

On learning adaptive acquisition policies for undersampled multi-coil MRI reconstruction

Tim Bakker^{*1}

¹ *University of Amsterdam*

T.B.BAKKER@UVA.NL

Matthew Muckley²

Adriana Romero-Soriano²

Michal Drozdal^{†2}

Luis Pineda^{†2}

² *Facebook AI Research*

MMUCKLEY@FB.COM

ADRIANARS@FB.COM

MDROZDZAL@FB.COM

LEP@FB.COM

Abstract

Most current approaches to undersampled multi-coil MRI reconstruction focus on learning the reconstruction model for a fixed, equidistant acquisition trajectory. In this paper, we study the problem of joint learning of the reconstruction model together with acquisition policies. To this end, we extend the End-to-End Variational Network with learnable acquisition policies that can adapt to different data points. We validate our model on a coil-compressed version of the large scale undersampled multi-coil fastMRI dataset using two undersampling factors: $4\times$ and $8\times$. Our experiments show on-par performance with the learnable non-adaptive and handcrafted equidistant strategies at $4\times$, and an observed improvement of more than 2% in SSIM at $8\times$ acceleration, suggesting that potentially-adaptive k -space acquisition trajectories can improve reconstructed image quality for larger acceleration factors. However, and perhaps surprisingly, our best performing policies learn to be explicitly non-adaptive.

Keywords: MRI reconstruction, undersampled multi-coil MRI, adaptive acquisition.

1. Introduction

Magnetic resonance imaging (MRI) is one of the best non-invasive methods for assessing soft-tissue structure in the clinic. However, widespread MRI adoption is limited due to its long acquisition times. Almost since its inception, substantial research has been done to reduce these acquisition times, yielding a variety of undersampled MRI reconstruction techniques such as parallel imaging (Sodickson and Manning, 1997; Pruessmann et al., 1999; Griswold et al., 2002), compressed sensing (Lustig et al., 2007) and deep learning (DL) (Schlemper et al., 2017; Hammernik et al., 2018; Aggarwal et al., 2018). Although the DL-based approaches have been shown to achieve the strongest results, they tend to use either fixed or random k -space sampling patterns that do not adapt to the data, which may be suboptimal and lead to underestimation of the maximum possible acceleration rates.

There is a substantial literature - going back decades - on designing sampling trajectories for MRI (e.g. Cao and Levin (1993); Gao and Reeves (2000); Seeger et al. (2009); Haldar

^{*} Majority of work was done while interning at Facebook AI Research.

[†] Contributed equally.

and Kim (2019)). Since the introduction of deep learning, researchers have attempted to further improve DL-based MRI reconstruction by learning k -space sampling patterns from the data. These learning-based approaches result in either *non-adaptive* or *adaptive* sampling patterns – often referred to as policies. Non-adaptive policies learn a dataset specific acquisition trajectory – *e.g.*, each data point in the dataset is reconstructed following the same learnt trajectory –, while adaptive policies are conditioned on the initial k -space measurements and, as a result, have the potential to produce different sampling trajectories per data point. Non-adaptive policies have been shown to outperform handcrafted sampling patterns for both single-coil (Bahadir et al., 2019; Weiss et al., 2020; Huijben et al., 2020) and multi-coil (Zhang et al., 2020; Wang et al., 2021; Zibetti et al., 2021) acquisition settings. However, adaptive policies have only been devised in the single-coil setting (Zhang et al., 2019; Jin et al., 2019; Sanchez et al., 2020; Pineda et al., 2020; Bakker et al., 2020; Yin et al., 2021; Van Gorp et al., 2021), showing promising results which in some cases outperform non-adaptive ones. Learning adaptive policies for deep multi-coil MRI remains, to the best of our knowledge, largely unexplored.

In this work, we are the first to devise such a model for joint learning of 2D deep learning MRI reconstruction together with adaptive k -space acquisition trajectories for the more clinically relevant *multi-coil* setting. In particular, we extend recent work that learns adaptive acquisition trajectories (Yin et al., 2021) to the multi-coil scenario and enhance the End-to-End Variational Network (E2E VarNet) (Sriram et al., 2020), a standard model for multi-coil reconstruction, with the ability to learn dataset-specific as well as potentially adaptive k -space sampling strategies. We perform extensive evaluation on Cartesian sampling for 2D MRI using the multi-coil fastMRI knee dataset (Knoll et al., 2020b) on $4\times$ and $8\times$ acceleration. We hope our effort provides the community a point of departure for further research into adaptive multi-coil acquisition. Our experiments¹ show that:

- On the $8\times$ setup, our learned policies improve $\sim 2\%$ in SSIM over the strongest baseline, highlighting the ability of potentially-adaptive k -space acquisition to improve MRI reconstruction under high acceleration factors.
- On the $4\times$ setup, the gain due to k -space trajectory optimisation reduces, with our policies performing on-par with the strongest competing method.
- Interestingly, our top performing policies learn to be explicitly non-adaptive, suggesting that adaptivity of the k -space acquisition trajectories may come at the expense of over-regularising the reconstruction model.

2. Preliminaries

2.1. Background

We consider a dataset \mathcal{D} of k -space measurements $\mathbf{y} \in \mathbb{C}^{N \times M}$ from which we can reconstruct MR images. In the *single-coil* setting, the reconstructed MR images can be obtained by the inverse Fourier transform \mathcal{F}^{-1} as $\mathbf{x} = \mathcal{F}^{-1}(\mathbf{y})$. However, modern scanners accelerate the acquisition of the k -space by using multiple receiver coils that are each sensitive to different

1. We’ve made our code and pre-trained models publicly available as part of the: [fastMRI repository](#).

regions of the anatomy, thus exploiting redundancies in k -space measurements (Sodickson and Manning, 1997; Pruessmann et al., 1999; Griswold et al., 2002). Hence, in the *multi-coil* setting, we define $\mathbf{y} \in \mathbb{C}^{N \times M \times K}$, where K is the number of coils and where $\mathbf{y}_i \in \mathbb{C}^{N \times M}$ represents the output of a measurement by the i -th coil. The reconstructed MR images can then be obtained as

$$\mathbf{x} = \sum_{i=1}^K \bar{S}_i \odot \mathcal{F}^{-1}(\mathbf{y}_i), \quad (1)$$

where \odot denotes element-wise multiplication, and \bar{S}_i is the complex-conjugate of the complex-valued sensitivity map associated with each receiver coil i , normalised such that $\sum_{i=1}^K \bar{S}_i S_i = 1$. These sensitivity maps encode how sensitive each coil is to each region in the anatomy, and can be estimated in an auto-calibrating fashion by fully sampling the center of the k -space, also known as the auto-calibration signal (ACS) region, with each coil. The acquisition of k -space measurements can be further accelerated by collecting fewer measurements and reconstructing the images using a partially observed k -space. To simulate partial observations of the k -space, we introduce a Cartesian binary sampling mask \mathbf{M} that selects $B < M$ measurements, including the ACS measurements, and define the *undersampled* k -space as $\tilde{\mathbf{y}}_i = \mathbf{M} \odot \mathbf{y}_i$, where \odot denotes element-wise multiplication. Note that the same mask is applied to the measurements from all coils. The reconstructed MR images can then be obtained by leveraging the undersampled k -space as

$$\tilde{\mathbf{x}} = \sum_{i=1}^K \bar{S}_i \odot \mathcal{F}^{-1}(\tilde{\mathbf{y}}_i). \quad (2)$$

This however results in blur or aliasing effects in the reconstructed images, which can be mitigated through the use of recent deep learning models, such as the End-to-End Variational Network (E2E VarNet) (Sriram et al., 2020). In particular, the E2E VarNet takes as input the partially observed k -space $\tilde{\mathbf{y}}$ along with the sampling mask \mathbf{M} decomposed into the mask of ACS measurements \mathbf{M}_{ACS} and the mask of the non-ACS measurements \mathbf{M}' , such that $\mathbf{M} = \mathbf{M}_{\text{ACS}} + \mathbf{M}'$. The model estimates the sensitivity maps from the ACS region, and uses a cascaded neural network, g , to produce a high fidelity image reconstruction, $\hat{\mathbf{x}} = g(\tilde{\mathbf{y}}, \mathbf{M}_{\text{ACS}}, \mathbf{M}'; \phi)$, where ϕ are learnable parameters. Note that \mathbf{M}' is commonly handcrafted to select equidistant measurements.

2.2. Problem formulation

Our goal is to adapt the sampling of measurements in \mathbf{M}' to each MR slice (image). To that end, we seek a policy that, given an initial undersampled k -space (*e.g.*, the ACS k -space, $\tilde{\mathbf{y}}_{\text{ACS}}$), predicts which measurements to acquire next. More precisely, we aim to learn an acquisition policy $\pi(\tilde{\mathbf{y}}_{\text{ACS}}, \mathbf{M}_{\text{ACS}}; \theta) \rightarrow \mathbf{M}'$, parameterised by θ , that selects the measurements to acquire in order to improve the image reconstruction process defined by g :

$$\phi^*, \theta^* = \arg \min_{\phi, \theta} \sum_j \mathcal{L}(\hat{\mathbf{x}}_j, \mathbf{x}_j), \quad (3)$$

where ϕ^* and θ^* are the optimised reconstruction and policy parameters respectively, j indexes the dataset, and \mathcal{L} is a loss function measuring the discrepancy between the model prediction $\hat{\mathbf{x}}$ and the image reconstructed from the fully sampled k -space \mathbf{x} .

3. Method

This section outlines our Policy network and details its integration with the E2E VarNet. An overview of the proposed system is depicted in Figure 1. From now on, we will use the capitalised ‘Policy’ to refer to our proposed model, while continuing to use ‘policy’ as general descriptor of ‘subsampling strategy’.

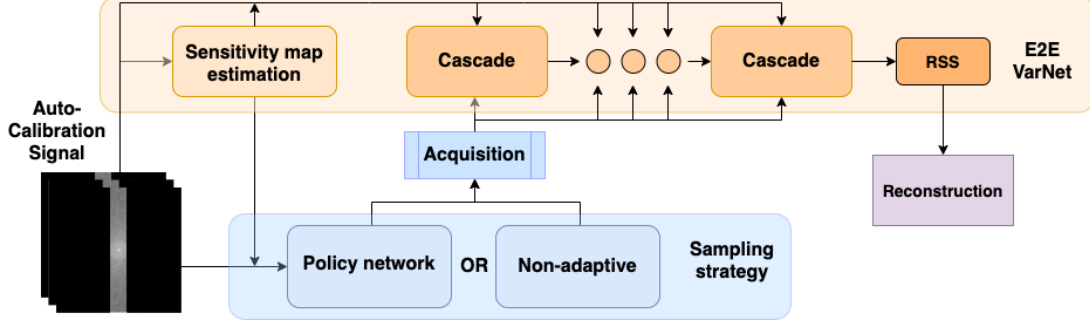


Figure 1: Overview of our system. The E2E VarNet computes sensitivity maps from the ACS, which are passed to the cascaded reconstruction model together with a subsampled k -space. The final reconstruction is reduced to a real-valued MR image by a root-sum-of-squares (RSS) operation. The subsampled k -space consists of the ACS and acquisitions suggested by a sampling strategy, which may be adaptive or non-adaptive.

3.1. Policy network

Our Policy is a neural network that takes as input the ACS k -space $\tilde{\mathbf{y}}_{\text{ACS}}$ and the mask of ACS measurements \mathbf{M}_{ACS} , and outputs a sampling probability for each measurement in k -space. From $\tilde{\mathbf{y}}_{\text{ACS}}$, we first estimate the sensitivity maps, and reconstruct a complex image $\tilde{\mathbf{x}}_{\text{ACS}}$ following Eq. 2. Using this image as input reduces the Policy network’s size while maintaining relevant information, and allows for taking inspiration from convolutional architectures that were successfully used in the single-coil adaptive MRI literature. In particular, we employ the Policy network of Bakker et al. (2020) and extend it to handle complex-valued inputs. The network is composed of a convolutional feature extractor, followed by a fully-connected block that outputs a heatmap encoding the relative salience of each potential k -space measurement in the acquisition step: see Appendix C.4 for details. It remains to normalise these values and then sample k -space measurements. To this end, Yin et al. (2021) have shown that straight-through estimation outperforms reinforcement learning based approaches for backpropagation through discrete sampling in the single-coil setting. We thus employ their formulation, which is as follows: a non-linearity (e.g. a softplus as in Yin et al. (2021) or a sigmoid as in Bahadir et al. (2019); Zhang et al. (2020)) is first applied to ensure non-negative values. To prevent already observed k -space measurements in the ACS region from being sampled again, we set their corresponding probabilities to 0. The resulting vector is normalised to obtain M independent realisations of a Bernoulli distribution from which to sample the measurements to be acquired. We employ rejection sampling to obtain exactly B measurements on each forward pass. As a

result of the sampling, we obtain the binary mask of measurements to be acquired, \mathbf{M}' . The aforementioned straight-through estimation - employed during backpropagation - is realised by treating the non-differentiable sampling discretisation as a sigmoid function with slope 10. In following Yin et al. (2021), we additionally enable a fairer comparison to the multi-coil baseline of Zhang et al. (2020) (see Section 4.2), which employs such straight-through estimation as well.

3.2. Integrating policy network and E2E VarNet

The original E2E VarNet takes as input a partially observed k -space $\tilde{\mathbf{y}}$, a binary sampling mask of ACS measurements \mathbf{M}_{ACS} , and a *predefined* binary sampling mask of non-ACS measurements \mathbf{M}' . In the adaptive acquisition setup \mathbf{M}' is instead *predicted* by the policy network, so we further decompose $\tilde{\mathbf{y}}$ into the ACS measurements $\tilde{\mathbf{y}}_{\text{ACS}}$ and the policy prediction $\tilde{\mathbf{y}}'$. As a result, our E2E VarNet takes as input $\tilde{\mathbf{y}}_{\text{ACS}}$, $\tilde{\mathbf{y}}'$, \mathbf{M}_{ACS} , and \mathbf{M}' . Following the original work, we first estimate sensitivity maps from the ACS measurements $\tilde{\mathbf{y}}_{\text{ACS}}$ by means of a U-Net (Ronneberger et al., 2015). The weights of the U-Net that predict the sensitivity maps are tied between the E2E VarNet and the Policy network, such that the sensitivity maps are re-used by both networks. Then, the output of the Policy network \mathbf{M}' is used to acquire $\tilde{\mathbf{y}}'$. As a result, we obtain a mask of observed measurements $\mathbf{M} = \mathbf{M}_{\text{ACS}} + \mathbf{M}'$ and a partially observed k -space $\tilde{\mathbf{y}} = \tilde{\mathbf{y}}_{\text{ACS}} + \tilde{\mathbf{y}}'$ that are used inside the E2E VarNet to produce a high fidelity MR image reconstruction. The original E2E VarNet is composed of cascaded modules that apply soft data consistency (DC) layers and refinement operations simultaneously. DC layers ensure that k -space predictions stay close to the observed k -space, while the refinement operations refine the k -space predictions by applying a U-Net to the corresponding complex-valued image-space representation. The cascaded and data consistency structure of the E2E VarNet offers several potential interface points with the policy model. After experimentation we chose to re-purpose the DC layers to input acquisitions made by the policy into the E2E VarNet pipeline. More precisely, we introduce hard DC layers into the E2E VarNet (Schlemper et al., 2017), which directly replace phantasised measurements in the k -space predictions with the observed measurements. To ensure that each cascade possesses all relevant information, we apply DC and refinement operations sequentially, rather than simultaneously. An E2E VarNet cascade now first applies a hard DC layer – which inputs acquired measurements, and restores changes to observed measurements due to a previous cascade –, followed by a refinement operation on the subsampled k -space $\tilde{\mathbf{y}}$.

4. Experiments

4.1. Data

We use the fastMRI multi-coil knee dataset (Zbontar et al., 2018) for all experiments, which contains 973 train volumes and 199 validation volumes of fully sampled k -space data. The test volumes are not fully sampled, and therefore cannot be used for our purposes. Our models treat every slice in a volume independently, resulting in 34,742 train slices and 7,135 validation slices to use as our dataset. For ease of experimentation, we reduce the number of coils by taking a Singular Value Decomposition and using the $K = 4$ coils with

the largest singular values (Buehrer et al., 2007). We further crop the MR slices to the center (128×128) region of k -space, see Appendix C.1. To simulate clinical conditions more closely, we create the ground truth target image from the *uncompressed* k -space by applying a coil-wise inverse Fourier transform followed by a root-sum-of-squares (RSS) on the resulting multi-coil image representation.

4.2. Baselines

We compare our method to two baselines: Equispaced (or Equisdistant) and LOUPE. Equispaced masks were shown by Hammernik et al. (2018) to be a strong hand-designed subsampling strategy for deep learning-reconstructed multi-coil MRI. Such masks have been long-used in the parallel imaging literature and are a current standard for clinical 2D imaging. At the same time, they depart from random patterns favoured by compressed sensing, which is applied for single-coil reconstruction. LOUPE is a dataset-specific (*i.e.*, non-adaptive) learned strategy that aims to optimise a single subsampling mask for the entire dataset, without conditioning on initial k -space measurements (Bahadir et al., 2019). Recently, Zhang et al. (2020) extended LOUPE to the multi-coil setting, and we use their method as implemented by Yin et al. (2021), which employs the same normalisation and straight-through estimation used by our Policy network.

4.3. Training details

All models are trained to optimise SSIM (Wang et al., 2004) using Adam (Kingma and Ba, 2014) for 50 epochs with a learning rate of 0.001, decaying it by a factor 10 on epoch 40 — the default fastMRI E2E VarNet training setting. We initialise the $4\times$ acceleration experiments with the 10 lowest frequency measurements k -space, and acquire 22 more measurements with our models. The $8\times$ acceleration experiments are initialised with the 4 lowest frequency measurements, and we acquire 12 more, instead. This initialisation corresponds, to the ACS used in the fastMRI E2E VarNet implementation to estimate sensitivity maps², and is thus a natural starting point for acquisition. See Appendix C for additional details. We empirically search over several hyperparameter settings for both the reconstruction model and the policy network. For the reconstructor, we consider either 5 or 7 cascades, and either 18 or 36 channels in the first layer of the refinement modules. We also explore the heatmap non-linearity mentioned in Section 3.1 and run experiments for both the sigmoid and softplus non-linearities, using slope $\in \{1, 5, 10\}$ and $\beta \in \{0.5, 1, 5\}$, respectively. Since LOUPE employs the same normalisation and straight-through estimation as our Policy network, we also explore these non-linearities for LOUPE. Unless otherwise specified, we report the best (averaged over seeds) run under the explored hyperparameters.

4.4. Results

We report our main results in Table 1, where we present validation SSIM for the best performing hyperparameter setting of each model. While the Policy network performs on-par with LOUPE at $4\times$ acceleration, it outperforms the best competing method at $8\times$ accelerations by 1.89 SSIM points. To further understand the gains obtained by the Policy

2. As of July 30th, 2021.

Table 1: Results on the undersampled multi-coil dataset. Policy performs on-par with the best competing model at $4\times$ and dominates performance at $8\times$.

	Policy	LOUPE	Equispaced
$4\times$	95.63 \pm 0.27	95.61 \pm 0.55	95.38 \pm 0.03
$8\times$	93.26 \pm 0.20	91.37 \pm 0.67	91.30 \pm 0.06

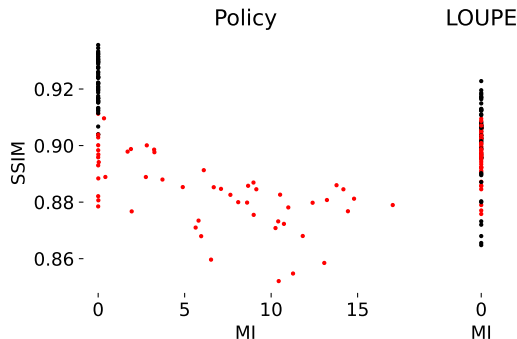


Figure 2: SSIM as a function of policy mutual information (MI), for the $8\times$ setting. Each dot is a single model. Red: softplus; black: sigmoid.

model, we inspect the adaptivity of the learned policies by plotting the SSIM score as a function of the mutual information (MI) between probability masks and images of all learned policies (Policy and LOUPE) — see Figure 2 for the $8\times$ acceleration results. We plot the policies that use the sigmoid non-linearity in black, and policies that use the softplus non-linearity in red. Surprisingly, we observe that the best performing Policy models exhibit zero mutual information, denoting no adaptivity — *i.e.*, the distribution of actions is constant for all data points in our dataset. We observe that early in the learning process all policies are adaptive and that some of them *learn to be non-adaptive*. Moreover, we observe that all sigmoid-based Policies end up being non-adaptive while the majority of the softplus-based Policies converge to adaptive strategies, suggesting that this non-linearity plays a crucial role in learning adaptivity. We hypothesise that it may be easier for Policies using the sigmoid non-linearity to learn non-adaptive strategies — which requires learning to ignore their input — given the saturation of the function for both very large positive and negative values. In contrast, the softplus non-linearity only saturates for large negative values, while the model is simultaneously encouraged to assign positive values to at least B actions when sampling from the distribution. Figure 3 displays some examples of image reconstructions and learnt subsamplings for the $8\times$ acceleration. The sigmoid-based Policy chooses a subset of actions to sample from with equal probability, whereas LOUPE appears to yield very-nearly sparse probabilities; exhibiting probability close to 1 for a limited set of actions, while most actions end up with a probability near 0. The more adaptive softplus-based Policy assigns less regular probability values, and seems to favour the center region less than both its sigmoid-based counterpart and LOUPE. Appendix D.2 contains additional qualitative results.

4.4.1. DISCUSSION

In this subsection, we outline three hypotheses as to why the best performing policies learn to be non-adaptive. The first one is model *generalisation*: the model may learn a strong adaptive Policy on the training set that does not generalise to the validation set. However,

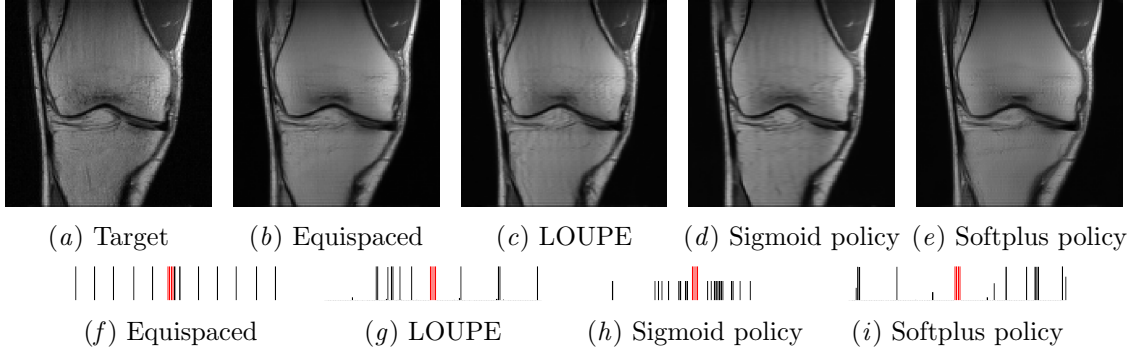


Figure 3: Qualitative results for $8\times$: (a) Ground truth, (b-e) reconstructions, and (f-g) visualisations of subsampling policies. Each policy depicts 128 probabilities; one per potential k -space measurement. The ACS region (red) has probability 1.

we examine the Policy adaptivity on the training set in Appendix D.1.2, and find no evidence that the lack of adaptivity is caused by overfitting the train data. The second hypothesis is connected to *amortisation* of the parameters of the reconstruction model over the large training dataset. Adaptivity can act as a regulariser on the reconstruction model, since a single reconstruction model needs to reconstruct MR images following multiple different k -space sampling patterns. In Appendix D.1.3, we report results for systems trained with reconstruction models of increasing capacities, and assess whether leveraging higher capacity reconstructors can help achieve Policy adaptivity. Although the higher capacity models lead to slight improvements in terms of SSIM, the best models are still not adaptive. Finally, we hypothesise that the estimation of sensitivity maps – a significant difference between single-coil and multi-coil reconstruction pipelines – may affect the adaptivity of the joint model. In the E2E VarNet, sensitivity maps are estimated independently per slice, and this may enable a form of adaptivity beyond the mask selection that we have explored here. However, the interaction between the sensitivity maps and the acquisition policies requires further investigation and is left as future work (see Appendix B). We display some learned sensitivity maps for all models in Appendix D.2.3.

5. Conclusion

We have explored the problem of jointly optimising an adaptive sampling strategy and a deep reconstruction model for multi-coil 2D MRI. To this end, we have proposed the first method for integrating a policy network with the E2E VarNet reconstruction model, and evaluated it on the large-scale, multi-coil fastMRI knee dataset. Our Policy networks outperform previous learning based non-adaptive approaches as well as the handcrafted equispaced masks at $8\times$ accelerations. Interestingly however, the Policies learn to be explicitly non-adaptive. The main limitation of our work is the analysis on coil-compressed (128×128) k -space acquisitions, rather than the full data. The relevance of our work for the original, un-cropped data should be empirically verified by future research.

Acknowledgments

T. Bakker’s PhD work is partially supported by the ‘Efficient Deep Learning’ (EDL, <https://efficientdeeplearning.nl>) research programme, which is financed by the Dutch Research Council (NWO) domain Applied and Engineering Sciences (TTW).

We are grateful to the Weights&Biases team ([Biewald, 2020](#)) for providing their experiment tracking software. We would also like to thank Edward Smith for useful discussions.

References

- Ben Adcock, Anders C Hansen, Clarice Poon, and Bogdan Roman. Breaking the coherence barrier: A new theory for compressed sensing. In *Forum of Mathematics, Sigma*, 2017.
- Hemant K Aggarwal, Merry P Mani, and Mathews Jacob. MoDL: Model-based deep learning architecture for inverse problems. *IEEE Transactions on Medical Imaging*, 2018.
- Cagla Deniz Bahadir, Adrian V. Dalca, and Mert R. Sabuncu. Learning-based optimization of the under-sampling pattern in MRI. *International Conference on Information Processing in Medical Imaging*, 2019.
- Tim Bakker, Herke van Hoof, and Max Welling. Experimental design for MRI by greedy policy search. In *NeurIPS*, 2020.
- Peter L. Bartlett and Jonathan Baxter. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 2001.
- M.A. Bernstein, K.F. King, and X.J. Zhou. *Handbook of MRI Pulse Sequences*. Academic Press, 2004.
- Lukas Biewald. Experiment tracking with Weights and Biases, 2020. URL <https://www.wandb.com/>. Software available from wandb.com.
- Martin Buehrer, Klaas P. Pruessmann, Peter Boesiger, and Sebastian Kozerke. Array compression for MRI with large coil arrays. *Magnetic Resonance in Medicine*, 2007.
- Y. Cao and D.N. Levin. Feature-recognizing MRI. *Magnetic Resonance in Medicine*, 1993.
- Nicolas Chauffert, Philippe Ciuciu, Jonas Kahn, and Pierre Weiss. Variable density sampling with continuous trajectories. *SIAM Journal on Imaging Sciences*, 2014.
- Jeffrey A. Fessler. Model-based image reconstruction for mri. *IEEE Signal Processing Magazine*, 2010.
- Y. Gao and S.J. Reeves. Optimal k-space sampling in MRSI for images with a limited region of support. *IEEE Transactions on Medical Imaging*, 2000.
- Mark A Griswold, Peter M Jakob, Robin M Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 2002.

- David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. *CoRR*, 2016.
- Justin P. Haldar and Daeun Kim. Oedipus: An experiment design framework for sparsity-constrained MRI. *IEEE Transactions on Medical Imaging*, 38:1545–1558, 2019.
- Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 2018.
- Iris AM Huijben, Bastiaan S Veeling, and Ruud JG van Sloun. Learning sampling and model-based signal recovery for compressed sensing MRI. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
- Kyong Hwan Jin, Michael Unser, and Kwang Moo Yi. Self-Supervised Deep Active Accelerated MRI. *arXiv e-prints*, 2019.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2014.
- Florian Knoll, Kerstin Hammernik, Chi Zhang, Steen Moeller, Thomas Pock, Daniel K. Sodickson, and Mehmet Akcakaya. Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues. *IEEE Signal Processing Magazine*, 2020a.
- Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2020b.
- Evan Levine and Brian Hargreaves. On-the-fly adaptive k-space sampling for linear MRI reconstruction using moment-based spectral analysis. *IEEE transactions on medical imaging*, 2017.
- Michael Lustig, David Donoho, and John Pauly. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 2007.
- Michael Lustig, David L. Donoho, Juan M. Santos, and John M. Pauly. Compressed sensing mri. *IEEE Signal Processing Magazine*, 2008.
- Andrew L. Maas, Awni Y. Hannun, and Andrew Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *in ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- Peter Mansfield. Multi-planar image formation using NMR spin echoes. *Journal of Physics C: Solid State Physics*, 1977.
- Vinod Nair and Geoffrey E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, 2010.

- Luis Pineda, Sumana Basu, Adriana Romero, Roberto Calandra, and Michal Drozdal. Active MR k-space sampling with reinforcement learning. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2020.
- Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 1999.
- Saiprasad Ravishankar and Yoram Bresler. Adaptive sampling design for compressed sensing MRI. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2011.
- Artem Razumov, Oleg Y Rogov, and Dmitry V Dylov. Optimal mri undersampling patterns for ultimate benefit of medical vision tasks. *arXiv preprint arXiv:2108.04914*, 2021.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 2015.
- Thomas Sanchez, Igor Krawczuk, Zhaodong Sun, and Volkan Cevher. Closed loop deep Bayesian inversion: Uncertainty driven acquisition for fast MRI. *OpenReview ICLR submission*, 2020.
- Jo Schlemper, Jose Caballero, Joseph V Hajnal, Anthony N Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Transactions on Medical Imaging*, 2017.
- Matthias Seeger, Hannes Nickisch, R. Pohmann, and Bernhard Schölkopf. Optimization of k-space trajectories for compressed sensing by Bayesian experimental design. *Magnetic Resonance in Medicine*, 63:116–26, 11 2009. doi: 10.1002/mrm.22180.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of Go without human knowledge. *Nature*, 2017.
- Daniel K Sodickson and Warren J Manning. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 1997.
- Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C. Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson. End-to-end variational networks for accelerated MRI reconstruction. In *MICCAI proceedings*, 2020.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. Instance normalization: The missing ingredient for fast stylization. *CoRR*, 2016.
- Hans Van Gorp, Iris Huijben, Bastiaan S Veeling, Nicola Pezzotti, and Ruud JG Van Sloun. Active deep probabilistic subsampling. In *International Conference on Machine Learning*, pages 10509–10518. PMLR, 2021.

- Hado Van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double Q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016.
- SS Vasanawala, MJ Murphy, Marcus T Alley, P Lai, Kurt Keutzer, John M Pauly, and Michael Lustig. Practical parallel imaging compressed sensing MRI: Summary of two years of experience in accelerating body MRI of pediatric patients. In *2011 IEEE International Symposium on Biomedical Imaging: From nano to macro*. IEEE, 2011.
- Guanhua Wang, Tianrui Luo, Jon-Fredrik Nielsen, Douglas C Noll, and Jeffrey A Fessler. B-spline parameterized joint optimization of reconstruction and k-space trajectories (BJORK) for accelerated 2D MRI. *arXiv preprint arXiv:2101.11369*, 2021.
- Zhou Wang, Alan Bovik, Hamid Sheikh, and Eero Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 2004.
- Tomer Weiss, Sanketh Vedula, Ortal Senouf, Alex Bronstein, Oleg Michailovich, and Michael Zibulevsky. Joint learning of Cartesian undersampling and reconstruction for accelerated MRI. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020.
- Katherine L. Wright, Jesse I. Hamilton, Mark A. Griswold, Vikas Gulani, and Nicole Seiberlich. Non-cartesian parallel imaging reconstruction. *Journal of Magnetic Resonance Imaging*, 2014.
- Tianwei Yin, Zihui Wu, He Sun, Adrian V. Dalca, Yisong Yue, and Katherine L. Bouman. End-to-end sequential sampling and reconstruction for MR imaging, 2021.
- Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: An open dataset and benchmarks for accelerated MRI. *CoRR*, 2018.
- Jinwei Zhang, Hang Zhang, Alan Wang, Qihao Zhang, Mert Sabuncu, Pascal Spincemaille, Thanh D Nguyen, and Yi Wang. Extending LOUPE for k-space under-sampling pattern optimization in multi-coil MRI. In *International Workshop on Machine Learning for Medical Image Reconstruction*, 2020.
- Zizhao Zhang, Adriana Romero, Matthew J Muckley, Pascal Vincent, Lin Yang, and Michal Drozdal. Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- Marcelo VW Zibetti, Florian Knoll, and Ravinder R Regatte. Alternating learning approach for variational networks and undersampling pattern in parallel MRI applications. *arXiv preprint arXiv:2110.14703*, 2021.

Supplementary material to “On learning adaptive acquisition policies for undersampled multi-coil MRI reconstruction”.

Appendix A. Related work

Our literature review focuses on online sampling design algorithms for gridded trajectories. As such, we do not cover other substantial fields of MRI acceleration work, such as on non-Cartesian trajectories, dynamic imaging, and 3D sampling (Bernstein et al., 2004; Lustig et al., 2008; Wright et al., 2014; Fessler, 2010; Knoll et al., 2020a).

Following the discovery of Fourier encoding (Mansfield, 1977), Cartesian MRI sampling patterns would collect the entire Fourier transform of the object prior to reconstruction. The advent of parallel imaging (Sodickson and Manning, 1997; Pruessmann et al., 1999; Griswold et al., 2002) enabled to collect less than the full Fourier transform for reconstruction. Sampling patterns with parallel imaging are regular (*i.e.*, collecting every other line) due to the beneficial noise properties of regular undersampling that arise from the smoothness of the coil sensitivity patterns. The introduction of compressed sensing (Lustig et al., 2007) altered this paradigm, as random patterns are beneficial for compressed sensing. Traditionally in MRI, this divergence has been settled with patterns designed based on heuristics, such as Poisson-disc sampling (Vasanawala et al., 2011). Motivation to improve over heuristics has led to a wealth of research on sampling pattern design, which can be broadly classified based on 1) their adaptivity and 2) whether they jointly optimise the reconstructor.

Non-adaptive sampling / Non-joint training: Several works have proposed optimising sampling patterns for sparse reconstruction. Some works considered incoherence between Fourier operators and sparse bases (Chauffert et al., 2014; Adcock et al., 2017). Optimising sampling patterns while considering sensitivity coils is more difficult because the coils are patient-specific, but some works have considered this case with on-the-fly methods (Levine and Hargreaves, 2017).

Non-adaptive sampling / Joint training: In the CS-based literature, the work by Ravishanker and Bresler (2011) investigated the use of alternate optimisation to design a fixed sampling pattern for specific reconstruction methods. More recently, Bahadir et al. (2019) introduced the LOUPE algorithm, which co-jointly trains a non-adaptive probabilistic sampling pattern and deep learning-based reconstruction model. Other variants of this joint training idea have been proposed, exploring different ways to model the sampling pattern and differentiating through the sampling operation (Weiss et al., 2020; Huijben et al., 2020; Wang et al., 2021). Recent work has also studied the use of alternate/iterative optimisation techniques (Razumov et al., 2021) in the context of deep learning-based reconstructions, optimising the samples for downstream tasks rather than for reconstruction only. However, in general, the majority of these work has focused on the single-coil setup.

Adaptive sampling / Non-joint training: A recent line of work has focused on adaptive methods for a pre-trained reconstructor. One approach formalises the MRI acquisition problem as a partially observable Markov decision process (POMDP), and proposed the use of reinforcement learning (Sutton and Barto, 2018) to solve it. In particular, Bakker et al. (2020) considered the use of policy gradients (Bartlett and Baxter, 2001), while Pineda et al. (2020) proposed using the DDQN algorithm (Van Hasselt et al., 2016). An alternative approach considers combining Bayesian inversion with posterior estimation to select k -space

points [Sanchez et al. \(2020\)](#). The non-joint nature of these models requires a pre-trained reconstruction model, which is usually effectively trained *off-policy* with respect to the subsampling strategy (as it is learned only after reconstruction training concludes). This may be detrimental to overall system performance due to 1) the reconstruction model expending capacity on images resulting from subsampling patterns that are unlikely under the learned policies, and 2) the policies learning under the constraint that its proposed subsampling masks are within distribution for the reconstruction model. These difficulties suggest the need for a jointly optimised system.

Adaptive sampling / joint training: Following the previous discussion, another recent trend involves jointly training the sampler and the reconstructor. [Jin et al. \(2019\)](#) trained the sampler to mimic Monte Carlo tree search ([Silver et al., 2017](#)), while the reconstructor was simultaneously trained on the patterns generated by the sampler, but without any direct gradient path from one to the other. [Zhang et al. \(2019\)](#) proposed an uncertainty criterion for action selection, using a GAN-inspired training objective between reconstructor and the sampler. [Yin et al. \(2021\)](#); [Van Gorp et al. \(2021\)](#) are the closest to our work, in that they co-design both the sampler and the reconstructor in a fully-differentiable framework, with [Yin et al. \(2021\)](#) specifically focusing on the question of joint vs non-joint optimisation and finding the former advantageous. However, in contrast to our work, all these approaches are evaluated only in the single-coil MRI setting.

Appendix B. Future work

The optimal degree of adaptivity is unknown for both single- and multi-coil MRI subsampling policies, but especially the multi-coil setting remains under-explored. We hope that our results, especially the discussion in Section 4.4.1 of the main text, open up promising avenues for future research.

We suggest that one such avenue may be to explore the potential over-regularisation effect of subsampling adaptivity, by employing reconstruction models better equipped to distinguish inputs resulting from varying k -space masks. This may, for instance, be achieved by conditioning the reconstruction model on the subsampling mask explicitly, either by adding the subsampling mask as input to the refinement module, or possibly through the use of a hypernetwork ([Ha et al., 2016](#)) that predicts refinement module parameters conditional on \mathbf{M} .

Another potentially fruitful research direction is to explore to what degree the per-slice sensitivity maps of the E2E VarNet confer an intrinsic adaptivity to the joint model. Sensitivity maps represent the primary conceptual difference between the single- and multi-coil settings, and may present a unique challenge in designing adaptive strategies for multi-coil MRI reconstruction.

Finally, beyond adaptive policies, a promising area of research may be to explore *active* multi-coil policies, *i.e.*, strategies that dynamically adapt to acquired measurements during the scanning process. While in our experiments we apply a single Policy network to the ACS input, our utilisation of hard DC layers sequentially with refinement allows for placing the Policy in-between cascades of the E2E VarNet as well. We briefly experimented with this, and found no significant differences in performance applying the Policy to the reconstruction resulting from the first cascade instead. However, we did not deeply explore the possibility

of interspersing multiple Policies among the cascades of the E2E VarNet: by conditioning the Policy on outputs of previous Policies in such a way, the model may be able to learn active subsampling strategies

Limitations of our work: We performed our analysis on coil-compressed (128×128) k -space acquisitions, rather than the full data. The relevance of our work for the original, un-cropped data should be empirically verified by future research.

Appendix C. Implementation details

In this section we provide additional implementation details.

C.1. Cropping

To properly crop the fastMRI knee k -space, we need to account for its oversampling in the frequency encoding direction, as well as preserve the correct aspect-ratio. To this end, we first transform the full MR k -space to the image domain, and crop to the `recon_size` attribute stored in the `.h5` file metadata. The result is then transformed to k -space, where we finally crop to (128×128) . The target is constructed from this cropped k -space.

C.2. Data range

Computing SSIM and PSNR values requires specifying a data range (or dynamic range) that denote the maximum value a pixel can take. Our coil-compressed implementation employs the maximum of the slice as the data range during training, but validation numbers presented in this paper are computed using the maximum of the volume as data range, as is standard in the MRI literature.

C.3. Early stopping

The default fastMRI E2E VarNet implementation employs a form of early stopping, which we also follow: models are trained for the full 50 epochs mentioned in Section 4.3 of the main text, but only the model that performs best on validation SSIM is saved as a checkpoint and used to report final results.

C.4. Policy network architecture

The Policy network starts with an initial (1×1) convolution with 2 (real and imaginary) input channels and 16 output channels, followed by instance normalisation (Ulyanov et al., 2016) and ReLU (Nair and Hinton, 2010) activation. We further employ four convolutional blocks, each consisting of a zero-padded (3×3) convolution layer that doubles the number of channels, followed by an instance normalisation, ReLU activation, and (2×2) max-pooling layer. Note that this architecture is reminiscent of the first half of the default fastMRI U-Net (Zbontar et al., 2018). The resulting tensor is flattened and fed through two dense layers each composed of 256 output units, followed by a leaky-ReLU (Maas et al., 2013) activation with slope 0.01, and one final dense layer with number of output units equal to the number of measurements available in the full k -space M . The Policy has 4,685,712 total parameters.

C.5. Sparse and multiplicative DC

As mentioned in Section 3.2 of the main text, our implementation of a cascade differs slightly from the original E2E VarNet. Specifically, within a cascade we employ a hard DC layer sequentially with the refinement module, rather than a soft DC layer simultaneously with the refinement module. When combined with learned sampling strategies, this change opens up an interesting choice in the DC layer implementation. The original DC layer is implemented using a sparse operation, namely `torch.where(mask, observed_kspace - phantasised_kspace, 0)` on the difference between the observed and phantasised k -space. While this is no different from a non-sparse operation (*e.g.*, multiplying the difference by the mask values directly) when the observed k -space is not part of training, as soon as the observed k -space requires gradients (such as when training the Policy network or LOUPE) this sparse operation has the effect of removing gradients for unobserved k -space measurements.

We have explored the effect of this choice by repeating all Policy and LOUPE experiments for both types of DC layer. All results presented in this paper include a search over DC layer type, as well as all the other hyperparameters mentioned. However, while we include this choice in our hyperparameter search, we find empirically that the choice makes no significant difference to the performance of any of our models.

Appendix D. Additional results

In this section we present and discuss additional quantitative (Appendix D.1) and qualitative results (Appendix D.2).

D.1. Quantitative results

Here we present some additional quantitative results. In Appendix D.1.1 we report our main results table extended with the PSNR and NMSE metrics. Appendix D.1.2 contains additional mutual information results, as well as a discussion on the hypothesis that generalisation may be preventing the system from learning strong adaptive Policies: we conclude that this is unlikely. Finally, in Appendix D.1.3 we present results for systems trained with reconstruction models of increasing capacities, and further discuss the amortisation hypothesis: we conclude that future research is needed.

D.1.1. PSNR AND NMSE OF MAIN RESULTS

In Table 2 we report our main results extended with PSNR and NMSE values. SSIM is multiplied by 100, and NMSE by 1000 for readability. Note that SSIM ranking correlates well with PSNR and NMSE ranking. We also split results by non-linearity (sigmoid or softplus). While sigmoid-based strategies slightly outperform softplus-based strategies in all settings, the performance gap is much larger in specifically the $8\times$ Policy setting.

D.1.2. MUTUAL INFORMATION

In this section we present further results on Policy adaptivity through analysis of the mutual information of the learned subsampling strategies, following Bakker et al. (2020). In our

Table 2: Extension of primary results. Reported are validation SSIM, PSNR and NMSE of the best performing models in each category, averaged over three seeds. The sigmoid strategies consistently outperform the softplus strategies.

Policy						
Accel	SSIM	Sigmoid PSNR	NMSE	SSIM	Softplus PSNR	NMSE
4×	95.63 \pm 0.27	37.95 \pm 0.49	6.040 \pm 0.60	94.98 \pm 0.30	37.05 \pm 0.35	7.433 \pm 0.58
8×	93.26 \pm 0.20	35.17 \pm 0.12	11.21 \pm 0.28	90.40 \pm 0.55	32.66 \pm 0.49	20.22 \pm 2.11

LOUPE						
Accel	SSIM	Sigmoid PSNR	NMSE	SSIM	Softplus PSNR	NMSE
4×	95.61 \pm 0.55	37.73 \pm 0.92	6.387 \pm 1.24	94.94 \pm 0.14	37.02 \pm 0.13	7.457 \pm 0.23
8×	91.37 \pm 0.67	33.34 \pm 0.45	17.32 \pm 1.84	90.76 \pm 0.17	32.86 \pm 0.09	19.21 \pm 0.40

Equispaced			
Accel	SSIM	PSNR	NMSE
4×	95.38 \pm 0.03	37.80 \pm 0.06	6.307 \pm 0.11
8×	91.30 \pm 0.06	33.70 \pm 0.07	15.77 \pm 0.22




Figure 4: Validation SSIM as a function of the policy adaptivity measured as the mutual information (MI) for the 4× acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.




Figure 5: Validation SSIM as a function of the policy adaptivity measured as the mutual information (MI) for the 8× acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.

case, the MI is computed under the assumption that the policy outputs M independent Bernoulli distributions; one for each potential measurement site. Due to the rejection sampling step – which guarantees the correct number of measurements B is acquired – the actual sampled measurements are not fully independent; thus, the independence assumption is an approximation. Note however that, due to the data-processing inequality, our MI approximation provides an upper bound on the MI of the joint acquisition distribution.

Figure 4 is the 4× counterpart of the 8× Figure 2 from the main text, here re-presented as Figure 5. The trends for the Policy models are similar: the highest performing Policies employ the sigmoid non-linearity and are non-adaptive, although the effect is less exaggerated




Figure 6: Train SSIM as a function of the difference (train MI - validation MI), for the $4\times$ acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.




Figure 7: Train SSIM as a function of the difference (train MI - validation MI), for the $8\times$ acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.




Figure 8: The difference (train SSIM - validation SSIM) as a function of the difference (train MI - validation MI), for the $4\times$ acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.




Figure 9: The difference (train SSIM - validation SSIM) as a function of the difference (train MI - validation MI), for the $8\times$ acceleration factor. Each dot is a single model. Red: softplus; black: sigmoid.

than in the $8\times$ setting. For LOUPE we note a qualitative difference in relative performance of the sigmoid and softplus models: for $4\times$ sigmoid models are clearly favoured while in the $8\times$ setting the difference is smaller.

In Figures 6 and 7 we present similar plots, where now *train* SSIM is placed against the difference between train and validation MI (labeled as ‘Delta MI’). We observe that all the best performing models are no more adaptive on the train data than on the validation data: given that these models are non-adaptive on the validation data, we conclude that they are non-adaptive on the train data as well. This observation provides evidence against the hypothesis that the lack of high-performing adaptive Policies is due to generalisation issues. In Section 4.4 we noted that non-adaptivity requires the Policy to learn to ignore the data: indeed, if the Policy has learned to ignore validation data, it must have learned to do so on the training set: this additionally suggests that the generalisation hypothesis is unlikely.

In Figures 8 and 9 we have repeated these figures with now the difference between train SSIM and validation SSIM on the y-axis (labeled as ‘Delta SSIM’). This visualisation

Table 3: Effect of hyperparameters on validation SSIM for Policy and LOUPE, averaged over three seeds. 5c/18ch: reconstructor with 5 cascades and 18 refinement channels. The sigmoid strategies consistently outperform the softplus strategies.

Capacity	Accel	Policy		LOUPE	
		Sigmoid	Softplus	Sigmoid	Softplus
5c/18ch	4×	95.42 \pm 0.19	94.31 \pm 0.90	95.37 \pm 0.19	94.21 \pm 1.05
	8×	92.55 \pm 0.79	89.28 \pm 0.93	90.46 \pm 0.80	89.76 \pm 0.45
7c/18ch	4×	95.51 \pm 0.13	94.59 \pm 0.58	95.61 \pm 0.55	94.94 \pm 0.14
	8×	92.91 \pm 0.69	89.20 \pm 1.89	91.37 \pm 0.67	90.76 \pm 0.17
5c/36ch	4×	95.63 \pm 0.27	94.98 \pm 0.30	95.55 \pm 0.22	94.74 \pm 0.43
	8×	93.26 \pm 0.20	90.40 \pm 0.55	91.28 \pm 0.21	90.15 \pm 0.47

shows that the softplus-based policies are more prone to overfitting than the sigmoid-based policies, especially at 8× accelerations, suggesting that overfitting may play a role in the overall lower performance of softplus-based policies. However, note that this effect is not just mediated through adaptivity, as a number of policies that show no significant change in MI between train and validation data overfit quite strongly.

D.1.3. RECONSTRUCTION MODEL CAPACITY

In Table 3 we present results for systems trained with reconstruction models of increasing capacities, as mentioned in Section 4.4.1 of the main text. Although the higher capacity models lead to slight improvements in terms of Policy method SSIM, we note that sigmoid models still dominate softplus models, especially in the 8× Policy setting.

These higher capacity models were already included in the adaptivity visualisations of Figures 2 and 4. Note that the best performing Policy models all show no adaptivity. If amortisation is an issue as hypothesised, these results suggest that increased capacity alone may not overcome it. Indeed, as the refinement module does not directly receive information about the subsampling mask as input, the model may be unable to distinguish inputs resulting from various masks, preventing it from using its increased capacity effectively. Further research is needed to explore the degree to which amortisation over the dataset poses a challenge for learning strong adaptive subsampling strategies. We suggest some directions specifically aimed at exploring the amortisation hypothesis in Appendix B.

D.2. Qualitative results

In this section we provide visualisations of our model outputs. Specifically, we visualise some of the learned subsampling strategies in Appendix D.2.1, reconstructions in D.2.2, and sensitivity maps in D.2.3.

D.2.1. SUBSAMPLING VISUALISATION

Figures 10 and 11 show sampling strategies (learned) by the various methods for respectively 4× and 8× acceleration. Note that softplus-based Policies assign nonzero probability to every action, while sigmoid-based Policies are more selective. This effect is also observed for LOUPE at 8×, although it vanishes at 4× acceleration.




Figure 10: Sampling strategies for the $4\times$ case. Every line corresponds to a possible k -space measurement. The height of a line denotes the probability of sampling that measurement. ACS measurements are shown in red and always have probability 1.

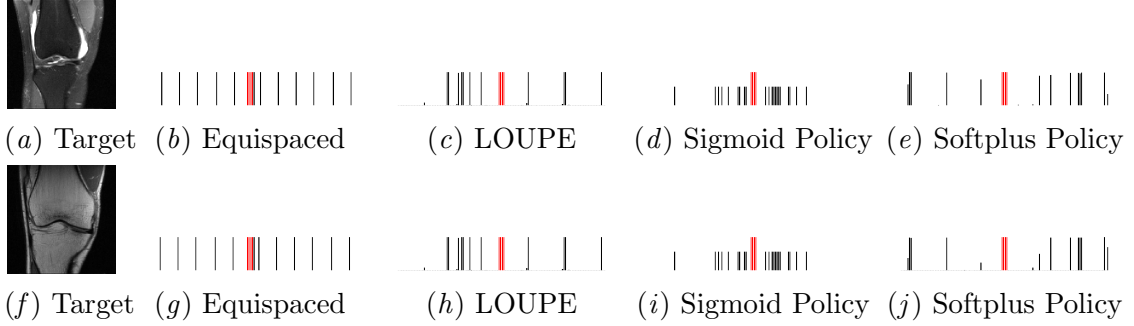


Figure 11: Sampling strategies for the $8\times$ case. Every line corresponds to a possible k -space measurement. The height of a line denotes the probability of sampling that measurement. ACS measurements are shown in red and always have probability 1.

D.2.2. RECONSTRUCTIONS

Figures 12 and 13 show selected reconstructions by the various methods for respectively $4\times$ and $8\times$ acceleration.

D.2.3. SENSITIVITY MAPS

Figures 14 and 15 show sensitivity maps learned by the various methods for respectively $4\times$ and $8\times$ acceleration. Further research is needed to explore the interaction between the sensitivity maps and the acquisition strategies.

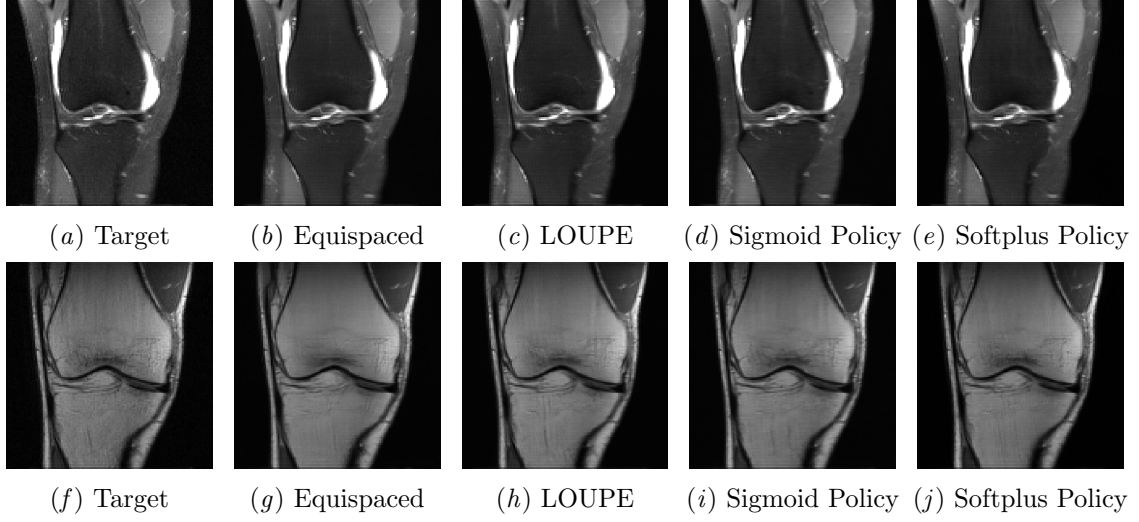


Figure 12: Examples of $4\times$ reconstructions.




Figure 13: Examples of $8\times$ reconstructions.




Figure 14: Examples of $4 \times$ sensitivity maps estimated by each model corresponding to the reconstructions displayed in Figure 12.




Figure 15: Examples of $8 \times$ sensitivity maps estimated by each model corresponding to the reconstructions displayed in Figure 13.

