Inter-Slice Acquisition Budget Allocation for Pathology Aware Accelerated MRI

Meva Himmetoglu[1][0000−0003−3228−7061] and Ender Konukoglu[1,2][0000−0002−2542−3611]

[1] Computer Vision Lab, ETH Zürich, Sternwartstrasse 7, Zürich, Switzerland
meva.himmetoglu@vision.ee.ethz.ch
[2] LOOP Zürich - Medical Research Center, Zürich, Switzerland

Abstract. Accelerated MRI acquisition reduces scan time but risks degrading image quality in clinically relevant regions. Existing adaptive sampling methods optimize k-space line selection within individual slices, ignoring the volumetric structure of clinical acquisitions where pathology spans only a subset of slices. We propose a volume-level Reinforcement Learning (RL) framework that learns to allocate acquisition budget *across slices* of a 3D MRI volume, prioritizing slices containing abnormalities. Our approach formulates the acquisition as a sequential decision process: a reinforcement learning agent observes per-slice lesion confidence scores and sampling fractions, then selects which slice should receive additional k-space lines at each step. We evaluate on the fastMRI knee dataset with various knee pathology annotations at 2×–8× acceleration, comparing against random uniform sampling and methods that optimize sampling based on reconstruction quality, detection rate, or disease severity. Our agent attains the highest lesion-region SSIM at the clinically common 2×–4× factors and remains competitive at 8×, improving over random uniform sampling by +0.041 to +0.078 SSIM without any access to lesion annotations at test time.

Keywords: active acquisition · reconstruction with pathologies · reinforcement learning.

1 Introduction

Magnetic Resonance Imaging (MRI) is a cornerstone of diagnosis in radiology, offering unparalleled soft-tissue contrast without ionizing radiation. However, MRI acquisitions are inherently slow due to sequential sampling of the spatial frequency domain. Accelerated MRI reduces scan time by acquiring fewer k-space measurements and relying on computational methods to reconstruct diagnostic quality images from undersampled data [6, 16]. The dominant paradigm in accelerated MRI applies a fixed undersampling pattern uniformly across all slices of a volume. While methods such as equidistant sampling with auto-calibration signal (ACS) lines provide robust acceleration, they treat every slice identically regardless of clinical relevance. In many diagnostic scenarios pathology is localized to a small subset of slices, so allocating equal budget to healthy and

pathological slices is a suboptimal use of limited acquisition time, and adaptive allocation can improve pathology visualization at the same acceleration rate.

Reconstruction driven adaptive sampling. Pineda et al. [7] pioneered the RL formulation for active MRI acquisition, modeling single-slice k-space column selection as a Partially Observable Markov Decision Process (POMDP) and solving with Double DQN (DDQN) [10]. Bakker et al. [2,3] framed the same problem as policy gradient experimental design. Beyond RL, Yin et al. [15] proposed a fully differentiable co-design that jointly trains a sequential sampler and reconstructor end-to-end, and Yang and Dong [12] jointly learned a sampler-reconstructor pair through a sparse reward POMDP. Yiasemis et al. [14] extended this co-optimization to dynamic multi-coil cardiac cine MRI, training an adaptive sampler that produces case-specific Cartesian patterns for a target acceleration factor together with a dynamic reconstruction network. Earlier, Bahadir et al. [1] learned a single population-level probabilistic mask jointly with an unrolled reconstruction network. All of these works rewarded global image fidelity rather than diagnostic content.

Diagnostic and task driven sampling. A second line of work optimizes sampling for the downstream task directly, often bypassing reconstruction. Du et al. [4] introduced patient level "image-less" active sampling, inferring disease directly from undersampled k-space with a greedy policy gradient agent rewarded by classification improvement, and extended it [5] to sequential diagnosis with a multi-objective step-weighted reward. Yen et al. [13] likewise learned an image-less PPO policy over k-space columns for pathology classification.

Closest to our diagnostic motivation, Xu and Oksuz [11] proposed a PPO (Proximal Policy Optimization) framework for column selection that bridges k-space and image domains. Their compound reward combines global and lesion-bounding-box SSIM with k-space fidelity in lesion regions, dynamically reweighted by the global and lesion SSIM gap. Despite their differing objectives (reconstruction, classification, or segmentation) and reward designs, all previous works adapted sampling at the single-slice level or, for dynamic acquisitions, across temporal frames, in both cases optimizing *which k-space columns* to acquire within each 2D frame. None allocate a shared sampling budget *across the slices of a 3D volume*. We identify a complementary and previously unexplored optimization axis: *inter-slice budget allocation*. Rather than sampling every slice equally, our method learns to concentrate budget on slices containing pathology, improving reconstruction quality precisely where diagnostic decisions are made. Our main contribution in this work is the formulation of volume-level MRI acquisition as a reinforcement learning problem where actions correspond to slice selection rather than k-space line / point selections within a slice, introducing inter-slice budget allocation as a new optimization dimension. We demonstrate improved reconstruction quality in lesion regions compared to random uniform sampling.

2 Method

We formulate volume-level MRI acquisition as a finite-horizon Markov Decision Process (MDP). At each time step an RL agent selects a slice to receive additional k-space lines, guided by the predicted pathology confidence and the sampling progress across the volume. Figure 1 illustrates the architecture.

2.1 Problem Formulation

Consider a 3D MRI volume with S slices, where each slice s has a 2D k-space $\mathbf{k}_s \in \mathbb{C}^{C\times H\times W}$ with C receiver coils. The total acquisition budget allows sampling B k-space lines distributed across slices, corresponding to a target acceleration factor R. We seek a policy π that allocates these B lines across slices to maximize reconstruction quality in regions potentially containing pathology. The agent repeatedly decides which slice receives the next block of k-space lines, placed at random within that slice. The state $\mathbf{s}_t \in \mathbb{R}^{S_{\max}\times 3}$ encodes three features per slice at time step t: $\mathbf{s}_t[i] = \left[\hat{p}_i^{(t)}, f_i^{(t)}, b_t\right], \quad i = 1, \ldots, S_{\max}$ where $\hat{p}_i^{(t)} \in [0,1]$ is the *raw predicted lesion probability* output by the classifier, $f_i^{(t)} = n_i^{(t)}/W$ is the fraction of k-space lines acquired for slice i out of a total of W phase-encoding lines, and b_t is the *global remaining budget fraction*, $b_t = B_t/(S \times W)$, where B_t is the number of k-space lines still to be allocated at step t. The total per-volume budget is set by the target acceleration factor, $B = \lfloor S \times W/R \rfloor$ (minus the fixed ACS lines).

The action $a_t \in \{0, 1, \ldots, S_{\max} - 1\}$ selects which slice receives additional k-space lines at step t. Upon selection, L lines are drawn uniformly at random from the unsampled positions of slice a_t's acquisition mask, where L is a fixed hyperparameter. Invalid actions (fully sampled slices) are masked before action selection, ensuring they are never chosen.

The reward measures reconstruction quality improvement within the lesion region of the selected slice: $r_t = \mathrm{SSIM}(\hat{\mathbf{x}}_{a_t}^{(t+1)}, \mathbf{x}_{a_t}^{\mathrm{gt}}; \mathcal{B}_{a_t}) - \mathrm{SSIM}(\hat{\mathbf{x}}_{a_t}^{(t)}, \mathbf{x}_{a_t}^{\mathrm{gt}}; \mathcal{B}_{a_t})$ where $\hat{\mathbf{x}}_{a_t}^{(t)}$ is the reconstruction of slice a_t at step t, $\mathbf{x}_{a_t}^{\mathrm{gt}}$ is the fully-sampled ground truth, and $\mathrm{SSIM}(\cdot, \cdot; \mathcal{B})$ computes SSIM within bounding box $\mathcal{B}_{a_t}$. For

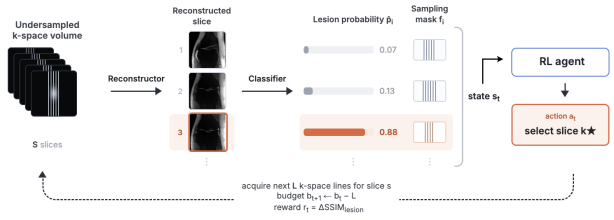

Fig. 1: Overview of our framework. At each step t, the RL agent observes per-slice features and selects which slice receives additional lines. The reward is the SSIM improvement within the lesion bounding box $\mathcal{B}$.

slices without lesion annotations, $r_t = 0$. This design directly aligns the optimization objective with clinical utility: the agent earns reward only by improving reconstruction quality where it matters for diagnosis.

2.2 Model Architecture

Reconstruction and Detection Model The detection backbone consists of a reconstructor and a pathology classifier module: an End-to-End Variational Network (E2E-VarNet) [9] with 12 cascades, followed by ResNet-18 classification head. The VarNet receives multi-coil k-space and a sampling mask, and outputs an image. This image is passed to the ResNet-18 classifier, which produces a binary probability $\hat{p}_i \in [0, 1]$ indicating whether slice i contains pathology of any of the targeted labels of knee pathologies. We normalize the reconstructed images before classification with the percentile statistics obtained from ACS only reconstructions to ensure consistent features across different accelerations.
Acquisition Agent Our agent uses a permutation-equivariant per-slice scorer network to encode the inductive bias that slices with identical features should receive identical scores regardless of their position in the volume. We apply a shared-weight MLP ϕ_θ independently to each slice's 3-dimensional feature vector, $\mathbf{h}_i = \phi_\theta(\mathbf{s}_t[i]) \in \mathbb{R}^{d_h}$, where ϕ_θ has two hidden layers of dimension $d_h = 256$ with ReLU activations. The score for action $a = i$ is then a learned linear projection of $\mathbf{h}_i$: $Q(s, i) = \mathbf{w}^\top \mathbf{h}_i + b$. We instantiate this scorer in an actor-critic form for policy gradient training. The critic head mean-pools the per-slice representations across all slices to obtain a permutation-invariant global state summary, then projects to a scalar value estimate.

2.3 Training Details

We trained E2E-VarNet at random acceleration factors $R \sim \text{Uniform}(1, 32)$ per sample to ensure the classifier remains informative across the range of sampling fractions the RL agent will encounter. The classifier head was trained with frozen VarNet in a second stage. This training used the AdamW optimizer ($\text{lr} = 10^{-3}$, weight decay $= 0$, step decay $\times 0.1$ at epoch 40) with dropout $p = 0.3$.

We trained the acquisition agent with PPO [8]. Advantages were estimated via Generalized Advantage Estimation (GAE): $\hat{A}_t = \sum_{k=0}^{T-t-1} (\gamma\lambda)^k \delta_{t+k}, \quad \delta_t = r_t + \gamma V(\mathbf{s}_{t+1}) - V(\mathbf{s}_t)$ with $\lambda = 0.95$. The clipped surrogate objective was: $\mathcal{L}_{\text{PPO}} = -\mathbb{E}\Big[\min\Big(\rho_t \hat{A}_t,\ \text{clip}(\rho_t, 1-\epsilon, 1+\epsilon)\hat{A}_t\Big)\Big]$ with $\rho_t = \pi_\theta(a_t|\mathbf{s}_t)/\pi_{\theta_{\text{old}}}(a_t|\mathbf{s}_t)$ and clipping parameter $\epsilon = 0.2$. The full loss contains a value function term and an entropy bonus: $\mathcal{L} = \mathcal{L}_{\text{PPO}} + 0.5 \cdot \mathcal{L}_{\text{value}} - 0.01 \cdot \mathcal{H}[\pi_\theta]$. The agent was trained for 5000 episodes, each corresponding to one randomly sampled training volume until full k-space exhaustion from 4% ACS initialization. PPO hyperparameters were: $\gamma = 0.99$, GAE $\lambda = 0.95$, $\epsilon = 0.2$, 4 PPO epochs per update, 4 episodes per update, $\text{lr} = 10^{-4}$. We swept $L \in \{16, 32, 64, 256\}$ and report each acceleration at its best setting ($L = 32, 16, 16, 64$ at $2\times$, $3\times$, $4\times$, $8\times$); performance is flat across $L \in \{16, 32, 64\}$, varying by at most 0.017 SSIM at any acceleration.

3 Experimental Setup

3.1 Dataset

We used the fastMRI multi-coil knee dataset [16] with fastMRI+ [17] annotations. We used 974 volumes for training including a held-out validation set. All results are reported on a disjoint test set of 147 lesion-bearing volumes, together with 51 held-out volumes carrying no annotation that we use for the no-lesion analysis; the comparison against Du et al. in Table 2 is restricted to the 37 test volumes that were also held out by their model.

The classifier is trained on 11 fastMRI+ knee pathology labels, spanning three tissue groups: meniscus tear, cartilage thickness losses, and ligament sprains. A slice is labeled positive if any of these pathologies is annotated within it.

3.2 Baselines and Compared Methods

We compared our volume-level RL approach against multiple baselines:
Equidistant Sampling. All slices receive identical sampling patterns consisting of ACS center lines plus equidistantly spaced outer lines, achieving the target acceleration factor R uniformly across the volume.
Oracle Lesion Greedy. Upper bound that uses ground truth annotations to always allocate budget to the least-sampled lesion slice. This represents the best achievable performance with perfect pathology localization.
Classifier Greedy. A prediction-driven greedy policy that, at each step, allocates the next block of L lines to the valid slice with the highest classifier lesion confidence $\hat{p}_i$, breaking ties at random. Unlike the oracle it relies solely on the detection head and has no access to ground-truth annotations.
Informed Greedy. A budget-aware variant of the classifier-greedy policy that ranks valid slices by the ratio $\hat{p}_i/f_i$, where f_i is the current sampling fraction of slice i, and selects the highest-scoring slice at each step (ties broken at random).
Pineda et al. [7] Designed for single-slice k-space line selection, applied independently to each slice with an equal per-slice budget; this represents intra-slice optimization operating within a uniform inter-slice allocation.
Xu et al. [11] Their PPO-based column selection method, trained on our annotated fastMRI knee data. Each slice is processed independently with a uniform per-slice budget: at 4× on typically 320-column k-space the agent selects 68 columns per slice (80 total minus 12 ACS lines), one per step.

4 Results

Table 1 reports lesion-region SSIM (structural similarity index measure) and PSNR (peak signal to noise ratio) across acceleration factors. Our agent obtained the highest lesion-region SSIM at 2×, 3× and 4× — the range commonly used in clinical knee MRI protocols — including over the intra-slice methods that optimize columns within each slice under a uniform inter-slice budget (Pineda

Table 1: Comparison of averages of lesion region metrics ($n = 147$), at different acceleration factors. For reference, the fully sampled acquisition ($1\times$) attains SSIM 0.997 ± 0.002 and PSNR 45.94 ± 2.70 dB. Best per column in **bold**. †: the best (bold) method is significantly higher than every other sampler in that column (paired Wilcoxon signed-rank test, $p < 0.001$). Our agent is reported at the best action granularity L per acceleration.

	2×		**3×**		**4×**		**8×**	
Method	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR
Oracle	0.997 ± 0.002	45.92 ± 2.69	0.995 ± 0.007	45.01 ± 3.39	0.985 ± 0.020	42.62 ± 5.11	0.939 ± 0.044	36.42 ± 5.92
Equidistant	0.952 ± 0.024	35.50 ± 3.05	0.923 ± 0.040	33.40 ± 3.27	0.909 ± 0.042	32.57 ± 3.51	0.811 ± 0.067	29.33 ± 3.37
Random	0.941 ± 0.027	34.54 ± 3.13	0.926 ± 0.033	33.51 ± 3.27	0.877 ± 0.046	31.31 ± 3.10	0.790 ± 0.076	28.79 ± 3.27
Informed greedy	0.962 ± 0.020	36.76 ± 3.44	0.946 ± 0.027	35.06 ± 3.59	0.908 ± 0.041	32.74 ± 3.30	0.852 ± 0.054	30.44 ± 3.15
Classifier greedy	0.979 ± 0.033	44.47 ± 3.42	0.944 ± 0.073	$\mathbf{41.82} \pm 4.66$	0.885 ± 0.097	$\mathbf{38.47}^{\dagger} \pm 5.30$	0.759 ± 0.121	31.37 ± 4.71
Pineda [7]	0.947 ± 0.026	35.23 ± 2.85	0.919 ± 0.037	33.33 ± 2.98	0.898 ± 0.044	32.20 ± 3.15	0.832 ± 0.067	29.91 ± 3.17
Xu [11]	0.956 ± 0.024	35.87 ± 3.48	0.898 ± 0.041	32.03 ± 3.18	0.923 ± 0.037	32.92 ± 4.03	$\mathbf{0.874} \pm 0.057$	31.09 ± 3.69
Ours	$\mathbf{0.987}^{\dagger} \pm 0.020$	$\mathbf{44.73} \pm 3.34$	$\mathbf{0.967}^{\dagger} \pm 0.036$	40.98 ± 4.67	$\mathbf{0.940}^{\dagger} \pm 0.048$	37.28 ± 4.16	0.868 ± 0.067	$\mathbf{31.67} \pm 3.32$

and Xu). Notably, these methods optimize an axis orthogonal to ours: within each selected slice our agent places its lines at random, whereas they select columns with a learned policy. The comparison therefore holds their intra-slice advantage fixed and isolates the value of inter-slice allocation. These gains bring lesion-region quality close to the ground-truth lesion-greedy oracle, despite using only predicted lesion confidence and no access to annotations at test time. The lesion-aware reward design directly optimizes the metric that matters for clinical diagnosis. By concentrating acquisition budget on these regions, our method improved exactly the information that drives diagnostic decisions, while accepting slightly lower quality on healthy tissue slices that may be less diagnostically relevant. Measured over whole images, our policy loses -0.059 SSIM at $4\times$ against random uniform sampling. This is intrinsic to lesion-directed sampling rather than a defect of the learned policy: the ground-truth oracle pays a comparable -0.052, and classifier greedy pays -0.082 for a smaller lesion-region gain. Ours therefore buys the largest lesion-region improvement per unit of whole-volume fidelity given up; whether that exchange is worthwhile is a protocol-level decision, favourable only when the diagnostic question is localized. While we evaluated the method on knee pathologies, the framework is agnostic to the specific pathology. The classifier could be retrained for any condition with slice-level labels and the RL agent requires no modification, as it operates on the abstract state of classifier confidence scores and sampling fractions. Since the state is driven by that classifier, we note that its discrimination is moderate but stable across sampling rates (AUC 0.840 at full sampling to 0.820 at $16\times$), which is what allows a single policy to span the whole acceleration range. Placing the L lines uniformly at random within the chosen slice is to isolate the inter-slice effect from any intra-slice advantage, and not a proposed acquisition pattern; a deployed system would combine our slice-level decisions with a scanner-compatible in-slice ordering.

A distinct line of work instead samples directly for a downstream diagnostic objective: Du et al. [5] learn a weighted sequential policy over k-space columns

Table 2: Lesion region reconstruction quality of our inter-slice policy versus the diagnostically-driven sampler of Du et al. [5], on the 37 volumes (308 lesion slices) held out for both methods.

	4×		**8×**	
Method	SSIM	PSNR	SSIM	PSNR
E2E-VarNet reconstruction (mean ± std)				
Du et al. [5]	0.915 ± 0.029	33.30 ± 3.67	$\mathbf{0.874 \pm 0.042}$	31.51 ± 3.77
Ours	$\mathbf{0.950 \pm 0.049}$	$\mathbf{38.68 \pm 4.65}$	0.851 ± 0.070	$\mathbf{32.07 \pm 4.00}$
Zero-filled IFFT reconstruction (median ± MAD)				
Du et al. [5]	0.506 ± 0.069	20.18 ± 3.43	0.449 ± 0.079	19.52 ± 3.39
Ours	$\mathbf{0.957 \pm 0.043}$	$\mathbf{35.45 \pm 8.53}$	$\mathbf{0.747 \pm 0.102}$	$\mathbf{26.87 \pm 3.35}$

that maximizes slice-level lesion classification without ever forming an image. To position our approach against this paradigm we feed both policies' masks through a shared reconstructor, a learned multi-coil E2E-VarNet and a parameter-free zero-filled IFFT, and measure lesion-region SSIM/PSNR on the 37 volumes held out for both models (Table 2). Under VarNet our allocation leads at 4× and is comparable at 8×; under zero-filled reconstruction it is ahead at both, reflecting that under the same budget our agent concentrates far more raw k-space on the lesion region. This is a proxy we impose: Du's method optimizes classification directly, and we make no claim of superior *global* reconstruction.

Figure 2 presents visual results on two slices. Within the annotated lesion region our reconstruction is visibly sharper than every baseline, and the lesion-region SSIM confirms the ordering seen in Table 1. Figure 3 shows *which* slices each policy chooses over the course of acquisition. Our agent repeatedly returned to the lesion-containing band from the earliest steps, which clarifies the source of the volume-level gains in Table 1: the advantage stems from concentrating the budget on the diagnostically relevant slices and from doing so early, which is precisely what matters when acquisition is halted under a tight time budget.

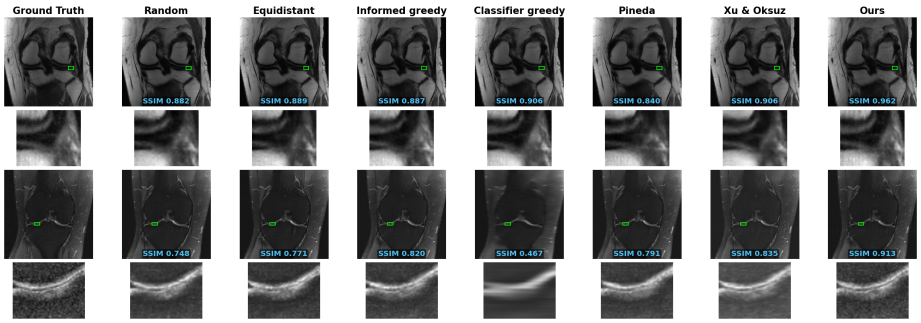

Fig. 2: Qualitative comparison on a slice with meniscus tear and one with thickness loss at 4× acceleration.

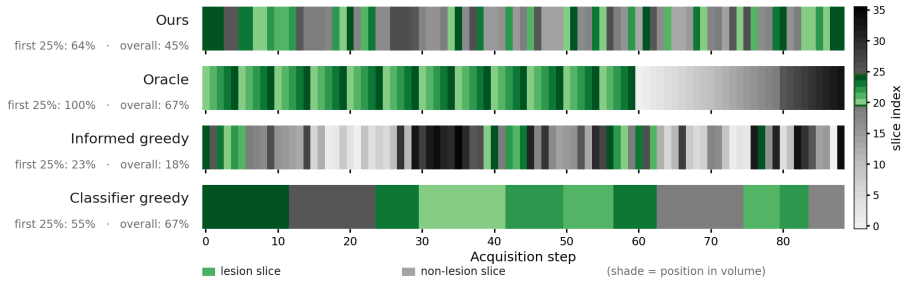

Fig. 3: Slice-selection timeline for a representative volume at 4× acceleration (36 slices, 5 with lesions). Each policy is shown as a horizontal strip over acquisition steps; every step is coloured with lightness encoding position in the volume.

Limitations. The gap between our best agent and the oracle indicates room for improvement. A natural extension is to combine inter-slice allocation with learned *intra-slice* column placement in a hierarchical agent that decides both which slice to sample and which k-space columns to acquire within it. A more fundamental limitation is that the agent is only as reliable as the detector that guides it: the state uses classifier scores, so the policy inherits the strengths and the blind spots. On volumes with no true pathology, rather than receiving a flat signal, the agent is misled by the learned prior, which tends to cluster in the central slices of the volume. As a result it continues to concentrate budget on this central band even when no lesion is present, which can hinder overall reconstruction quality relative to a uniform allocation that would have been preferable. Improving the calibration and robustness of the detection head and propagating its uncertainty or a global metric into the reward, are therefore likely to translate directly into better acquisition policies, as we see as our next steps. Finally, our evaluation uses a single anatomy (knee); broader validation across anatomies and pathologies is needed before clinical deployment.

5 Conclusion

We presented a reinforcement learning framework for adaptive MRI acquisition that introduces a previously unexplored optimization axis: *inter-slice budget allocation*, deciding which slice of a volume receives additional k-space lines rather than which columns to acquire within a single slice. On the fastMRI knee dataset, the agent achieves strong gains at clinically relevant acceleration factors. By concentrating acquisition budget on pathological slices, the method improves the regions that could drive diagnostic decisions. Our framework is modular: the reconstruction model, classifier, and RL agent can be independently improved. The most promising direction is combining inter-slice allocation with learned intra-slice column placement in a single hierarchical agent.

Acknowledgments. This project is funded by the Swiss National Science Foundation (SNSF) (205320_200877).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bahadir, C.D., Dalca, A.V., Sabuncu, M.R.: Learning-based optimization of the under-sampling pattern in mri. In: international conference on information processing in medical imaging. pp. 780–792. Springer (2019)
2. Bakker, T., van Hoof, H., Welling, M.: Experimental design for mri by greedy policy search. Advances in Neural Information Processing Systems **33**, 18954–18966 (2020)
3. Bakker, T., Muckley, M., Romero-Soriano, A., Drozdzal, M., Pineda, L.: On learning adaptive acquisition policies for undersampled multi-coil mri reconstruction. In: International Conference on Medical Imaging with Deep Learning. pp. 63–85. PMLR (2022)
4. Du, Y., Dharmakumar, R., Tsaftaris, S.A.: The mri scanner as a diagnostic: imageless active sampling. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 467–476. Springer (2024)
5. Du, Y., Liu, J., Dharmakumar, R., Tsaftaris, S.A.: Active sampling for mri-based sequential decision making. arXiv preprint arXiv:2505.04586 (2025)
6. Lustig, M., Donoho, D., Pauly, J.M.: Sparse mri: The application of compressed sensing for rapid mr imaging. Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine **58**(6), 1182–1195 (2007)
7. Pineda, L., Basu, S., Romero, A., Calandra, R., Drozdzal, M.: Active mr k-space sampling with reinforcement learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 23–33. Springer (2020)
8. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
9. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-end variational networks for accelerated mri reconstruction. In: International conference on medical image computing and computer-assisted intervention. pp. 64–73. Springer (2020)
10. Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: Proceedings of the AAAI conference on artificial intelligence. vol. 30 (2016)
11. Xu, R., Oksuz, I.: A reinforcement learning approach for optimized mri sampling with region-specific fidelity. Neurocomputing **638**, 130116 (2025)
12. Yang, P., Dong, B.: L2sr: learning to sample and reconstruct for accelerated mri via reinforcement learning. Inverse Problems **40**(5), 055015 (2024)
13. Yen, C.Y., Singhal, R., Sharma, U., Ranganath, R., Chopra, S., Pinto, L.: Adaptive sampling of k-space in magnetic resonance for rapid pathology prediction. In: Proceedings of the 41st International Conference on Machine Learning. ICML'24, JMLR.org (2024)

14. Yiasemis, G., Sonke, J.J., Teuwen, J.: End-to-end co-optimization of adaptive k-space sampling and reconstruction for dynamic mri. In: Medical Imaging with Deep Learning (MIDL). Proceedings of Machine Learning Research, vol. 315, pp. 1285–1324. PMLR (2026)
15. Yin, T., Wu, Z., Sun, H., Dalca, A.V., Yue, Y., Bouman, K.L.: End-to-end sequential sampling and reconstruction for mri. In: Roy, S., Pfohl, S., Rocheteau, E., Tadesse, G.A., Oala, L., Falck, F., Zhou, Y., Shen, L., Zamzmi, G., Mugambi, P., Zirikly, A., McDermott, M.B.A., Alsentzer, E. (eds.) Proceedings of Machine Learning for Health. Proceedings of Machine Learning Research, vol. 158, pp. 261–281. PMLR (04 Dec 2021)
16. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., et al.: fastmri: An open dataset and benchmarks for accelerated mri. arXiv preprint arXiv:1811.08839 (2018)
17. Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y.W., Hansen, M.S., Lungren, M.P.: fastmri+, clinical pathology annotations for knee and brain fully sampled magnetic resonance imaging data. Scientific Data **9**(1), 152 (2022)