

The MRI Scanner as a Diagnostic: Image-less Active Sampling

Yuning Du¹, Rohan Dharmakumar², and Sotirios A.Tsaftaris^{1,3}

¹ School of Engineering, The University of Edinburgh, Edinburgh, EH9 3FG, UK

² Krannert Cardiovascular Research Center, Indiana University School of Medicine,
Indianapolis, Indiana, USA

³ The Alan Turing Institute, London, NW1 2DB, UK

Abstract. Despite the high diagnostic accuracy of Magnetic Resonance Imaging (MRI), using MRI as a Point-of-Care (POC) disease identification tool poses significant accessibility challenges due to the use of high magnetic field strength and lengthy acquisition times. We ask a simple question: *Can we dynamically optimise acquired samples, at the patient level, according to an (automated) downstream decision task, while discounting image reconstruction?* We propose an ML-based framework that learns an active sampling strategy, via reinforcement learning, at a *patient-level* to directly infer disease from undersampled k -space. We validate our approach by inferring Meniscus Tear in undersampled knee MRI data, where we achieve diagnostic performance comparable with ML-based diagnosis, using fully sampled k -space data. We analyse task-specific sampling policies, showcasing the adaptability of our active sampling approach. The introduced frugal sampling strategies have the potential to reduce high field strength requirements that in turn strengthen the viability of MRI-based POC disease identification and associated preliminary screening tools.

Keywords: Point-of-Care Diagnosis · Active Sampling Strategy · Reinforcement Learning · Magnetic Resonance Imaging.

1 Introduction

Despite the proliferation of Magnetic Resonance Imaging (MRI), its role as a Point-Of-Care (POC) diagnostic tool is muted due to poor accessibility caused by long acquisition time and cumbersome equipment [4]. Consequently, advancements have been made in assisting the diagnostic process with machine learning methods to sample less k -space data and reduce acquisition time in turn [8].

We first illustrate how learning-based strategies can help to reduce acquisition time by considering the depiction of conventional MRI-based diagnostic processes. In Figure 1a, an MRI scanner samples the full k -space, resulting in high-fidelity images, which are further interpreted by professional radiologists to identify biomarkers and provide a diagnosis. Alternatively, machine-learning (ML) based diagnostic processes can enable acquisition time savings by reconstructing a high-fidelity image from undersampled k -space (Figure 1b). Previous work




Fig. 1. Conventional vs. ML-based Diagnostic Processes.

has focused on optimising both reconstruction [3] and k -space sampling strategies [1]. To find the best k -space sampling pattern, various methods [12, 17, 18] optimise the mask given pre-defined sample rates, resulting in *population-level* masks, while reinforcement-learning-based methods [2, 11] can optimise an active sampling strategy at population or patient-level. A re-interpretation of such conventional use of MRI can lead to considerable savings, which will largely enable a low-field MRI future opening the road to POC and bedside imaging.

Considering that image reconstruction is a shared stage in both conventional (Figure 1a) and ML-based (Figure 1b) diagnostic processes, one ponders *Is the image necessary for automated, diagnostic, inferences?* One option is to perform direct inference on undersampled k -space as shown in Figure 1c. In fact, [14] have shown that it is possible to obtain biomarkers directly from k -space data. However, one then immediately ponders *Can AI agents learn effective k -space sampling strategies by considering expected diagnostic performance?*

A recent realisation of this concept is presented in [15]. The work formulates the described process as a classification task with an optimised sampling strategy. However, the sampling strategy is designed at the *population-level*, resulting in the same mask for all patients. All approaches that find optimal patient-level active sampling strategies use image reconstruction [2, 11, 20]. To the best of our knowledge, active sampling strategies for patient-level disease inference, from undersampled k -space, are lacking. Motivated by this, our **contributions** include:

- We propose, the first, patient-level active sampling using a reinforcement learning framework, aiming for direct disease inference from k -space.
- We test feasibility to infer Meniscus Tear presence in undersampled MRI..
- We investigate how different policies make decisions and how different starting points alter behaviour.

2 MRI Undersampling Preliminaries

Instead of directly imaging the human anatomy, MRI captures the electromagnetic activity in the body after exposure to magnetic fields and radiofrequency




Fig. 2. Our framework for active MRI sampling for point-of-care diagnosis.

pulses, which can be measured in k -space (i.e., the frequency domain). Considering the single coil measurement, the k -space data can be represented as a 2-dimensional complex-valued matrix $x \in \mathbb{C}^{r \times c}$, where r is the number of rows and c is the number of columns. The spatial image I can be obtained by applying the inverse Fourier Transform to x , denoted as $I = \mathcal{F}^{-1}(x)$. The undersampled k -space can be therefore represented as $x_s = U_L \circ x$, where U_L can be viewed as a binary mask $U \in \{0, 1\}^{r \times c}$ with L measurements from k -space. In our work, we exclusively consider the Cartesian mask for MRI undersampling. Consequently, the undersampled image is denoted as $I_s = \mathcal{F}^{-1}(x_s)$.

3 Methods

Overview: Our framework, illustrated in Figure 2, aims to reduce acquisition time by selectively sampling the k -space according to diagnostic significance, in a progressive fashion. The process begins with a small subset of randomly sampled k -space with L measurements denoted as $x_{s_0} = U_L \circ x$, resulting in a low-quality image I_{s_0} . I_{s_0} is then input into a pre-trained classification network f to generate the initial prediction y_0 . The extracted feature maps m_0 are processed by the *Feature Operator* O , producing high-level features $h_0 = O(m_0)$, which are then fed into the *Active Sampler* S . The policy network generates a sampling policy $p_\phi(h_0)$ parameterised by ϕ , guiding the selective sampling of diagnostically significant lines from the k -space. These sampled lines are subsequently added to U_L . The updated undersampled k -space data at step $t \in [0, T]$ are denoted as $x_{s_t} = U_t \circ x$, and are then fed back into the classification network with inverse Fourier transform. The iterative process continues until a sampling budget T is exhausted or user-defined reliability criteria are satisfied. During the training stage, the predictions y_t , accompanied by the ground truth diagnostic label g ,

are used as the criterion $r_t = r(y_t, g)$ to supervise the S . While in inference stage, there is no ground truth diagnostic label g provided.

Classification Network and Image Operator: To improve stability, the classification network is pre-trained with undersampled k -space data, ensuring an advantageous reward for the active sampler during training. This network also functions as the feature extractor in our setting. In the classification network, the earlier layers tend to learn low-level features such as edges, textures, and simple patterns, while deeper layers learn more complex and high-level features that are useful for discriminating between different classes. Therefore, utilising the feature maps from the model, the high-level features are selected and processed by the feature operator as further inputs.

Active k -space Sampler with Greedy Policy: Inspired by [2, 11], the sequential selection of k -space can be formalised as a Partially Observable Markov Decision Process (POMDP) [16]. Greedy Policy is used to maximise the expected return $J(\phi)$ of a policy p_ϕ parameterised by ϕ in such a POMDP. At each step, the classification improvement can be calculated using $R(r_t, r_{t+1}) = r(y_{t+1}, g) - r(y_t, g)$, where the criterion r is the cross-entropy.

During inference, the agent will be sampling one-line at a time. However, this will slow during training. Hence, the policy network is trained by sampling several lines in parallel the rewards of which are averaged [6]. Formally, we sample q lines at every time step, for a reward $R_{i,t}$ as the reward obtained from sample i at time step t , to obtain the following estimator:

$$\nabla_\phi J(\phi) \approx \frac{1}{q-1} \mathbb{E}_x \sum_{i=1}^q \sum_{t=L}^{T-1} \left[\nabla_\phi \log p_\phi(h_t) \left(R_{i,t} - \frac{1}{q} \sum_{j=1}^q R_{j,t} \right) \right]. \quad (1)$$

Evaluation Metrics: To assess the network’s classification performance, we employ metrics such as Recall, Area Under Curve (AUC), and Specificity.

4 Experiments

4.1 Dataset and Pre-processing

Dataset: We used single-coil k -space data and slice-level labels from the publicly available fastMRI dataset [19] and fastMRI+ dataset [21]. Randomly selecting 1100 annotated volumes (39125 slices) from the fastMRI Knee dataset. Our diagnostic task is to identification Meniscus Tear (MT) in each slice. Thereby, there are 32035 train slices (11.8% with MT), 5249 validation slices (9.3% with MT), and 1841 test slices (11.1% with MT).

Data Pre-processing: Since the k -space data have various sizes, we first use inverse Fourier transform to the fully sampled k -space data to get the ground truth image and crop it to size (320×320) for computation convenience. Thereby, the fully sampled k -space data of uniform size can be obtained by applying Fourier transform to the ground truth image. Notably, there is a severe class imbalance regarding the MT identification task. During training the classification network, we oversample the data to avoid overfitting on the majority class and poor generalisation on the minority class. During training the policy network, we undersample the data for computational expedience as [2].

Table 1. Diagnosis Support and Data Access of Implemented Methods.

Method	Diagnostic Support		Data Used by the Model		
	Sampling Optimization	Patient-level Strategy	Full k -space	Diagnostic Label	Undersampled k -space
<i>Oracle</i>			✓	✓	
<i>Undersampled</i>				✓	✓
<i>Policy Reconstruction</i> [2]	✓	✓	✓	✓	✓
<i>Policy Classifier</i> (Ours)	✓	✓		✓	✓

4.2 Implementation Details

We devise inference benchmarks to allow us to evaluate our approach in a fair fashion. Two benchmarks have the access to fully sampled k -space (high-fidelity images), and two do not optimise at the patient level the sampling strategy. The diagnosis support and data access are shown in Table 1.

Fully Sampled (*Oracle*): This serves as an benchmark estimator of classifier performance on image input obtained by fully sampled k -space data, and hence no sampling optimisation occurs. We trained the classifier with the ground truth image as input which is transformed from fully sampled k -space data and supervised by the diagnostic label using cross-entropy loss. We use as classification backbone a ResNet-50 [5] and to address the class imbalance in the training set, we add extra dropout layers overfitting, resulting in a total of 25.6M parameters.

Undersampled: This classifier serves as a baseline of performance when simple inverse Fourier is used to transform the under-sampled k -space data, without any sampling optimisation. It has the same backbone as the *Oracle*, and is trained with undersampled images with various sample rates (4 to 20) and center fraction (0 to 0.10), supervised as before.

Policy (via) Reconstruction [2]: We compare with a model that optimises patient-level sampling strategy with image reconstruction error as rewards. Notably this method has access to fully sampled k -space, and hence has access to more information during training. We pre-trained a reconstruction network using a U-Net [13] as the backbone with a first feature map size of 16 and 4 pooling cascades, resulting in a total of 837K parameters. The reconstruction network is trained with various sample rates (4 to 20) and center fraction (0 to 0.10) supervised using the ℓ_1 loss. We train the active sampler for reconstruction with reconstructed images from pre-trained model as inputs, and Structural Similarity Index Measure as the criterion to provide rewards, resulting in a total of 26.7M parameters. The reconstructed images obtained with the policy network and pre-trained reconstruction model are evaluated with the *Oracle*, referred to as the Policy-based Reconstruction Network (*Policy Reconstruction*).

Proposed Policy Classifier: For our method, a backbone similar to the *Undersampled* is used as pre-trained classification network. The reward is driven by the predictions of the network and its feature maps are used to train the policy. The Feature Operator uses the last 2 layers’ feature maps as input and the global average pooling function to achieve a (80, 512) output to feed in the policy network. The policy network (11M parameters) uses the cross entropy from classification network to provide rewards. It is trained with an initial sam-




Fig. 3. Classification performance varying sample rates (horizontal axis). More metrics are reported in the Supplementary.

ple rate of 5% with multiple center fraction and samples 64 lines to reach the sample rate of 25% as the sampling budget. The parallel acquisition $q = 8$.

For all methods, we employ the Adam optimiser with a learning rate of 10^{-4} and a step-based scheduler with a decay gamma of 0.1 for all model training. All classification and reconstruction models are trained for 30 epochs, and the policy networks of the active sampler are trained for 20 epochs. Our experimental setup uses the PyTorch framework, and all computations are conducted on NVIDIA A100 Tensor Core GPUs. Our code⁴ is available.

4.3 Results

Making Diagnostic Decisions with Undersampled Data

Figure 3 compares the performance at various sample rates for three strategies that use undersampled masks, namely the *Undersampled*, *Policy Reconstruction* and *Policy Classifier*. The first randomly samples k -space lines progressively; the other two start with a randomly initialised mask with a sample rate of 5% and optimally decide using their respective rewards and policies (see Section 3).

Our method, the *Policy Classifier*, consistently outperforms the *Undersampled* baseline across all metrics, showing that optimal sampling via the policy network helps in identifying diagnostically relevant k -space lines. This leads to considerable savings in data sampled (and consequently scan time). Our model approximates well the performance of the *Oracle*, which has been trained on high-fidelity data, and reaches optimal performance quickly. Taking AUC as an example, our approach reaches an AUC of 0.780 with 7% of the sampled k -space.

Our method closely approximates the performance of the *Policy Reconstruction*, which we highlight has been trained with access to fully sampled k -space and its label while our policy uses noisy reward from classifier as supervision.⁵

k -space Sampling Behaviour of the Policies

The results of the previous paragraph are reported over a coarse percentage of

⁴ https://github.com/vios-s/MRI_Active_Sampling

⁵ We observe that the *Policy Reconstruction* performs better than the *Oracle*. This policy reconstructs an image which is given to the *Oracle* for classification. Some denoising and smoothing are happening at the reconstruction which in turn acts as a regulariser for the *Oracle* classifier explaining this slightly improved performance.




Fig. 4. k -space behaviour of the two policies. The horizontal axis indicates the cumulative lines acquired while sampling. The initial 16 lines are randomly sampled.




Fig. 5. k -space preference as a function of sample rate (vertical axis downward) and for 3 center fraction scenarios. Colorbar indicates the possibility of being sampled.

k -space lines acquired. It is worth looking into how the two different policies behave when asking each policy to progressively make decisions on which line to acquire in a line-by-line manner. The results of this exercise are shown in Figure 4. The behaviour of the two methods is evidently different. Taking Recall and AUC as examples, the *Policy Classifier* plateaus quickly, reaching excellent performance with 22 sampled lines and continuing to make, small, improvements. The *Policy Reconstruction* over the same interval makes sub-optimal (to classification) decisions before it plateaus. This behaviour can be explained by the fact that a reconstruction policy may not favour lines useful for classification. The *Policy Classifier* appears to change behaviour between 19 and 21 lines acquired. This can be attributed to a drift from how the classifier was pre-trained with randomly sampled lines. We discuss solutions to this later.

Task-specific ‘Coarse-to-Fine’ Sampling Policy

It would be interesting to see which k -space lines are favoured by our policy and how altering the starting point (namely different percentage of center fraction which represents the amount of low-frequency k -space intentionally sampled from the center of the k -space) modulates this behaviour. Figure 5 illustrates the average masks provided by three policy networks. Non-optimised random sampling mask (population-level mask) is also included as a comparison. It is evident that non-optimised and optimised masks vary in Figure 5, indicating

that a prefixed population-level mask would either over-acquire k -space lines for some patients or under-acquire for others compared to the optimised patient-level mask. Compared with the sampling policy for reconstruction in supplementary, this analysis reveals that early on the policy favours both low-frequency and some high-frequency data. However, later on the policy favours specific diagnostically-relevant k -space lines of high-frequency. Hence we see the policy capturing ‘Coarse’ features such as anatomical structures or essential patterns first and later ‘Fine’ details from the high-frequency data. When center fraction increases, the policy samples even more high-frequency information early on. While this might be obvious it is actually driven by low AUC performance when forcing the model to sample more center lines from the start (the AUC drops from 0.794 to 0.743 when the center fraction is set to 0.05, see also more metrics in the supplementary). By forcing the model to sample redundant low-frequency data, the policy tries to recover to identify diagnostically significant data within the k -space.

5 Discussion

Clinical Relevance: Our method actively samples k -space data whilst simultaneously conducting disease inference, enabling real-time diagnosis during scanning. This approach eliminates the need for sampling full k -space, thereby reducing acquisition time. In addition, our approach does not require high-fidelity fully sampled data for policy training, which is a hard requirement of reconstruction policies. We believe this has the potential to further accelerate the development of low-field MR. We envision applications in pre-screening by offering preliminary diagnostic results to aid in resource allocation, particularly in regions with limited medical resources.

Limitations: Our proof of concept is based on the fastMRI knee data and specifically MT. One could readily see the application of a portable low-field scanner in sports medicine [10] but it is a less compelling clinical application. Ideally, a dataset of e.g. stroke and acute brain trauma [9], which would probe diverse, and not only structural, MR signal characteristics would be more compelling. At this moment such dataset is currently lacking. Our policy network relies on a pre-trained classifier, which can propagate bias and domain shifts to the policy decisions. It is of great interest to train the classifier simultaneously with the policy network (e.g., by fine-tuning the classifier) to address domain shifts. Finally, our policy network does not leverage physical priors of k -space utility which are implicitly leveraged by reconstruction policies using a pre-trained reconstruction network. Such prior also may encode physical limitations [7] of the employed sequence in making sampling decisions.

6 Conclusion

Our proposed framework presents a novel approach to enhance the efficiency and accessibility of MRI as a Point of Care (POC) diagnostic tool by optimising

sampling patterns in a learnable fashion. Our approach distinctly stipulates that k -space sampling acquisition decisions are undertaken by an agent optimising a dynamic per-patient classification task. Results reporting direct inference from undersampled k -space data, concerning the presence of Meniscus Tear, showed that our approach achieves the comparable performance as an *Oracle* Classifier whilst reducing acquisition time with only 25% k -space usage. Furthermore, our analysis of task-specific sampling policies revealed that the policy network adapted its sampling strategy based on the nature of the task, demonstrating the adaptability of the active sampling approach.

Acknowledgments. This work was supported in part by National Institutes of Health (NIH) grant 7R01HL148788-03. Y. Du thanks additional financial support from the School of Engineering, the University of Edinburgh. S.A. Tsiftaris also acknowledges the support of Canon Medical and the Royal Academy of Engineering and the Research Chairs and Senior Research Fellowships scheme (grant RCSR1819\8\25), and the UK’s Engineering and Physical Sciences Research Council (EPSRC) support via grant EP/X017680/1. The authors would like to thank Dr. McDonagh and Dr. Moroshko for inspirational discussions and assistance.

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., Wang, A.Q., Dalca, A.V., Sabuncu, M.R.: Deep-learning-based optimization of the under-sampling pattern in mri. *IEEE Transactions on Computational Imaging* **6**, 1139–1152 (2020)
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. Cai, L., Gao, J., Zhao, D.: A review of the application of deep learning in medical image classification and segmentation. *Annals of translational medicine* **8**(11) (2020)
4. Geethanath, S., Vaughan Jr, J.T.: Accessible magnetic resonance imaging: a review. *Journal of Magnetic Resonance Imaging* **49**(7), e65–e77 (2019)
5. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV* 14. pp. 630–645. Springer (2016)
6. Kool, W., van Hoof, H., Welling, M.: Buy 4 reinforce samples, get a baseline for free! (2019)
7. Kuperman, V.: *Magnetic resonance imaging: physical principles and applications*. Elsevier (2000)
8. Lin, D.J., Johnson, P.M., Knoll, F., Lui, Y.W.: Artificial intelligence for mr image reconstruction: an overview for clinicians. *Journal of Magnetic Resonance Imaging* **53**(4), 1015–1028 (2021)
9. Lyu, M., Mei, L., Huang, S., Liu, S., Li, Y., Yang, K., Liu, Y., Dong, Y., Dong, L., Wu, E.X.: M4raw: A multi-contrast, multi-repetition, multi-channel mri k-space dataset for low-field mri research. *Scientific Data* **10**(1), 264 (2023)
10. Massimiliano, L., Giuseppe, G., Michela, B., Michele, A., Silvia, R., Alessio, P., Francesco, P., Lia, R., Alessandro, S., Federico, A.G., et al.: Role of low field mri in detecting knee lesions. *Acta Bio Medica: Atenei Parmensis* **90**(Suppl 1), 116 (2019)
11. Pineda, L., Basu, S., Romero, A., Calandra, R., Drozdal, M.: Active mr k-space sampling with reinforcement learning. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II* 23. pp. 23–33. Springer (2020)
12. Ravula, S., Levac, B., Jalal, A., Tamir, J.I., Dimakis, A.G.: Optimizing sampling patterns for compressed sensing mri with diffusion generative models. *arXiv preprint arXiv:2306.03284* (2023)
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. pp. 234–241. Springer (2015)
14. Schlemper, J., Oktay, O., Bai, W., Castro, D.C., Duan, J., Qin, C., Hajnal, J.V., Rueckert, D.: Cardiac mr segmentation from undersampled k-space using deep latent representation learning. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I* pp. 259–267. Springer (2018)
15. Singhal, R., Sudarshan, M., Mahishi, A., Kaushik, S., Ginocchio, L., Tong, A., Chandarana, H., Sodickson, D.K., Ranganath, R., Chopra, S.: On the feasibility of machine learning augmented magnetic resonance for point-of-care identification of disease. *arXiv preprint arXiv:2301.11962* (2023)

16. Sondik, E.J.: The optimal control of partially observable Markov processes. Stanford University (1971)
17. Wang, J., Yang, Q., Yang, Q., Xu, L., Cai, C., Cai, S.: Joint optimization of cartesian sampling patterns and reconstruction for single-contrast and multi-contrast fast magnetic resonance imaging. *Computer Methods and Programs in Biomedicine* **226**, 107150 (2022)
18. Xuan, K., Sun, S., Xue, Z., Wang, Q., Liao, S.: Learning mri k-space subsampling pattern using progressive weight pruning. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II* 23. pp. 178–187. Springer (2020)
19. 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)
20. Zhang, Z., Romero, A., Muckley, M.J., Vincent, P., Yang, L., Drozdal, M.: Reducing uncertainty in undersampled mri reconstruction with active acquisition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2049–2058 (2019)
21. 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 multi-coil mri data. *arXiv preprint arXiv:2109.03812* (2021)