

Progressive Subsampling for Oversampled Data - Application to Quantitative MRI

Stefano B. Blumberg^{1*}, Hongxiang Lin^{1,3}, Francesco Grussu^{1,2},
Yukun Zhou¹, Matteo Figini¹, and Daniel C. Alexander¹

¹ University College London (UCL)

² Vall d'Hebron Barcelona Hospital

³ Zhejiang Lab

hxlin@zhejianglab.edu.cn, stefano.blumberg.17@ucl.ac.uk

Abstract. We present PROSUB: PROgressive SUBsampling, a deep learning based, automated methodology that subsamples an oversampled data set (e.g. channels of multi-channelled 3D images) with minimal loss of information. We build upon a state-of-the-art dual-network approach that won the MICCAI MUlTI-Diffusion (MUDI) quantitative MRI (qMRI) measurement sampling-reconstruction challenge, but suffers from deep learning training instability, by subsampling with a hard decision boundary. PROSUB uses the paradigm of recursive feature elimination (RFE) and progressively subsamples measurements during deep learning training, improving optimization stability. PROSUB also integrates a neural architecture search (NAS) paradigm, allowing the network architecture hyperparameters to respond to the subsampling process. We show PROSUB outperforms the winner of the MUDI MICCAI challenge, producing large improvements $>18\%$ MSE on the MUDI challenge sub-tasks and qualitative improvements on downstream processes useful for clinical applications. We also show the benefits of incorporating NAS and analyze the effect of PROSUB's components. As our method generalizes beyond MRI measurement selection-reconstruction, to problems that subsample and reconstruct multi-channelled data, our code is [7].

Keywords: Magnetic Resonance Imaging (MRI) Protocol Design, Recursive Feature Elimination, Neural Architecture Search

1 Introduction

Multi-modal medical imaging gives unprecedented insight into the microstructural composition of living tissues, and provides non-invasive biomarkers that hold promise in several clinical contexts. In particular, quantitative MRI fits a model in each pixel of a multi-channel acquisition consisting of multiple images each with unique contrast obtained by varying multiple MRI acquisition parameters, see e.g. [22]. This provides pixel-wise estimates of biophysical tissue properties [16]. In spite of this potential, comprehensively sampling high-dimensional acquisition

* Corresponding Author

spaces leads to prohibitively long acquisition times, which is a key barrier to more widespread adoption of qMRI in clinical use.

The MUlti-Diffusion (MUDI) MRI challenge [1,30] addressed this by providing data covering a densely-sampled MRI acquisition space (3D brain images with 1344 channels). The task was to reconstruct the full set of measurements from participant-chosen measurements from a small subsample, i.e. to obtain economical, but maximally informative acquisition protocols for any model that the full data set supports. That involves two sub-tasks: selecting the most informative measurements, and reconstructing the full data set from them. The challenge winner was SARDU-Net [17,18,30] with a dual-network strategy, that respectively subsamples the measurements, then reconstructs the full dataset from the subsampled data. However, SARDU-Net selects different sets of measurements with a hard decision boundary on each training batch, altering the second network’s input across different batches. This can cause instability, suupl. mat. fig. 3 and [8,25] show similar issues produce training instability. Furthermore, the popularity of paradigms such as recursive feature elimination (RFE), suggests that subsampling all of the measurements required immediately, is suboptimal. These two issues may lead to substandard performance.

We propose PROSUB, a novel automated methodology that selects then reconstructs measurements from oversampled data. Unlike classical approaches to experiment design [4], we approach the MUDI challenge in a new model-independent way. PROSUB builds upon the SARDU-Net by (i) using a form of RFE which progressively removes measurements across successive RFE steps and (ii) learning an average measurement score across RFE steps, which chooses the measurements to remove or preserve. This enhances the stability of our optimization procedure. Within each RFE step, PROSUB (iii) progressively subsamples the required measurements during deep learning training, building upon [8,25], that improves training stability. Also, PROSUB (iv) incorporates a generic neural architecture search (NAS) paradigm in concurrence to the RFE – so the architectures may respond to the measurement subsampling process.

Our implementation [7] is based on AutoKeras [24], KerasTuner [29]. PROSUB outperforms the SARDU-Net and SARDU-Net with AutoKeras NAS by $> 18\%$ MSE on the publicly available MUDI challenge data [2,30]. We show qualitative improvements on downstream processes: T2*,FA,T1,Tractography useful in clinical applications [5,10,14,20,27]. We examine the effect of how PROSUB’s components, including NAS, improve performance. We release the code [7] as PROSUB is not limited to subsampling MRI data for microstructure imaging.

2 Related Work and Preliminaries

Problem Setting Suppose we have an oversampled dataset $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbb{R}^{n \times N}$ where each sample has N measurements $\mathbf{x}_i \in \mathbb{R}^N$. The aim is to subsample $M < N$ measurements $\tilde{\mathbf{x}} = \{\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_n\} \in \mathbb{R}^{n \times M}$, $\tilde{\mathbf{x}}_i \in \mathbb{R}^M$, with the same M elements of each \mathbf{x}_i in each $\tilde{\mathbf{x}}_i$. We aim to lose as little information as possible when choosing $\tilde{\mathbf{x}}$, thereby enabling the best recovery of the full data

set \mathbf{x} . We therefore have two interconnected problems i) choosing which measurements to subsample, ii) reconstructing the original measurements from the subsampled measurements. We achieve this by (i) constructing a binary mask m containing M ones and $N - M$ zeros so $\tilde{\mathbf{x}} = m \cdot \mathbf{x}$, (ii) with a neural network \mathcal{R} .

SARDU-Net and Dual-Network Approaches The SARDU-Net [17,18,30], which is used for model-free quantitative MRI protocol design and won the MUDI challenge [30], has two stacked neural networks, trained in unison. The first network learns weight w from \mathbf{x} . $N - M$ smallest values of w are clamped to 0 and the first network subsamples and selects the measurements, by outputting $\mathbf{x} \cdot w$. The second network then predicts the original data from $\mathbf{x} \cdot w$. Related dual-network approaches include [11], which processed large point clouds, differing from our problem, as we do not assume our data has a spatial structure. We build upon the SARDU-Net and we use it as a baseline in our experiments.

Recursive Feature Elimination (RFE) One of the most common paradigms for feature selection is RFE, which has a long history in machine learning [3,34]. Recursively over steps $t = 1, \dots, T$, RFE prunes the least important features based on some task-specific importance score, successively analyzing less and less features over successive steps. We use a form of RFE in PROSUB.

Neural Architecture Search (NAS) Selecting neural network architecture hyperparameters e.g. number of layers and hidden units, is a task-data-dependent problem, where the most common strategies are random search [6] or grid search [26]. NAS approaches, see e.g. [12], outperform classical approaches with respect to time required to obtain high-performing models and can broadly be seen as a subfield of Automated Machine Learning (AutoML) [21]. PROSUB uses a generic NAS paradigm which optimizes network architectures over successive steps $t = 1, \dots, T$ in an outer loop. In an inner loop, with fixed architecture (and fixed t), we perform standard deep learning training across epochs $e = 1, \dots, E$, caching network training and validation performance r_t^e after each epoch. At the end of step t , the previous losses $\{r_j^i : i \leq E, j \leq t\}$ are used to update the network architectures for step $t + 1$. Our implementation is based on AutoKeras [24], with KerasTuner [29], which has good documentation and functionality.

3 Methods

We address the interdependency of the sampling-reconstruction problem with a dual-network strategy in section 3.1, illustrated in fig. 1. In section 3.2 we progressively construct our mask m , used to subsample the measurements. PROSUB has an outer loop: steps $t = 1, \dots, T$ where we simultaneously perform NAS and RFE, choosing the measurements to remove via a score, averaged across the steps, whilst simultaneously updating the network architecture hyperparameters. For fixed t , we perform deep learning training as an inner loop across epochs




Fig. 1. The computational graph of PROSUB.

$e = 1, \dots, E$, where we learn the aforementioned score and also progressively subsample the measurements. We summarize PROSUB in algorithm 1.

3.1 Scoring-Reconstruction Networks

Inspired by [11,17], we use two neural networks $\mathcal{S}_t, \mathcal{R}_t$, trained in unison, to address the interdependency of our two interconnected problems.

Scoring Network The first network \mathcal{S}_t learns a preliminary score with a sigmoid activation in its last layer, to weight each measurement:

$$s_t = \mathcal{S}_t(\mathbf{x}) \quad s_t \in (0, 2)^{n \times N}. \quad (1)$$

Mask As described in section 2, we use an array $m_t \in [0, 1]^N$ as a mask

$$\tilde{\mathbf{x}}_t = m_t \cdot \mathbf{x} \quad (2)$$

to subsample the measurements. We describe in section 3.2, how we progressively and manually set m_t to have $N - M$ entries with 0.

Average Measurement Score To score each measurement in step t , we use an exponential moving average, commonly used in time-series analysis (e.g [19]), across the scores of previous steps $\bar{s}_1, \dots, \bar{s}_{t-1}$ and s_t , to obtain a better estimate of the score and reduce the effect of the current learnt score s_t , if network performance is poor. With moving average coefficient hyperparameter α_t we calculate

$$\bar{s}_t = \alpha_t \cdot s_t + (1 - \alpha_t) \cdot \bar{s}_{t-1} \quad (3)$$

and we use $\alpha_t = \frac{T-t}{T-1}$. The averaged score \bar{s}_t is used to weight the subsampled measurements and to construct the mask (described in section 3.2).

Reconstruction Network The second network \mathcal{R}_t takes the weighted subsampled measurements to estimate \mathbf{x} with $\hat{\mathbf{x}}$, then passed through *Loss* (we use MSE)

$$L = \text{Loss}(\hat{\mathbf{x}}, \mathbf{x}), \quad \hat{\mathbf{x}} = \mathcal{R}_t(\tilde{\mathbf{x}}_t \cdot \bar{s}_t) \quad (4)$$

and the gradients from L are then backpropagated through $\mathcal{R}_t, \mathcal{S}_t$.

Algorithm 1 PROSUB: PROgressive SUBsampling for Oversampled Data

Data and Task: $\mathbf{x} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathbb{R}^N$, $M < N$
Training and NAS: $1 \leq E_d \leq E$, $1 < T_1 < T$, $NAS \leftarrow AutoKeras$
Scoring and RFE: $\alpha_t \leftarrow \frac{T-t}{T-1}$, $D_t \leftarrow \approx \frac{N-M}{T-T_1+1}$
Initialize: $m_1 \leftarrow [1]^N$, $\bar{s}_0 \leftarrow [0]^N$

- 1: **for** $t \leftarrow 1, \dots, T_1, \dots, T$ **do** ▷ RFE and NAS steps
- 2: **if** $1 \leq t < T_1$ **then**
- 3: $D = \emptyset$ ▷ No measurements to subsample
- 4: **else if** $T_1 \leq t \leq T$ **then** ▷ Subsampling stage
- 5: $D = \underset{j=1, \dots, D_t}{\operatorname{argmin}} \{\bar{s}_{t-1}[j] : m_t[j] = 1\}$ ▷ Measurements to subsample Eq. 5
- 6: **end if**
- 7: **for** $e \leftarrow 1, \dots, E_d, \dots, E$ **do** ▷ Training and validation epoch
- 8: $m_t^e \leftarrow \max\{m_t - \frac{(e-E_d)\mathbb{1}_{e \geq E_d}}{E_d} \cdot \mathbb{1}_{i \in D}(i), 0\}$ ▷ Compute mask Eq. 6
- 9: $s_t^e = \mathcal{S}(\mathbf{x})$, $\tilde{\mathbf{x}}_t = m_t^e \cdot \mathbf{x}_t$ ▷ Forward pass Eq. 1,2
- 10: $\bar{s}_t = \alpha_t \cdot s_t + (1 - \alpha_t) \cdot \bar{s}_{t-1}$ ▷ Average measurement score Eq. 3
- 11: $r_t^e \leftarrow L(\tilde{\mathbf{x}}, \mathbf{x})$, $L = L(\tilde{\mathbf{x}}, \mathbf{x})$, $\hat{\mathbf{x}} = \mathcal{R}_t(\tilde{\mathbf{x}} \cdot \bar{s}_t)$ ▷ Forward/backward pass Eq.4
- 12: **end for**
- 13: Use NAS, $\{r_j^i : i \leq E, j \leq t\}$, to calculate $\mathcal{R}_{t+1}, \mathcal{S}_{t+1}$ ▷ Update architectures
- 14: $m_{t+1} \leftarrow m_t^E$, *cache* \bar{s}_t
- 15: **end for**
- 16: **return** $m_T, \bar{s}_T, \mathcal{R}_T$ – use as described in section 2

3.2 Constructing the Mask to Subsample the Measurements

We construct a mask m_t^e , used to subsample the measurements in section 2 and eq. 2. We progressively set $N - M$ entries of m_t^e to zero across NAS and RFE outer loop $t = 1, \dots, T$ and deep learning inner loop $e = 1, \dots, E$. We refer to algorithm 1 for clarity.

Outer Loop: Choosing the Measurements to Remove Following standard practise in RFE e.g. [3,34], we remove the measurements recursively, in our case, across steps $t = 1, \dots, T$ in alg.-line 1. We split the RFE in two stages, by choosing a dividing step, hyperparameter $1 < T_1 < T$.

In the first stage $t = 1, \dots, T_1$ the optimization procedure learns scores \bar{s}_t and optimizes the network architectures via NAS. In alg.-line 3, we choose no measurements to subsample ($D = \emptyset$) thus the mask $m_t^e = m_t$ is fixed in alg.-line 8.

In the second stage $t = T_1, \dots, T$, we perform standard RFE. We first choose a hyperparameter $D_t \in \mathbb{N}$ – the number of measurements to subsample in step t . In this paper, we remove the same number of measurements per step, so $D_t \approx \frac{N-M}{T-T_1+1}$. In alg.-line 5 we then choose the measurements to remove in RFE step t , which correspond to those with the lowest scores in the previous step

$$D = \underset{j=1, \dots, D_t}{\operatorname{argmin}} \{\bar{s}_{t-1}[j] : m_t[j] = 1\}, \quad m_t \in \{0, 1\}^N \quad (5)$$

where here m_t indicates whether the measurement has been removed in previous steps $< t$. Our rationale is since the subsampled measurements are weighted

by the score, used as inputs to the reconstruction network in eq. 4, setting lowest-scored values to 0 may have small effect on the performance (in eq. 4).

Inner Loop: Progressively Subsampling the Measurements by Altering the Mask During Training Given D – computed in the outer loop alg.-line 5 we progressively, manually, alter the mask m_t^e in the inner loop of deep learning training alg.-line 7, i.e. gradually setting the value of these measurements to 0 in $\tilde{\mathbf{x}}_t$. We are inspired by [8,25] which used a similar approach to improve training stability. We alter m_t^e across chosen epochs $e = E_d, \dots, 2 \cdot E_d - 1 \leq E$ for hyperparameter $E_d < \frac{E}{2}$, used in alg.-line 8, for indicator function \mathbb{I} :

$$m_t^e = \max \left\{ m_t - \frac{(e - E_d) \mathbb{I}_{e \geq E_d}}{E_d} \cdot \mathbb{I}_{i \in D}(i), 0 \right\}. \quad (6)$$

4 Experiments and Results

MUDI Dataset and Task Data of images from 5 subjects are from the MUDI challenge [1,30], publicly available at [2]. Data features a variety of diffusion and relaxometry (i.e. T1 and T2*) contrasts, and were acquired with the ZEBRA MRI technique [22]. The total acquisition time for these oversampled data sets was $\approx 1h$, corresponding to the acquisition of $N = 1344$ measurements in this dense parameter space, resulting in 5, 3D brain images with 1344 channels (here unique diffusion- T2* and T1- weighting contrasts), with $n \approx 558K$ brain voxels. Detailed information is in [22,30]. We used the same task as the MICCAI MUDI challenge [30], where the participants were asked to find the most informative subsets of size $M = 500, 250, 100, 50$ out of N , while also estimating the fully-sampled signals from each of these subsets, and the evaluation is MRI signal prediction MSE. The winner of the original challenge [1,30] was the aforementioned SARDU-Net [17,30]. In this paper, we also consider smaller subsets $M = 40, 30, 20, 10$.

Experimental Settings We did five-fold cross validation using two separate subjects for validation and testing. We compare PROSUB and PROSUB w/o NAS with four baselines: i) SARDU-Net-v1: winner of the MUDI challenge [17,30]; ii) SARDU-Net-v2: latest official implementation of (i) [15,18]; iii) SARDU-Net-v2-BOF: five runs of (ii) with different initializations, choosing the best model from the validation set; iv) SARDU-Net-v2-NAS: integrating (ii) with AutoKeras NAS. To reduce total computational time with NAS techniques, we performed all of the tasks in succession. We first use algorithm 1 with $T_1, T, M = 4, 8, 500$, then take the final model, as initialization for algorithm 1 with $T_1, T, M = 1, 5, 250$, performing this recursively for $M = 100, 50, 40, 30, 20, 10$, using the best model for each different M . Consequently, SARDU-Net-v2-BOF and the NAS techniques in table 1 are trained for approximately the same number of epochs. We performed a brief search for NAS hyperparameters. In figs. 2,3 we examined the effect of PROSUB’s components and present all hyperparameters.

Main Results We present quantitative results in table 1 and note PROSUB’s

Table 1. Whole brain Mean-Squared-Error between $N = 1344$ reconstructed measurements and N ground-truth measurements, on leave-one-out cross validation on five Multi-Diffusion (MUDI) challenge subjects. The SARDU-Net won the MUDI challenge.

		MUDI Challenge Subsamples M for $N = 1344$			
		500	250	100	50
SARDU-Net-v1 [17,30]	Baseline	1.45 ± 0.14	1.72 ± 0.15	4.73 ± 0.57	5.15 ± 0.63
SARDU-Net-v2 [15,18]	Baseline	0.88 ± 0.10	0.89 ± 0.01	1.36 ± 0.14	1.66 ± 0.10
SARDU-Net-v2-BOF [15,18]	Baseline	0.83 ± 0.10	0.86 ± 0.10	1.30 ± 0.12	1.67 ± 0.12
SARDU-Net-v2-NAS	Baseline	0.82 ± 0.13	0.99 ± 0.12	1.34 ± 0.26	1.76 ± 0.24
PROSUB w/o NAS	Ours	0.66 ± 0.08	0.67 ± 0.09	0.88 ± 0.07	1.54 ± 0.11
PROSUB	Ours	0.49 ± 0.07	0.61 ± 0.11	0.89 ± 0.11	1.35 ± 0.11
		$M = 40$	30	20	10
SARDU-Net-v1 [17,30]	Baseline	6.10 ± 0.79	21.0 ± 6.07	19.8 ± 9.26	22.8 ± 6.57
SARDU-Net-v2 [15,18]	Baseline	1.95 ± 0.12	2.27 ± 0.20	3.01 ± 0.45	4.41 ± 1.39
SARDU-Net-v2-BOF [15,18]	Baseline	1.86 ± 0.18	2.15 ± 0.23	2.61 ± 0.24	3.74 ± 0.66
SARDU-Net-v2-NAS	Baseline	2.23 ± 0.22	6.00 ± 7.14	2.82 ± 0.41	4.27 ± 1.66
PROSUB w/o NAS	Ours	1.81 ± 0.18	2.18 ± 0.17	2.72 ± 0.34	3.91 ± 0.22
PROSUB	Ours	1.53 ± 0.05	1.87 ± 0.19	2.50 ± 0.40	3.48 ± 0.55

large improvements $>18\%$ MSE over all four baselines on the MUDI challenge sub-tasks. Using the Wilcoxon one-sided signed-rank test, a non-parametric statistical test comparing paired brain samples, our methods improvements have p-values of $7.14\text{E}-08$, $9.29\text{E}-07$, $3.20\text{E}-06$, $9.29\text{E}-07$ over the four respective baselines with Bonferroni correction, thus are statistically significant. We provide qualitative comparisons on a random test subject in fig. 2 on downstream processes (T2*, FA, T1, Tractography), useful in clinical applications [5,10,14,20,27].

Discussion Without explicitly optimizing PROSUB’s network architecture and training hyperparameters (fixing PROSUB’s hyperparameters to the SARDU-Net hyperparameters at $M = 500$), the PROSUB still outperforms all SARDU-Net baselines for MUDI Challenge M . Concerning NAS, we note the SARDU-Net-v2-NAS generally underperforms the SARDU-Net w/o NAS. Examining network performance during NAS (e.g. fig. 3), this is due to SARDU-Net performance being unstable due to its hard measurement selection, thus its performance is vulnerable to changes in architecture. Passing poor results to the NAS then reduces the effectiveness of the NAS in identifying high-performing architectures for small M . In contrast, PROSUB’s progressive subsampling allows the NAS to identify better architectures than the SARDU-Net-v2-NAS. PROSUB also outperforms the PROSUB w/o NAS i.e. the NAS is able to identify better performing architectures than the architecture chosen by the SARDU-Net for $M = 500$. In tables 2,4,5, we analyze the effect of PROSUB’s components, note removing any of the three non-NAS contributions worsens performance; we also show increasing network capacity in the PROSUB does not necessarily improve performance; we also list the architectures chosen by the NAS for SARDU-Net-v2-NAS and PROSUB.




Fig. 2. Qualitative comparison of downstream processes with useful clinical applications [5,10,14,20,27], of reconstructed $N = 1344$ measurements from chosen $M = 50$ samples, on a random test subject. PROSUB results are visually closer to the gold standard than the baseline. As MUDI data provides combined diffusion and relaxometry information, to evaluate the practical impact of the different reconstructions we estimated T1, T2* values and DTI parameters with a dictionary-based approach in [13]. We show whole-brain probabilistic tractography examining reconstructed fibre tracts, colors correspond to direction, on multi-shell/tissue constrained spherical deconvolution via iFOD2 in [23]. Quantitative improvements are, Baseline - Ours MSE, e.g. FA: 0.006 - 0.004, on NODDI parametric maps [33]: 0.022 - 0.007 FICVF, 0.023 - 0.005 FISO, 0.020 - 0.013 ODI.

5 Future Work

In future work, we could add an additional cost function to address the cost of obtaining specific combination of measurements from sets of MRI acquisition parameters; and develop a novel NAS algorithm to account for the concurrence of the subsampling process and architecture optimization. Our approach extends to many other quantitative MRI applications e.g. [4,9,16], other imaging problems e.g. [31], and wider feature selection/experiment design problems e.g. [28,32].

Acknowledgements

We thank: Tristan Clark and the HPC team (James O’Connor, Edward Martin); MUDI Organizers (Marco Pizzolato, Jana Hutter, Fan Zhang); Amy Chapman, Luca Franceschi, Marcela Konanova and Shinichi Tamura; NVIDIA for donating the GPU for winning 2019 MUDI. Funding details SB: EPSRC and Microsoft scholarship, EPSRC grants M020533 R006032 R014019, NIHR UCLH Biomedical Research Centre, HL: Research Initiation Project of Zhejiang Lab (No.2021ND0PI02), FG: Fellowships Programme Beatriu de Pinós (2020 BP 00117), Secretary of Universities and Research (Government of Catalonia).

References

1. Multi-dimensional Diffusion (MUDI) MRI challenge 2019. <https://web.archive.org/web/20200209111131/http://cmic.cs.ucl.ac.uk/cdmri>
2. Multi-dimensional Diffusion (MUDI) MRI challenge 2019 data. <https://www.developingbrain.co.uk/data/>
3. Scikit-learn recursive feature elimination. https://scikit-learn.org/stable/modules/generated/sklearn.feature_selection.RFE.html
4. Alexander, D.C.: A general framework for experiment design in diffusion MRI and its application in measuring direct tissue-microstructure features. *Magnetic Resonance in Medicine* **60**(2), 439–448 (2008)
5. Andica, C., Kamagata, K., Hatano, T., Saito, Y., Ogaki, K., Hattori, N., Aoki, S.: MR biomarkers of degenerative brain disorders derived from diffusion imaging. *Journal of Magnetic Resonance Imaging* **52**(6), 1620–1636 (2020)
6. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. *Journal of Machine Learning Research* **13**, 281–305 (2012)
7. Blumberg, S.B.: PROSUB code. <https://github.com/sbb-gh/PROSUB>
8. Blumberg, S.B., Palombo, M., Khoo, C.S., Tax, C.M.W., Tanno, R., Alexander, D.C.: Multi-stage prediction networks for data harmonization. In: *Medical Image Computing and Computer Assisted Intervention (MICCAI)* (2019)
9. Brihuega-Moreno, O., Heese, F.P., Hall, L.D.: Optimization of diffusion measurements using cramer-rao lower bound theory and its application to articular cartilage. *Magnetic Resonance in Medicine* **50** (2003)
10. Deoni, S.C.L.: Quantitative relaxometry of the brain. *Topics in Magnetic Resonance Imaging* **21**(2), 101–113 (2010)
11. Dovrat, O., Lang, I., Avidan, S.: Learning to sample. In: *Computer Vision and Pattern Recognition (CVPR)* (2019)

12. Elsken, T., Metzen, J.H., Hutter, F.: Neural architecture search: A survey. *Journal of Machine Learning Research* **20**, 1–21 (2019)
13. Garyfallidis, E., et al.: DIPY, a library for the analysis of diffusion MRI data. *Frontiers in Neuroinformatics* **8** (2014)
14. Granziera, C., et al.: Quantitative magnetic resonance imaging towards clinical application in multiple sclerosis. *Brain* **144**(5), 1296–1311 (2021)
15. Grussu, F.: SARDU-Net code. <https://github.com/fragrussu/sardunet>
16. Grussu, F., Battiston, M., Veraart, J., Schneider, T., Cohen-Adad, J., Shepherd, T.M., Alexander, D.C., Fieremans, E., Novikov, D.S., Wheeler-Kingshott, C.A.G.: Multi-parametric quantitative in vivo spinal cord MRI with unified signal readout and image denoising. *NeuroImage* **217** (2020)
17. Grussu, F., Blumberg, S.B., Battiston, M., İanuş, A., Singh, S., Gong, F., Whitaker, H., Atkinson, D., Wheeler-Kingshott, C.A.M.G., Punwani, S., Panagiotaki, E., Mertzaniidou, T., Alexander, D.C.: SARDU-Net: a new method for model-free, data-driven experiment design in quantitative MRI. In: *International Society for Magnetic Resonance in Medicine (ISMRM)* (2020)
18. Grussu, F., Blumberg, S.B., Battiston, M., Kakkar, L.S., Lin, H., İanuş, A., Schneider, T., Singh, S., Bourne, R., Punwani, S., Atkinson, D., Gandini Wheeler-Kingshott, C.A.M., Panagiotaki, E., Mertzaniidou, T., Alexander, D.C.: Feasibility of data-driven, model-free quantitative MRI protocol design: Application to brain and prostate diffusion-relaxation imaging. *Frontiers in Physics* **9** (2021)
19. Hamilton, J.D.: *Time Series Analysis*. Princeton University Press (1994)
20. Henderson, F., Abdullah, K.G., Verma, R., Brem, S.: Tractography and the connectome in neurosurgical treatment of gliomas: the premise, the progress, and the potential. *Neurosurgical Focus* **48**(2) (2020)
21. Hutter, F., Kotthoff, L., Vanschoren, J. (eds.): *Automated Machine Learning - Methods, Systems, Challenges*. Springer (2019)
22. Hutter, J., Slator, P.J., Christiaens, D., Teixeira, R.P.A., Roberts, T., Jackson, L., Price, A.N., Malik, S., Hajnal, J.V.: Integrated and efficient diffusion-relaxometry using ZEBRA. *Scientific reports* **8**(1), 1–13 (2018)
23. Jeurissen, B., Tournier, J.D., Dhollander, T., Connelly, A., Sijbers, J.: Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data. *NeuroImage* **103**, 411–426 (2014)
24. Jin, H., Song, Q., Hu, X.: Auto-keras: An efficient neural architecture search system. In: *International Conference on Knowledge Discovery & Data Mining (KDD)* (2019)
25. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of GANs for improved quality, stability, and variation. In: *International Conference on Learning Representations (ICLR)* (2018)
26. Larochelle, H., Erhan, D., Courville, A., Bergstra, J., Bengio, Y.: An empirical evaluation of deep architectures on problems with many factors of variation. *International conference on Machine learning (ICML)* (2007)
27. Lehericy, S., C, E.R., Goizet, Mochel, F.: MRI of neurodegeneration with brain iron accumulation. *Current Opinion in Neurology* **33**(4), 462–473 (2020)
28. Marinescu, R.V., et al.: TADPOLE challenge: Accurate alzheimer’s disease prediction through crowdsourced forecasting of future data. In: *Predictive Intelligence in Medicine*. pp. 1–10. Springer International Publishing, Cham (2019)
29. O’Malley, T., Bursztein, E., Long, J., Chollet, F., Jin, H., Invernizzi, L., et al.: *KerasTuner* (2019)
30. Pizzolato, M., Palombo, M., Bonet-Carne, E., Grussu, F., İanus, A., Bogusz, F., Pieciak, T., Ning, L., Blumberg, S.B., Mertzaniidou, T., Alexander, D.C., Afzali,

- M., Aja-Fernández, S., Jones, D.K., Westin, C.F., Rathi, Y., Baete, S.H., Cordero-Grande, L., Ladner, T., Slator, P.J., Christiaens, D., Thiran, J.P., Price, A.N., Seppehrband, F., Zhang, F., Hutter, J.: Acquiring and predicting MUlti-dimensional Diffusion (MUDI) data: an open challenge. In: Computational Diffusion MRI (CDMRI) Workshop in Medical Image Computing and Computer Assisted Intervention (MICCAI) (2019)
31. Prevost, R., Buckley, D.L., Alexander, D.C.: Optimization of the DCE-CT protocol using active imaging. In: 2010 IEEE International Symposium on Biomedical Imaging (ISBI): From Nano to Macro pp. 776–779 (2010)
 32. van der Putten, P., van Somere, M.: CoIL challenge 2000: The insurance company case. Institute of Advanced Computer Science Technical Report (2000)
 33. Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C.: NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. *NeuroImage* **61**(4) (2012)
 34. Zheng, A., Casari, A.: Feature Engineering for Machine Learning: Principles and Techniques for Data Scientists. O'Reilly (2018)

Supplementary Materials




Fig. 3. Left: Training losses of the SARDU-Net $M = 10$ for different seeds/initializations (colors), the SARDU-Net selects different sets of measurements with a hard decision boundary on each training batch, altering the second network’s input across different batches, producing instability. Right: Progressively subsampling measurements during training eq. 6, compared to hard measurement selection (table 2-row-3).

Table 2. Ablation study of PROSUB’s components (top-to-bottom) i) not computing an average measurement score, ii) not using RFE, iii) not progressively subsampling measurements during training (also fig. 3 right). No NAS, across five random seeds / initializations on a single validation subject, $T_1, T, M = 2, 6, 50$.

Formula Change	Val MSE
Alg.-line 10 is $\bar{s}_t = s_t, s_t \in \mathbb{R}^N$	3.02 ± 0.01
$T_1 = T = 6, D[T_1] = N - M$	1.86 ± 0.03
Alg.-line 8 is $m_t^e = \max \{m_t - (e - E_d) \mathbb{I}_{e \geq E_d} \cdot \mathbb{I}_{i \in D}(i), 0\}$	1.52 ± 0.03
PROSUB w/o NAS	1.47 ± 0.08

Table 3. All hyperparameters. We conducted a brief search for the NAS hyperparameters, which ex. dropout are the same between PROSUB and SARDU-Net+NAS. SARDU-Net w/o NAS baselines are from official implementation. 1st/2nd net. refers to respective fist/second network in the dual-network implementations. ‘choices in’ means NAS can choose one of the following values e.g. #Layers in 1st net = 2 choices in $\{1, 2, 3\} + 1$ means the first layer in PROSUB S_t is replaced by 1, 2, 3 layers.

Hyperparameters	Value	SARDU-Net	SARDU-Net-NAS	PROSUB
E , Optimizer, Learning Rate	200, <i>ADAM</i> , $1E-3$	✓	✓	✓
Batch Size, Weight Initialization	1500, He Normal	✓	✓	✓
Data Normalization	Max-99% " measurement-wise	✓	✓	✓
T_1, T	4, 36 + 9		✓	✓
α_t, D_t, E_d	$\frac{T-t}{T-1}, \frac{N-M}{T-t+1}, 20$		✓	✓
NAS adapted from	AutoKeras w. Greedy (default) Strategy		✓	✓
#Layers in 1st/2nd net	3 2 choices in $\{1, 2, 3\} + 1$	✓	✓	✓
#Units in layer 1,2 1st net.	1063, 781 ($M = 500$) 417, 333 ($M < 500$) 2 choices in $\{128, 256, \dots, 2048\}$	✓	✓	✓
#Units in layer 1,2 2nd net.	781, 1063 2 choices in $\{128, 256, \dots, 2048\}$	✓	✓	✓
Dropout	0.2 4 choices in $\{0, 0.1, 0.2, 0.3, 0.4\}$ 0	✓	✓	✓

Table 4. PROSUB results with the same experimental settings as table 1, at maximum network size of the NAS search space (see table 3). Increasing the network capacity does not necessarily improve performance.

M=500	250	100	50	40	30	20	10
1.38 ± 0.14	1.46 ± 0.19	1.79 ± 0.13	2.34 ± 0.26	2.61 ± 0.38	2.92 ± 0.30	3.39 ± 0.33	4.68 ± 1.36

Table 5. Architectures chosen by the AutoKeras Neural Architecture Search (NAS) in table 1 for respective network: Selector/Scoring | Prediction/Reconstruction, omitting first and last $N = 1344$ units in both networks. The (No-NAS) SARDU-Net and PROSUB w/o NAS have architectures $N \rightarrow 1063 \rightarrow 781 \rightarrow N$, $N \rightarrow 781 \rightarrow 1063 \rightarrow N$, the SARDU-Net used a smaller selector network: $N \rightarrow 417 \rightarrow 333 \rightarrow N$ for $M < 500$.

Split	Network	M=500	250	100	50
1	SARDU-Net-NAS	417,333 781,1063	417,333 781,1063	417 ² ,333 781,1063	417,333 781,1063
	PROSUB	1063,896 781 ² ,1063	1063,896 781 ² ,1063	1063,896 781 ² ,1063	1063,512,896 781 ² ,1063
2	SARDU-Net-NAS	417,333 781,1063	417,333 ² 781,1063	417,333 781,1063	417,333 781,1063
	PROSUB	1063,896,781 ² 781,1063	1063,640 781,1063	1063,640 781,1063	640,128,640 781,1063
3	SARDU-Net-NAS	417,256 781,1063	417,256 781,1063	417 ² ,256 781,1063	512,256 781,1063
	PROSUB	1063,2048 781,1063	1063,1536 781,1063	1063,2048 781,1063	1024,2048 781,1063
4	SARDU-Net-NAS	417,333 781,1063	417,333,333 781,1063	417,333 781,1063	417,1920,417,333 781,1063
	PROSUB	1063,1024,781 ² 781 ² ,1063	1063,1024,781 ² 781,1063	1063,1024,781,1920 781 ² ,1063	1063 ³ ,1024,781 ² 781 ² ,1063
5	SARDU-Net-NAS	417,1152 781,1063	417,333 ³ 781,1063	417,333 781,1063	417 ² ,333 781,1063
	PROSUB	1063,781 ² 781,1063	1063,781 ² 781,1063	1063,781,896 781,1063	1063,1664,1152,781 ² 781,1063
		M=40	30	20	10
1	SARDU-Net-NAS	417,333 781,1063 ³	417,333 781,1063	417,333 781,1063	417,333 781,1063
	PROSUB	1063,896 781,7781,1063	1063,1280,896 781 ² ,1063	1063,384 781 ² ,1063	1063,1408 781 ² ,1063
2	SARDU-Net-NAS	417,333 781,1063 ²	1024,333 781,1063	417,333 781,1063	417,333 781,1063
	PROSUB	1063,640 781,1063	1063,640 781 ³ ,1063	1063,640 781,1063	1063,640 781,1063
3	SARDU-Net-NAS	417,256 781,1063 ²	512,256 781,1063	417,333 781,1063	417,384 781,1063
	PROSUB	1063,2048 781,1063	1063,2048 781,1063	1063,384 781,1063	1063,2048 781,1063
4	SARDU-Net-NAS	417,256 781,1063	1920,333 781,1063	417,333 781,1063	640,256,333 781,1063
	PROSUB	1063,1024,781 ² 781,1063	1063,1024,781 ² 781 ² ,1063	1063,1024,781 ² 781 ² ,1063 ³	1063,1024,781 ² 781 ² ,1063
5	SARDU-Net-NAS	417,333 781,1063 ²	417,333 781,1063	417,333 781,1063	417,333 781,1063
	PROSUB	1063,781 ² 781,1063	1063,781 ² 781 ² ,1063	1063,781 ² 781,1063	1063,781 ² 781,1063