-2}, ms)]	(0,55), (500,55), (1000,55), (1500,55), (0,87), (500,87), (1000,87), (1500,87), (0,121), (500,121), (0,150), (500,150)	(0,55), (500,55), (1000,55), (1500,55), (0,87), (500,87), (1000,87), (1500,87), (0,121), (500,121), (0,150), (500,150) (same as fold 1)	(0,55), (500,55), (1000,55), (1500,55), (0,87), (500,87), (1000,87), (1500,87), (0,121), (0,150), (500,150), (1000,150)
(*b*, TE) measurements selected for $D = 9$ [(s mm^{-2}, ms)]	(0,55), (500,55), (1000,55), (0,87), (1500,87), (0,121), (500,121), (0,150), (500,150)	(0,55), (500,55), (1000,55), (0,87), (500,87), (1000,87), (500,121), (1000,121), (0,150)	(0,55), (1000,55), (1500,55), (0,87), (500,87), (1000,87), (500,121), (0,150), (500,150)

References

1. Cercignani M, Bouyagoub S. Brain microstructure by multi-modal MRI: Is the whole greater than the sum of its parts? Neuroimage 2018;182:117–127 doi: 10.1016/j.neuroimage.2017.10.052.
2. Laule C, Leung E, Li DKB, et al. Myelin water imaging in multiple sclerosis: quantitative correlations with histopathology. Mult. Scler. J. 2006;12:747–753 doi: 10.1177/1352458506070928.
3. Sabouri S, Chang SD, Savdie R, et al. Luminal Water Imaging: A New MR Imaging T2 Mapping Technique for Prostate Cancer Diagnosis. Radiology 2017;284:451–459 doi: 10.1148/radiol.2017161687.
4. Zhang H, Schneider T, Wheeler-Kingshott C, Alexander D. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. Neuroimage 2012.
5. Grussu F, Schneider T, Zhang H, Alexander DC, Wheeler–Kingshott CAM. Neurite orientation dispersion and density imaging of the healthy cervical spinal cord in vivo. Neuroimage 2015;111:590–601 doi: 10.1016/j.neuroimage.2015.01.045.
6. Duval T, McNab JA, Setsompop K, et al. In vivo mapping of human spinal cord microstructure at 300 mT/m. Neuroimage 2015;118:494–507 doi: 10.1016/j.neuroimage.2015.06.038.
7. Panagiotaki E, Chan RW, Dikaios N, et al. Microstructural Characterization of Normal and Malignant Human Prostate Tissue With Vascular, Extracellular, and Restricted Diffusion for Cytometry in Tumours Magnetic Resonance Imaging. Invest. Radiol. 2015;50:218–227 doi: 10.1097/RLI.0000000000000115.
8. Novikov DS, Veraart J, Jelescu IO, Fieremans E. Rotationally-invariant mapping of scalar and orientational metrics of neuronal microstructure with diffusion MRI. Neuroimage 2018;174:518–538 doi: 10.1016/j.neuroimage.2018.03.006.
9. Lampinen B, Szczepankiewicz F, Mårtensson J, et al. Towards unconstrained compartment modeling in white matter using diffusion-relaxation MRI with tensor-valued diffusion encoding. Magn. Reson. Med. 2020;early view doi: 10.1002/mrm.28216.
10. Alsop DC, Detre JA, Golay X, et al. Recommended implementation of arterial spin-labeled perfusion MRI for clinical applications: A consensus of the ISMRM perfusion study group and the European consortium for ASL in dementia. Magn. Reson. Med. 2015;73:102–116 doi: 10.1002/mrm.25197.
11. Rouvière O, Yin M, Dresner MA, et al. MR Elastography of the Liver: Preliminary Results. Radiology 2006;240:440–448 doi: 10.1148/radiol.2402050606.
12. Poorter J De, Wagter C De, Deene Y De, Thomsen C, Ståhlberg F, Achten E. Noninvasive MRI Thermometry with the Proton Resonance Frequency (PRF) Method:In Vivo

Results in Human Muscle. Magn. Reson. Med. 1995;33:74–81 doi: 10.1002/mrm.1910330111.
13. Giorgio A, De Stefano N. Effective Utilization of MRI in the Diagnosis and Management of Multiple Sclerosis. Neurol. Clin. 2018;36:27–34 doi: 10.1016/j.ncl.2017.08.013.
14. Enzinger C, Barkhof F, Ciccarelli O, et al. Nonconventional MRI and microstructural cerebral changes in multiple sclerosis. Nat. Rev. Neurol. 2015;11:676–686 doi: 10.1038/nrneurol.2015.194.
15. Feinberg DA, Setsompop K. Ultra-fast MRI of the human brain with simultaneous multi-slice imaging. J. Magn. Reson. 2013;229:90–100 doi: 10.1016/j.jmr.2013.02.002.
16. Barth M, Breuer F, Koopmans PJ, Norris DG, Poser BA. Simultaneous multislice (SMS) imaging techniques. Magn. Reson. Med. 2016;75:63–81 doi: 10.1002/mrm.25897.
17. Hutter J, Slator PJ, Christiaens D, et al. Integrated and efficient diffusion-relaxometry using ZEBRA. Sci. Rep. 2018;8:15138 doi: 10.1038/s41598-018-33463-2.
18. Benjamini D, Basser PJ. Use of marginal distributions constrained optimization (MADCO) for accelerated 2D MRI relaxometry and diffusometry. J. Magn. Reson. 2016;271:40–45 doi: 10.1016/j.jmr.2016.08.004.
19. Kim D, Doyle EK, Wisnowski JL, Kim JH, Haldar JP. Diffusion-relaxation correlation spectroscopic imaging: A multidimensional approach for probing microstructure. Magn. Reson. Med. 2017;78:2236–2249 doi: 10.1002/mrm.26629.
20. Ning L, Gagoski B, Szczepankiewicz F, Westin C-F, Rathi Y. Joint RElaxation-Diffusion Imaging Moments to Probe Neurite Microstructure. IEEE Trans. Med. Imaging 2020;39:668–677 doi: 10.1109/TMI.2019.2933982.
21. Veraart J, Novikov DS, Fieremans E. TE dependent Diffusion Imaging (TEdDI) distinguishes between compartmental T2 relaxation times. Neuroimage 2018;182:360–369 doi: 10.1016/j.neuroimage.2017.09.030.
22. Gong T, Tong Q, He H, Sun Y, Zhong J, Zhang H. MTE-NODDI: Multi-TE NODDI for disentangling non-T2-weighted signal fractions from compartment-specific T2 relaxation times. Neuroimage 2020;217:116906 doi: 10.1016/j.neuroimage.2020.116906.
23. Zhang H, Hubbard P, Parker G, Alexander D. Axon diameter mapping in the presence of orientation dispersion with diffusion MRI Jiang T, Navab N, Pluim JPW, Viergever MA, editors. Neuroimage 2011;56:640–647.
24. Jones DK, Horsfield MA, Simmons A. Optimal strategies for measuring diffusion in anisotropic systems by magnetic resonance imaging. Magn. Reson. Med. 1999;42:515–525 doi: 10.1002/(SICI)1522-2594(199909)42:3<515::AID-MRM14>3.0.CO;2-Q.
25. Drobnjak I, Siow B, Alexander D. Optimizing gradient waveforms for microstructure sensitivity in diffusion-weighted MR. J. Magn. Reson. 2010;206:41–51.
26. Drobnjak I, Alexander D. Optimising time-varying gradient orientation for microstructure

sensitivity in diffusion-weighted MR. J. Magn. Reson. 2011;212:344–354.
27. Grussu F, Schneider T, Zhang H, Alexander DC, Wheeler–Kingshott CAM. Neurite orientation dispersion and density imaging of the healthy cervical spinal cord in vivo. Neuroimage 2015;111:590–601 doi: 10.1016/j.neuroimage.2015.01.045.
28. Freeman A., Gowland P., Mansfield P. Optimization of the ultrafast look-locker echo-planar imaging T1 mapping sequence. Magn. Reson. Imaging 1998;16:765–772 doi: 10.1016/S0730-725X(98)00011-3.
29. Anastasiou A, Hall LD. Optimisation of T2 and M0 measurements of bi-exponential systems. Magn. Reson. Imaging 2004;22:67–80 doi: 10.1016/j.mri.2003.05.005.
30. Alexander DC. A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. Magn. Reson. Med. 2008;60:439–448 doi: 10.1002/mrm.21646.
31. Lemke A, Stieltjes B, Schad LR, Laun FB. Toward an optimal distribution of b values for intravoxel incoherent motion imaging. Magn. Reson. Imaging 2011;29:766–776 doi: 10.1016/j.mri.2011.03.004.
32. Knutsson H. Towards Optimal Sampling in Diffusion MRI. In: Proc of 2018 MICCAI workshop on Computational Diffusion MRI. ; pp. 3–18. doi: 10.1007/978-3-030-05831-9_1.
33. Turner R, Oros-Peusquens A-M, Romanzetti S, Zilles K, Shah NJ. Optimised in vivo visualisation of cortical structures in the human brain at 3 T using IR-TSE. Magn. Reson. Imaging 2008;26:935–942 doi: 10.1016/j.mri.2008.01.043.
34. Filipiak P, Fick R, Petiet A, et al. Reducing the number of samples in spatiotemporal dMRI acquisition design. Magn. Reson. Med. 2019;81:3218–3233 doi: 10.1002/mrm.27601.
35. Battiston M, Grussu F, Ianus A, et al. An optimized framework for quantitative magnetization transfer imaging of the cervical spinal cord in vivo. Magn. Reson. Med. 2018;79:2576–2588 doi: 10.1002/mrm.26909.
36. Chuhutin A, Hansen B, Jespersen SN. Precision and accuracy of diffusion kurtosis estimation and the influence of b-value selection. NMR Biomed. 2017;30:e3777 doi: 10.1002/nbm.3777.
37. Gudbjartsson H, Patz S. The rician distribution of noisy mri data. Magn. Reson. Med. 1995;34:910–914 doi: 10.1002/mrm.1910340618.
38. Hutchinson EB, Avram A V., Irfanoglu MO, et al. Analysis of the effects of noise, DWI sampling, and value of assumed parameters in diffusion MRI models. Magn. Reson. Med. 2017;78:1767–1780 doi: 10.1002/mrm.26575.
39. Andre JB, Bresnahan BW, Mossa-Basha M, et al. Toward Quantifying the Prevalence, Severity, and Cost Associated With Patient Motion During Clinical MR Examinations. J. Am. Coll. Radiol. 2015;12:689–695 doi: 10.1016/j.jacr.2015.03.007.
40. Grussu F, Blumberg SB, Battiston M, et al. SARDU-Net: a new method for model-free,

data-driven experiment design in quantitative MRI. In: Proc of International Society for Magnetic Resonance in Medicine (ISMRM). ; 2020.
41. Wang S, Peng Y, Medved M, et al. Hybrid multidimensional T2 and diffusion-weighted MRI for prostate cancer detection. J. Magn. Reson. Imaging 2014;39:781–788 doi: 10.1002/jmri.24212.
42. Chatterjee A, Bourne RM, Wang S, et al. Diagnosis of Prostate Cancer with Noninvasive Estimation of Prostate Tissue Composition by Using Hybrid Multidimensional MR Imaging: A Feasibility Study. Radiology 2018;287:864–873 doi: 10.1148/radiol.2018171130.
43. Lemberskiy G, Fieremans E, Veraart J, Deng F-M, Rosenkrantz AB, Novikov DS. Characterization of Prostate Microstructure Using Water Diffusion and NMR Relaxation. Front. Phys. 2018;6 doi: 10.3389/fphy.2018.00091.
44. Shrestha A, Mahmood A. Review of Deep Learning Algorithms and Architectures. IEEE Access 2019;7:53040–53065 doi: 10.1109/ACCESS.2019.2912200.
45. Ravi D, Wong C, Deligianni F, et al. Deep Learning for Health Informatics. IEEE J. Biomed. Heal. Informatics 2017;21:4–21 doi: 10.1109/JBHI.2016.2636665.
46. Bahadir CD, Dalca A V., Sabuncu MR. Learning-Based Optimization of the Under-Sampling Pattern in MRI. In: Proc of the International Conference on Information Processing in Medical Imaging (IPMI). ; 2019. pp. 780–792. doi: 10.1007/978-3-030-20351-1_61.
47. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. Nature 1986;323:533–536 doi: 10.1038/323533a0.
48. Kingma DP, Ba J. ADAM: A Method for Stochastic Optimization. Proc 3rd Int. Conf. Learn. Represent. 2015.
49. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 2014;15:1929–1958.
50. Usman M, Kakkar L, Kirkham A, Arridge S, Atkinson D. Model-based reconstruction framework for correction of signal pile-up and geometric distortions in prostate diffusion MRI. Magn. Reson. Med. 2019;81:1979–1992 doi: 10.1002/mrm.27547.
51. Veraart J, Novikov DS, Christiaens D, Ades-aron B, Sijbers J, Fieremans E. Denoising of diffusion MRI using random matrix theory. Neuroimage 2016;142:394–406 doi: 10.1016/j.neuroimage.2016.08.016.
52. Koay CG, Basser PJ. Analytically exact correction scheme for signal extraction from noisy magnitude MR signals. J. Magn. Reson. 2006;179:317–322 doi: 10.1016/j.jmr.2006.01.016.
53. Jenkinson M, Beckmann CF, Behrens TEJ, Woolrich MW, Smith SM. FSL. Neuroimage 2012;62:782–790 doi: 10.1016/j.neuroimage.2011.09.015.
54. Grussu F, Battiston M, Veraart J, et al. Multi-parametric quantitative in vivo spinal cord MRI with unified signal readout and image denoising. Neuroimage 2020 doi:

10.1016/j.neuroimage.2020.116884.
55. Golkov V, Dosovitskiy A, Sperl JI, et al. q-Space Deep Learning: Twelve-Fold Shorter and Model-Free Diffusion MRI Scans. IEEE Trans. Med. Imaging 2016;35:1344–1351 doi: 10.1109/TMI.2016.2551324.
56. Ye C. Tissue microstructure estimation using a deep network inspired by a dictionary-based framework. Med. Image Anal. 2017;42:288–299 doi: 10.1016/j.media.2017.09.001.
57. Palombo M, Singh S, Whitaker H, Punwani S, Alexander DC, Panagiotaki E. Relaxed-VERDICT: decoupling relaxation and diffusion for comprehensive microstructure characterization of prostate cancer. In: Proc of International Society for Magnetic Resonance in Medicine (ISMRM). ; 2020.
58. Liu C, Zoph B, Neumann M, et al. Progressive Neural Architecture Search. In: Proc of 2018 European Conference on Computer Vision (ECCV). ; 2018. pp. 19–35. doi: 10.1007/978-3-030-01246-5_2.