

Adaptive Sampling of k -space in Magnetic Resonance for Rapid Pathology Prediction

Chen-Yu Yen^{*1} Raghav Singhal^{*1} Umang Sharma¹ Rajesh Ranganath¹ Sumit Chopra¹ Lerrel Pinto¹

Abstract

Magnetic Resonance (MR) imaging, despite its proven diagnostic utility, remains an inaccessible imaging modality for disease surveillance at the population level. A major factor rendering MR inaccessible is lengthy scan times. An MR scanner collects measurements associated with the underlying anatomy in the Fourier space, also known as the k -space. Creating a high-fidelity image requires collecting large quantities of such measurements, increasing the scan time. Traditionally to accelerate an MR scan, image reconstruction from under-sampled k -space data is the method of choice. However, recent works show the feasibility of bypassing image reconstruction and directly learning to detect disease directly from a sparser learned subset of the k -space measurements. In this work, we propose Adaptive Sampling for MR (ASMR), a sampling method that learns an adaptive policy to sequentially select k -space samples to optimize for target disease detection. On 6 out of 8 pathology classification tasks spanning the Knee, Brain, and Prostate MR scans, ASMR reaches within 2% of the performance of a fully sampled classifier while using only 8% of the k -space, as well as outperforming prior state-of-the-art work in k -space sampling such as EMRT, LOUPE, and DPS.

1. Introduction

Magnetic Resonance (MR) imaging is a diagnostic medical imaging tool that exhibits superior soft-tissue contrast in its generated images. This has led to MR scans being the gold standard for diagnosing a range of cancers (Winawer et al., 1997; Ilic et al., 2018; Elmore et al., 1998), and mus-

^{*}Equal contribution ¹New York University. Correspondence to: Chen-Yu Yen <chenyu.yen@nyu.edu>, Raghav Singhal <rsinghal@nyu.edu>.

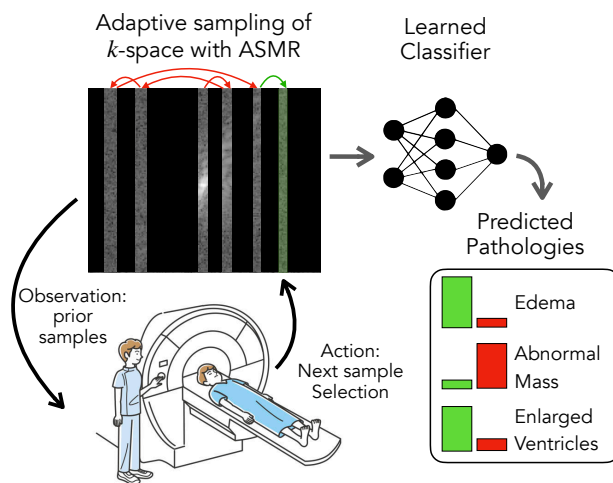


Figure 1. **Overview of ASMR.** Given a set of prior samples from the MR scanner, ASMR proposes the next sample to collect. These sampling steps are repeated for a fixed set of iterations T , at which point the collected samples are used by a classifier to predict the presence of a fixed set of pathologies. ASMR is trained to optimize classification performance and sidesteps reconstruction altogether.

culoskeletal disorders (Dean Deyle, 2011; Elvenes et al., 2000). An MR machine sequentially measures the responses to radio-frequency pulses of a human body placed in a magnetic field to generate cross-sectional images (slices) of the anatomy (Liang & Lauterbur, 2000; Zbontar et al., 2018). The measurements are acquired in the Fourier space, also known as the k -space (Liang & Lauterbur, 2000) in the MR community. The spatial image is then constructed by applying a multi-dimensional Fourier transform to the k -space data. High-resolution MR images require a large number of k -space measurements, leading to increased scan times (up to 40 minutes, depending on the anatomy) (Zbontar et al., 2018).

Reducing MR scan times is an active area of research. A common approach is to under-sample the k -space and then use image reconstruction techniques such as compressed sensing (Lustig et al., 2008; Zbontar et al., 2018; Donoho, 2006) and deep learning (Sriram et al., 2020; Chung & Ye, 2022) to reconstruct the underlying image accurately. A

recent challenge organized by NYU Langone (Knoll et al., 2020) shows that at a 12.5% sampling rate, reconstructed images started missing pathologies. This is understandable as reconstruction metrics, such as SSIM (Wang et al., 2004), measure the similarity between moments of localized patches rather than optimizing for diagnostic performance. To overcome this, recent work (Singhal et al., 2023; Tolpadi et al., 2023) shows that we can directly detect the presence of a pathology and bypass reconstruction altogether. Furthermore, Singhal et al. (2023) show that directly classifying from a sparse subset of the k -space yields similar performance to using the entire k -space. However, both Singhal et al. (2023); Tolpadi et al. (2023) use a heuristic approach to find the sampling pattern.

Typically, the sampling patterns are heuristic or hand-crafted (Lustig et al., 2007), based on the assumption that natural images have rapidly decaying Fourier coefficients (Lustig et al., 2007). However, Bahadir et al. (2019); Jin et al. (2019); Pineda et al. (2020); Huijben et al. (2019) show that learning the sampling pattern to optimize for the underlying task offers significant improvements over heuristic sampling patterns. In practice, a heuristic sampling pattern, such as an equi-spaced pattern (Johnson et al., 2022), at a fixed sampling rate is used for each patient (Zbontar et al., 2018). However, a heuristic pattern cannot adapt to the collection of previously acquired k -space samples.

In this work, we introduce Adaptive Sampling for Magnetic Resonance (ASMR). Using the fact that the k -space data collection is a sequential process, we pose the problem of learning the sampling pattern as a decision-making problem, which can then be optimized using reinforcement learning (RL). Concretely, at every sampling time step, the collection of previous samples serves as the observation, the choice of the next sample serves as the action, and the reward is computed by measuring the log-likelihood between the labels and the sub-sampled k -space. See Figure 1 for an overview of ASMR. The policy for selecting actions is optimized using Proximal Policy Optimization. Unlike prior work that uses RL for reconstruction (Bakker et al., 2020; Jin et al., 2019; Pineda et al., 2020), ASMR is trained to maximize the classification performance.

To understand the performance of ASMR, we run experiments on the FASTMRI dataset (Mathieu et al., 2013; Zhao et al., 2021), consisting of volumetric brain, knee and abdominal scans. For simplicity, we use Cartesian sampling (Liang & Lauterbur, 2000; Lustig et al., 2008) that selects columns of the 2D complex-valued k -space matrix. Our experiments reveal the following findings:

1. ASMR outperforms state-of-the-art non-adaptive sampling patterns such as EMRT (Singhal et al., 2023) on 6 out of 8 classification tasks. Compared to learned probabilistic sampling patterns like LOUPE (Bahadir

et al., 2019) and DPS, ASMR achieves an improvement of at least 2.5% in the AUC metrics on 7 out of 8 tasks (See Figure 5).

2. ASMR outperforms adaptive sampling methods (Bakker et al., 2020) that are optimized for reconstruction, providing an absolute gain of 2.88% and 7% in the AUC metrics on the Knee and Brain datasets respectively (See Figure 6).
3. Compared to an image-based classifier that uses the entire k -space data as input, ASMR comes within 2% of its performance on 6 out of 8 tasks with a 12.5 \times acceleration factor (See Figure 5).

Our code can be found at [adaptive-sampling-mr.github.io](https://github.com/adaptive-sampling-mr).

2. Related Work

ASMR builds upon several prior works in k -space-based prediction, mask sampling strategies, and adaptive learning.

Pathology prediction from sparse k -space: Singhal et al. (2023) show that directly classifying from sparsely acquired k -space data yields classification performance similar to an image-based model that uses as input images reconstructed using the entire k -space data. The sampling pattern they use to acquire the sparse k -space data is optimized using a non-adaptive variable density sampling prior (Lustig et al., 2007). Similarly, Tolpadi et al. (2023) show that directly generating segmentations as the main objective with a secondary reconstruction objective can also accelerate pathology prediction without any significant drop in performance. Interestingly, Tolpadi et al. (2023) show that there was no correlation between segmentation and reconstruction performance. In both works, the sampling patterns used were heuristic-based, non-adaptive, and not optimized for the downstream task.

Learned non-adaptive sampling methods: Several prior works (Bahadir et al., 2019; Huijben et al., 2019; Weiss et al., 2020) have looked at learning a probabilistic model over sampling patterns jointly with a reconstruction model. Gözcü et al. (2019); Weiss et al. (2019) learn a fixed sampling pattern using a pre-trained reconstruction model and the reconstruction error as a scoring rule for sampling patterns. The sampling patterns learned in these methods are non-adaptive and only produce sampling patterns for a fixed sampling budget. In Figure 4 we compare ASMR against sampling patterns optimized for classification using LOUPE (Bahadir et al., 2019) and DPS (Huijben et al., 2019).

Learned adaptive sampling methods: More recently, Bakker et al. (2020); Pineda et al. (2020); Zhang et al. (2019); Jin et al. (2019) pose learning an adaptive sampling pattern as an RL problem. For example Bakker et al.

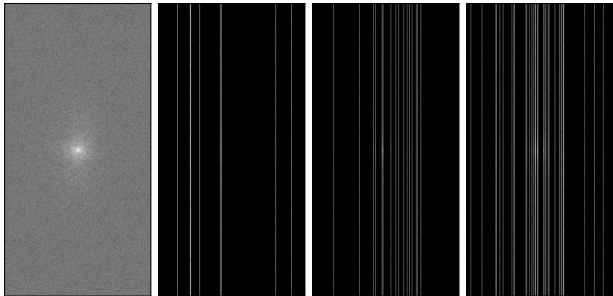


Figure 2. Cartesian Sampling on k -space at different sampling rates. The leftmost image shows the complete k -space followed 2.5%, 5% and 10% sampling rates (from left to right)

(2020) first trains a reconstruction model and then uses the reconstruction loss as a reward to train the policy. Zhang et al. (2019) also presents an active acquisition method, where instead of learning a policy model, they jointly train a reconstruction model and evaluator model. Similarly, Jin et al. (2019) jointly learn a reconstruction model and policy network using Monte-Carlo tree search, similar to AlphaGo (Silver et al., 2016). All these works use reconstruction metrics as the reward, which is often sub-optimal for pathology classification (Singhal et al., 2023). Figure 6 shows that ASMR, which is directly optimized for classification improves upon reconstruction-based sampling.

Reinforcement learning for adaptive sampling: Reinforcement learning for adaptive sampling is a promising area of research in the fields of gaming (Silver et al., 2016), finance (Hambly et al., 2023), and chatbots (Achiam et al., 2023). For instance, in the robotics domain, it has been used to improve controllers (Zucker et al., 2008; Pan et al., 2022), and path planning (Pan et al., 2022). Additionally, it has applications in autonomous balloon navigation (Bellemare et al., 2020), chip design and hardware optimization (Kumar et al., 2022; Chen et al., 2023). However, unlike applications where a large amount of data is available or new data can be easily synthesized, one of our challenges is the limited dataset size. Unlike scenarios where the reward is well-defined, another challenge is the choice of reward suitable for fast pathology classification.

3. Background and Setup

In this section, we provide background for understanding ASMR, which includes understanding how an MR scanner works, details on k -space classification, and adaptive sampling with reinforcement learning.

3.1. Magnetic Resonance (MR) Environment

MR imaging functions indirectly, acquiring frequency and phase measurements of the anatomy being imaged instead of direct spatial data. A receiver coil, situated near the target area, acquires a collection of measurements encoding the human body’s response to a combination of applied electromagnetic fields, including a static spatially homogeneous magnetic field, spatially varying magnetic fields (gradients), and radio-frequency pulses (the observed measurements) (Liang & Lauterbur, 2000). These measurements of the body’s response to different radio-frequency pulses correspond to points in the k -space and are stored in a matrix form called the k -space matrix. A k -space matrix is a complex-valued matrix $\mathbf{x} \in \mathbb{C}^{d_r \times d_c}$, which is the Fourier Transform of the underlying 2D MR image. Hence, an image can be reconstructed using an Inverse Fourier Transform of the k -space.

We represent the k -space sampling pattern using a two-dimensional binary mask $\mathbf{s} \in \prod_{i,j=1}^{d_r, d_c} \{0, 1\}$, such that $s_{i,j} = 1$ indicates that $\mathbf{x}_{i,j}$ has been sampled. Typically, the k -space is retrospectively under-sampled by multiplying the full k -space \mathbf{x} with the binary mask \mathbf{s} (Zbontar et al., 2018), which we denote as $\mathbf{x}_s = \mathbf{x} \odot \mathbf{s}$, where \odot stands for element-wise multiplication.

The k -space measurements are acquired *sequentially*. A MR scanner has physical restraints on how it can sample. Some examples of physically feasible trajectories include Cartesian, spiral, and radial (Zbontar et al., 2018; Weiss et al., 2019). In this work, we consider Cartesian sampling, which involves sampling columns of the k -space matrix sequentially. See Figure 2 for examples of Cartesian sampling on k -space. It can take a long time to acquire measurements sufficient for reconstructing a diagnostic quality image, thereby increasing the cost of the scan and making them infeasible for population-level screenings.

In the next section, we discuss recent work that bypasses image reconstruction and directly predicts the presence or absence of a set of pathologies.

3.2. Classification from k -space data

Given a labeled dataset ($\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$) consisting of k -space data \mathbf{x} and pathology labels \mathbf{y} , we follow Singhal et al. (2023) and directly train a classifier from the retrospectively under-sampled k -space \mathbf{x}_s using the KSPACE-NET classifier $q_\phi(\mathbf{y} | \mathbf{x}_s)$. Importantly, the KSPACE-NET classifier q_ϕ takes the complex-valued k -space as input and does not rely on reconstructed images. Similar to Jethani et al. (2021; 2023) and Singhal et al. (2023) we train the KSPACE-NET classifier using randomly masked out k -space data, where the masks are drawn independently from the variable density prior Lustig et al. (2007). Models trained in such a way can

then be used to score masks. For instance, we can estimate the mutual information (Cover, 1999) between \mathbf{x}_s and the label \mathbf{y} , upto a constant, using the classifier q_θ :

$$V_\theta(\mathbf{s}) = \mathbb{E}_{q(\mathbf{x}_s)q(\mathbf{y}|\mathbf{x}_s)} \log q_\theta(\mathbf{y} | \mathbf{x}_s) \quad (1)$$

In the Section 4, we show how this classifier can be used to design rewards for the policy.

3.3. Reinforcement Learning

The reinforcement learning problem formulates the decision-making problem as Markov Decision Process (MDP) $(\mathcal{S}, \mathcal{A}, p_0(s), p(s'|s, a), R(s, a), \gamma, T)$, where a is an action, s is the state, \mathcal{S} is a state space, \mathcal{A} is an action space, $p_0(s)$ is a distribution of initial states, $p(s'|s, a)$ is the environment dynamics, $R(s, a)$ is a reward function, γ is a discount factor, and T is the task horizon. The agent interacts with the MDP according to a policy $\pi(a|s)$. The goal of reinforcement learning is to obtain a policy that maximizes the cumulative discounted returns: $\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{t=0}^T \gamma^t R(s_t, a_t) \right]$. For more details on RL we refer the reader to Sutton & Barto (2018).

Proximal Policy Optimization (PPO): PPO (Schulman et al., 2017) is an RL algorithm that learns a policy π_θ and a value function V_ϕ with the goal of finding an optimal policy for an MDP. An actor proposes an action to the environment, and a critic estimates a value for the action proposed by the actor. To this end, PPO optimizes a surrogate objective

$$L = \mathbb{E} \left[\min(l_t(\theta) \hat{A}_t, \text{clip}(l_t(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_t) \right] \quad (2)$$

where $l_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\text{old}}(a_t|s_t)}$ denotes the likelihood ratio between the new and old policies, and $\hat{A}_t = R_t + \gamma V_\phi(s_{t+1}) - V_\phi(s_t)$ denotes an advantage function at timestep t . During training, PPO penalizes large changes to the policy to improve the stability of learning. For this work, we build our agent on top of an open-source implementation of PPO in (Huang et al., 2022).

4. Adaptive Sampling for Magnetic Resonance

In this section, we introduce Adaptive Sampling for Magnetic Resonance (ASMR). ASMR learns an adaptive sampling method for direct k -space classification by training an RL agent to select k -space samples sequentially. Unlike standard RL settings with access to an online environment, we retrospectively under-sample the k -space by applying a binary mask \mathbf{s} to the fully-sampled k -space \mathbf{x} . Furthermore, the action space and reward function require additional consideration given the constraints of the MR pathology prediction problem, such as data imbalance and the input being in the Fourier domain.

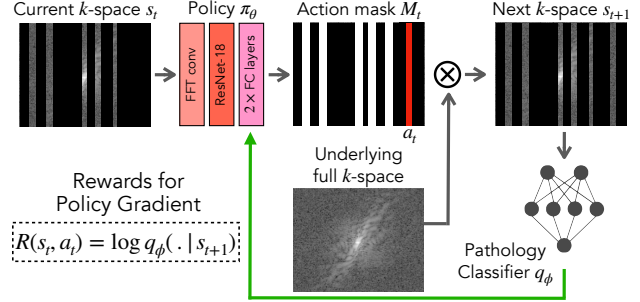


Figure 3. ASMR Training. ASMR takes an initial sub-sampled k -space s_t as input and proposes the next sample \mathbf{a}_t to generate the next state s_{t+1} . This state s_{t+1} is used by the reward model to compute the log-likelihood $q_\phi(\mathbf{y} | \mathbf{x}_{s_{t+1}})$, the reward for the actor. We repeat these steps for a fixed number of iterations T .

4.1. Formulating k -space selection as an RL Problem

To frame the selection of k -space measurements as a RL problem, we simulate environments using a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$, where each k -space sample $\mathbf{x}_i \in \mathbb{C}^{d_r \times d_c}$ represents one environment and has an associated pathology label \mathbf{y}_i . We employ a Cartesian sampling pattern, where an acquisition corresponds to a column of the k -space matrix.

Within each environment, the goal of the RL agent is to sequentially select a set of $T = |d_c| \times \alpha$ k -space columns to maximize classification performance, where T represents the sampling budget and serves as our episode length. The agent is tasked to take action $a_t \in \{0, 1, \dots, d_c\}$ at each time step t , which corresponds to selecting a column from the k -space matrix. From an implementation standpoint, we zero-fill columns that have not been selected by the agent yet. Hence, the task for the RL agent can also be framed as a mask selection problem where the objective is to obtain a mask M_T that maximizes classification performance. Starting with an initial empty mask M_0 , the final mask M_T is obtained by iteratively updating the masks as $M_{t-1}[a_t] = 1$. The actions a_t are produced by the policy π_θ , which acts on the current state s_t . We next describe the state and action spaces.

State space for ASMR: Let $\mathbf{x}^F \in \mathbb{C}^{d_r \times d_c}$ represent the fully acquired k -space data for a random sample from \mathcal{D} . The input state at time t is defined as $s_t = M_t \odot \mathbf{x}^F$, where M_t is the mask, a binary vector defined as $M_t \in \prod_{i=1}^{d_c} \{0, 1\}$. Initially, at $t = 0$ we set $M_0 = 0$ and $s_0 = 0$. We emphasize that \mathbf{x}^F is utilized merely for notational simplicity. The full k -space data, \mathbf{x}^F , is never disclosed to the agent, who can only observe the acquired data. Naturally, during inference, the agent can only observe the k -space data it has acquired till step t . Given an action a_t and state

s_{t+1} , the next state at time $t + 1$ is updated as:

$$s_{t+1} = M_{t+1} \odot \mathbf{x}^F \quad (3)$$

Action space for ASMR: We use a discrete action space, which is denoted as the set $\mathcal{A} = \{1, \dots, d_c\}$ where each element i of \mathcal{A} corresponds to selecting the i^{th} k -space column from a given sample. Unlike common RL settings, our action space dynamically shrinks as we iterate through the episode. To handle this shrinkage, we need to iteratively mask out already selected actions, which is discussed next.

Dynamic masking of action space: To account for the reduced action space after each selection, we maintain a set $\Omega_{t-1} = \{a_1, \dots, a_{t-1}\}$ to remove actions already selected in previous time steps of the current episode. At the step t , we update the set $\Omega_t = \Omega_{t-1} \cup \{a_t\}$ and reduce the action space $\mathcal{A}_t = \mathcal{A}_{t-1} \setminus \{\Omega_{t-1}\}$. We generate a mask M_t from the set Ω_t as:

$$M_t = (m_{ij})_{i=1, j=1}^{d_r, d_c} \quad \text{where} \quad m_{ij} = \begin{cases} 1 & \text{if } j \in \Omega_t \\ 0 & \text{otherwise} \end{cases}$$

An alternative approach to address the dynamic masking issue is to assign a negative reward for selecting a previously acquired measurement. However, this is inefficient in practice, and our experiments in Section 5.4 indicate a strong advantage for our explicit masking technique.

4.2. Reward Model

Typically, RL problem settings have access to an oracle reward model directly as part of the environment, for instance, the game score in Atari (Silver et al., 2016). The agent then uses the rewards to optimize the policy parameters θ . In the ASMR setting, we do not have ‘oracle’ rewards but instead have access to underlying pathology labels. However, it is not immediately obvious how to use them to provide intermediate rewards for our agent.

To address the lack of an oracle reward model, we learn a mapping between the undersampled k -space data and a quantifiable measure of how well the acquired undersampled k -space relates to the labels. As such, we repurpose the models used for direct k -space classification by EMRT (Singhal et al., 2023) to compute the log-likelihood between the state \mathbf{s}_t and the label \mathbf{y} .

Since our agent sequentially acquires k -space data we need our reward model to be able to provide rewards across varying undersampling rates. While a different classifier for each sampling rate can be used, following Jethani et al. (2021; 2023) we train the classifier using random sampling patterns across sampling rates. Such a training approach allows us to train a single classifier that is able to evaluate samples

across all rates. For a given state s_t at time t , the reward model is defined as

$$R(s_t, a_t) = \log q_\phi(\mathbf{y}|s_{t+1})$$

where q_ϕ is a classifier and s_{t+1} is obtained as in Equation (3). The classifier is frozen while training the policy.

Policy architecture: Following the PPO algorithm (Schulman et al., 2017), we adopt the actor-critic architecture. Since our input is in the Fourier domain, we adapt the KSPACE-NET in Singhal et al. (2023) for the actor and critic models. The KSPACE-NET first performs a convolution in the Fourier domain (Mathieu et al., 2013). The output of the convolution, which is complex-valued, is then transformed into real values by taking the magnitude of its inverse Fourier transform. These real-valued features are then fed into a backbone network, where any modern convolutional network such as a ResNet (He et al., 2016a;b) can be plugged in. In this work, we use a ResNet-18 as the backbone. The features from the backbone are used as input to the actor and critic networks. The actor-network processes these features into a 2-layer feed-forward layer to produce a categorical distribution over the action space.

Balancing the training environment: Traditionally, training environments are sampled uniformly; however, medical datasets can have a highly imbalanced label distribution. For classification problems, this label imbalance can cause the classifier to predict the majority class and still achieve good classification performance. In such situations, a common remedy is to either use importance weighting or up-sampling. In this work, we train the policy by sampling environments in a manner that ensures the agent sees data from all classes with equal probability. For a given dataset labels $\{\mathbf{y}_i\}_{i=1}^N$, we sample an environment \mathbf{x} with probability proportional to the inverse of its label frequency.

5. Experiments

Our experiments are designed to answer the following questions:

- How does ASMR compare to non-adaptive learned sampling methods?
- How do ASMR policies compare to policies optimized for image reconstruction?
- How important are the design choices to ASMR’s performance?

Datasets: We test ASMR on three different datasets consisting of knee, brain and prostate MR scans:

Adaptive Sampling of k-Space in Magnetic Resonance for Rapid Pathology Prediction

	Knee Dataset			Brain Dataset			Prostate Dataset	
	ACL	Mensc. Tear	Abnormal	Edema	Enlg. Ventricles	Mass	Abnormal	CS-PCA
Train slices	29100 (3.6%)	29100 (11.55%)	29100 (15.9%)	9420 (1.89%)	9420 (1.66%)	9420 (1.83%)	9420 (24.06%)	6649 (5%)
Validation slices	6298 (2.49%)	6298 (11.08%)	6298 (15.39%)	3148 (2.06%)	3148 (1.43%)	3148 (2.6%)	3148 (20.71%)	1431 (4.5%)
Test slices	6281 (3.58%)	6281 (11.94%)	6281 (16.75%)	3162 (2.44%)	3162 (2.3%)	3162 (2.4%)	3162 (23.88%)	1462 (6%)

Table 1. Number of slices in the training, validation and test splits for each task. The number in the bracket is the percentage of slices with a pathology.

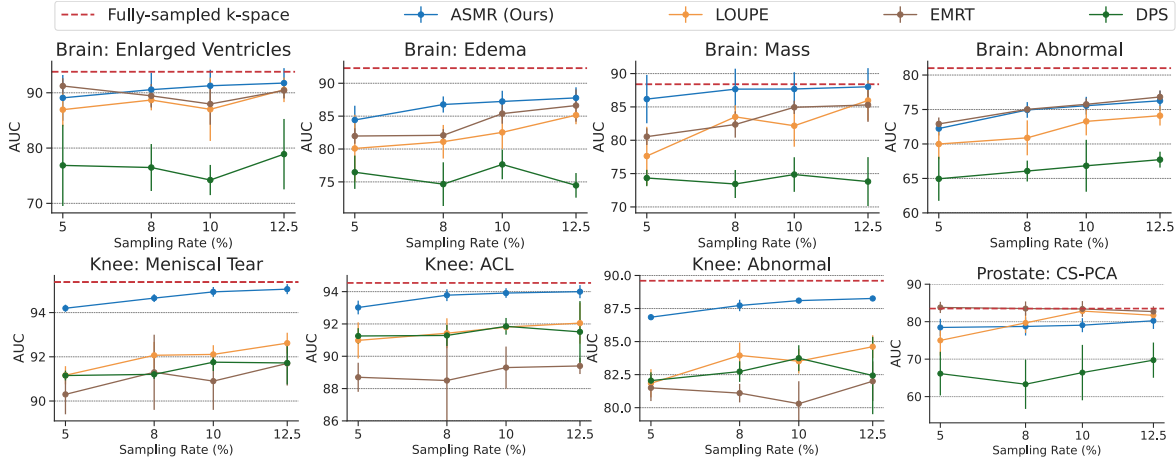


Figure 4. AUROCs obtained by ASMR compared to Learned Non-Adaptive Methods; the horizontal dotted line in red denotes the performance of an image-based classifier (which uses the entire k -space data). ASMR outperforms LOUPE and DPS for 6 out of 8 tasks, and outperforms EMRT on 7 out of 8 tasks. All results are computed over 5 seeds, and plotted with their means and standard deviations.

- Knee scans:** For the knee dataset, we use the FASTMRI knee dataset (Zbontar et al., 2018) with slice-level labels provided by Zhao et al. (2021). The training, validation, and test splits contain 816, 176, and 175 volumes, respectively. Following Bien et al. (2018); Singhal et al. (2023), we predict the presence of Meniscal Tears and ACL sprains in each slice. Additionally, we also predict an extra category called abnormal, which is a collection of pathologies that are less frequently observed (Bien et al., 2018) in the dataset.
- Brain scans:** We use the FASTMRI brain dataset with labels from Zhao et al. (2021). Our training, validation, and test splits contain 600, 200, and 201 volumes, respectively. We predict the presence of Enlarged Ventricles, Mass, Edema, and an abnormal category.
- Prostate scans:** For the prostate MR scans, we detect the presence of clinically significant prostate cancer (CS-PCA). CS-PCA is defined as a lesion within the prostate for which a radiologist assigns a Prostate Imaging Reporting and Data System (PI-RADS) score (Weinreb et al., 2016) greater or equal to 3. The data was split into 218, 48, and 47 volumes for training, validation, and testing, respectively.

For each dataset, we predict the presence of pathologies in each slice. To prevent label leakage we split each dataset at

the volume level, ensuring slices from the same volume do not spill across different splits. Slice level splits along with positivity rates are provided in Table 1

Due to computational constraints, we use the single-coil k -space generated from the multi-coil k -space matrix. Similar to Zbontar et al. (2018), we use the emulated single-coil (Tygert & Zbontar, 2020) to convert the multi-coil brain and prostate into single-coil. The emulated single-coil method takes a complex-valued linear combination of coils to produce a single matrix, where the combination of weights is learned per sample. For details see (Tygert & Zbontar, 2020; Zbontar et al., 2018). With additional computational resources, ASMR can be extended to multi-coil by just adapting the KSPACE-NET to accept multi-coil k -space as input.

Mask Distributions: In Appendix E, we plot the distribution of masks selected by ASMR on each of the datasets.

Evaluation metrics. For comparisons on imbalanced datasets, a commonly used metric is the area under the receiver operator curve (AUROC). In Appendix A, we provide classification metrics other than AUROC.

Baselines: We compare ASMR against several methods, including both adaptive and non-adaptive methods.

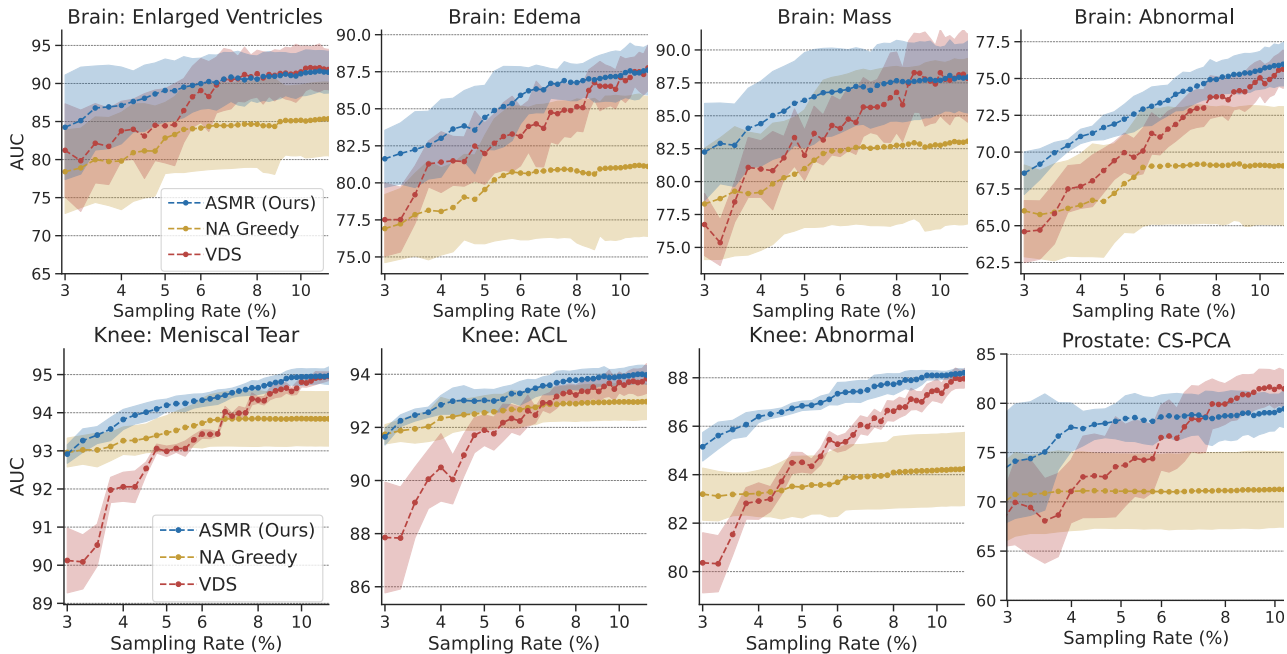


Figure 5. AUROCs obtained by ASMR compared to sequential sampling methods such as VDS random and non-adaptive greedy sequence. In the low sampling regime, ASMR consistently outperforms both baselines, while either outperforming or matching their performance at higher sampling rates. All results are computed over 5 seeds, and plotted with their means and standard deviations.

1. **EMRT:** EMRT learns a single classifier and a non-adaptive pattern for each sampling rate.
2. **Fully-Sampled Image-based Classifier:** A classifier trained on images generated from the fully-sampled k -space. This is the standard input for DL models performing image analysis.
3. **Learned Non-Adaptive Probabilistic Methods:** For a fixed sampling rate, we repurpose LOUPE (Bahadir et al., 2019) and DPS (Huijben et al., 2019) to learn a classifier and a probabilistic sampling mask jointly.
4. **Random Policy:** We use a random policy that samples k -space columns using the variable density sampling (VDS) prior (Lustig et al., 2007).
5. **Reconstruction Optimized Policy:** We use the greedy policy method in (Bakker et al., 2020), using the same 16 channel U-Net (Ronneberger et al., 2015) as the reconstruction model using single-coil data.
6. **Non-Adaptive Greedy Sequence:** A sequential non-adaptive method that uses the same KSPACE-NET used in ASMR to do forward greedy selection Macedo et al. (2019), producing a sequence of masks s_1, \dots, s_d .

Pathology classifier: For classifiers that use the k -space as input we use the KSPACE-NET architecture (Singhal et al.,

2023) with either a RESNET-18 or RESNET-50 as the backbone architecture. For methods that take an image as input, we use a RESNET-50.

5.1. Comparing ASMR to non-adaptive methods

Following (Singhal et al., 2023), we train EMRT with the KSPACE-NET classifier using random masks sampled from the VDS prior. Then, using the scoring rule defined in Equation (1) we sample $K = 100$ masks at a fixed sampling rate from the VDS prior and select the mask that maximizes the score rule. For gradient-based mask learning, we adapt DPS and LOUPE to learn a probabilistic mask and a classifier jointly. All these methods learn a classifier and a non-adaptive sampling pattern for a fixed sampling rate. We train EMRT, DPS, and LOUPE at 5%, 8%, 10%, and 12.5% sampling rates. Figure 4 compares the performance of ASMR over these non-adaptive methods. As an upper bound, we also provide the performance of an image-based classifier that uses the Inverse Fourier Transform of the fully sampled k -space as input. We note that ASMR approaches the performance of a fully sampled classifier using just 8% of the samples.

5.2. Comparing ASMR to sequential sampling methods

ASMR sequentially collects k -space samples, unlike non-adaptive methods like DPS and LOUPE. As such, we compare ASMR to other methods that perform sequen-

tial sampling. Following Macedo et al. (2019), we use the pre-trained classifier q_ϕ as in Section 4.2 and the scoring rule in Equation (1) to iteratively build a sequence $\mathbf{s}_{t+1} = \mathbf{s}_t + \arg \max_j \mathbf{V}(\mathbf{s}_t + \mathbf{e}_j)$, starting with $\mathbf{s}_0 = 0$. To evaluate the scoring rule, we use a balanced subset of the validation set to build this sequence. This is similar to the greedy non-adaptive oracle considered in Bakker et al. (2020). The greedy sequence suffers from high variance in classification metrics as can be observed across all datasets in Figure 5, and plateaus beyond a certain sampling rate.

On average, ASMR achieves 1.87%, 7.01%, and 9.82% absolute gains on AUROC over greedy solutions across the knee, brain, and prostate datasets, respectively. The VDS random policy, greedy sequence, and ASMR use the same pre-trained classifier for evaluation. Figure 5 shows a significant performance gap between ASMR and VDS, particularly in the low sampling rate regime ($\leq 8\%$ sampling rate). On average, we see a 1.33%, 2%, and 2.25% gain for ASMR over VDS on the knee, brain, and prostate datasets, respectively.

5.3. Comparing ASMR to reconstruction optimized policies

To show the benefits of learning policies for classification as opposed to image reconstruction, we follow the open source implementation of Bakker et al. (2020) to train a policy model to reconstruct scan images from undersampled k -space data. Since Bakker et al. (2020) requires specifying a target sampling rate, we train separate reconstruction-based policies for each sampling rate. Masks selected by the reconstruction-based policy are then evaluated using the same classifiers used to evaluate ASMR. Figure 6 compares the two policies in terms of the classification performance of different sampling rates. ASMR consistently outperforms reconstruction-based policies across all four pathologies in the brain dataset, achieving an average of 7% performance gains. Additional results for knee data and prostate data are included in Appendix B.1.

5.4. Importance of design decisions for ASMR

To understand the important design decisions of ASMR, we perform two ablations on our training regime choices, using a label-balanced training environment and dynamically masking the action space. We perform these experiments on the knee dataset only and evaluate selected samples by each ablated policy using the same set of classifiers that are used by ASMR. First, we investigate the effect of using a label-balanced training environment. For this, we implement a ‘‘Naive Environment Sampling’’, that samples the dataset uniformly during the training. Next, we ablate on the dynamic action space masking during training. Instead, we implement a reward penalty of -1 when the policy selects a

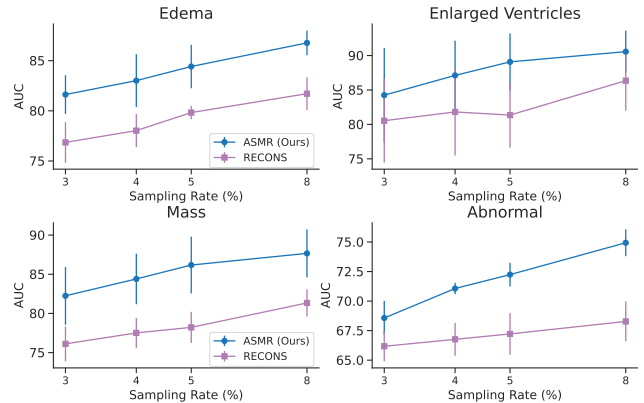


Figure 6. AUROCs obtained by ASMR compared to a reconstruction-based policy on the Brain test dataset. ASMR significantly outperforms these policies across all sampling rates, indicating the limitations of reconstruction objectives towards pathology prediction.

previously acquired k -space measurement. Figure 7 shows the results of our ablations. We hypothesize that without the label-balanced training environment, the learned policy is biased towards normal patients, causing performance drops across all sampling rates evaluated. Our design choices are crucial to improve the performance of ASMR over a vanilla PPO framework.

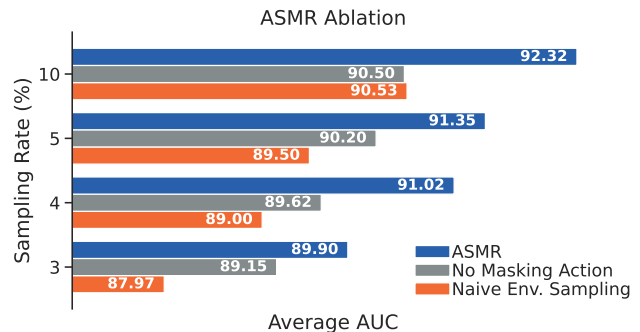


Figure 7. Ablation analysis on the Knee dataset for different components of ASMR, where significant performance gain is achieved.

6. Conclusion and Limitations

We have presented ASMR, an RL-based adaptive sampling technique for pathology prediction. ASMR improves over prior state-of-the-art methods in MR sampling and is a step towards enabling MR-based screening at a population scale as it offers strong pathology prediction using $12\times$ fewer k -space samples. However, several challenges still remain before widespread adoption.

Limitations: ASMR currently only uses single-coil k -space data while many modern MR scanners now collect multi-coil k -space data (Deshmane et al., 2012). We believe that ASMR can also be applied to the multi-coil setting as the under-sampling pattern for each coil is the same. Next, we note that due to limited volumetric data, we worked with slice-level classifiers rather than volume-level classifiers. Extending ASMR to volumetric classification or pathology segmentation is an exciting future direction.

Impact Statement

Methods that accelerate MR scans and enable wide-spread use of MR technology for screening need to be thoroughly validated prior to clinical adoption. Suddenly increasing the screening population can lead to false positives causing patient trauma and significant cost increases to the healthcare system (Kilpeläinen et al., 2010; Lafata et al., 2004). Therefore, prior to clinical adoption, methods such as ASMR need to be improved and undergo extensive clinical trials.

Acknowledgements

The computational requirements for this work were supported in part by the resources and personnel of the NYU Langone High Performance Computing (HPC) Core. LP is supported by the Packard Fellowship. We also thank Shenglong Wang and the NYU High Performance Computing for their support.

References

- Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Bahadir, C. D., Dalca, A. V., and 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.
- Bakker, T., van Hoof, H., and Welling, M. Experimental design for mri by greedy policy search. *Advances in Neural Information Processing Systems*, 33, 2020.
- Bellemare, M. G., Candido, S., Castro, P. S., Gong, J., Machado, M. C., Moitra, S., Ponda, S. S., and Wang, Z. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836):77–82, 2020.
- Bien, N., Rajpurkar, P., Ball, R. L., Irvin, J., Park, A., Jones, E., Bereket, M., Patel, B. N., Yeom, K. W., Shpanskaya, K., et al. Deep-learning-assisted diagnosis for knee magnetic resonance imaging: development and retrospective validation of mrnet. *PLoS medicine*, 15(11):e1002699, 2018.
- Chen, H., Zhang, X., Huang, K., and Koushanfar, F. Adatest: Reinforcement learning and adaptive sampling for on-chip hardware trojan detection. *ACM Transactions on Embedded Computing Systems*, 22(2):1–23, 2023.
- Chung, H. and Ye, J. C. Score-based diffusion models for accelerated mri. *Medical image analysis*, 80:102479, 2022.
- Cover, T. M. *Elements of information theory*. John Wiley & Sons, 1999.
- Dean Deyle, G. The role of mri in musculoskeletal practice: a clinical perspective. *Journal of Manual & Manipulative Therapy*, 19(3):152–161, 2011.
- Deshmane, A., Gulani, V., Griswold, M. A., and Seiberlich, N. Parallel mr imaging. *Journal of Magnetic Resonance Imaging*, 36(1):55–72, 2012.
- Donoho, D. L. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- Elmore, J. G., Barton, M. B., Mocerri, V. M., Polk, S., Arena, P. J., and Fletcher, S. W. Ten-year risk of false positive screening mammograms and clinical breast examinations. *New England Journal of Medicine*, 338(16):1089–1096, 1998.
- Elvenes, J., Jerome, C., Reikerås, O., and Johansen, O. Magnetic resonance imaging as a screening procedure to avoid arthroscopy for meniscal tears. *Archives of orthopaedic and trauma surgery*, 120(1-2):14–16, 2000.
- Gözcü, B., Sanchez, T., and Cevher, V. Rethinking sampling in parallel mri: A data-driven approach. In *2019 27th European Signal Processing Conference (EUSIPCO)*, pp. 1–5. IEEE, 2019.
- Hambly, B., Xu, R., and Yang, H. Recent advances in reinforcement learning in finance. *Mathematical Finance*, 33(3):437–503, 2023. doi: <https://doi.org/10.1111/mafi.12382>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/mafi.12382>.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016a.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016b.

- Huang, S., Dossa, R. F. J., Ye, C., Braga, J., Chakraborty, D., Mehta, K., and Araújo, J. G. Cleanrl: High-quality single-file implementations of deep reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274): 1–18, 2022. URL <http://jmlr.org/papers/v23/21-1342.html>.
- Huijben, I. A., Veeling, B. S., and van Sloun, R. J. Deep probabilistic subsampling for task-adaptive compressed sensing. In *International Conference on Learning Representations*, 2019.
- Ilic, D., Djulbegovic, M., Jung, J. H., Hwang, E. C., Zhou, Q., Cleves, A., Agoritsas, T., and Dahm, P. Prostate cancer screening with prostate-specific antigen (psa) test: a systematic review and meta-analysis. *Bmj*, 362, 2018.
- Jethani, N., Sudarshan, M., Aphinyanaphongs, Y., and Ranganath, R. Have we learned to explain?: How interpretability methods can learn to encode predictions in their interpretations. In *International Conference on Artificial Intelligence and Statistics*, pp. 1459–1467. PMLR, 2021.
- Jethani, N., Saporta, A., and Ranganath, R. Don’t be fooled: label leakage in explanation methods and the importance of their quantitative evaluation. In *International Conference on Artificial Intelligence and Statistics*, pp. 8925–8953. PMLR, 2023.
- Jin, K. H., Unser, M., and Yi, K. M. Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547*, 2019.
- Johnson, P. M., Tong, A., Donthireddy, A., Melamud, K., Petrocelli, R., Smereka, P., Qian, K., Keerthivasan, M. B., Chandarana, H., and Knoll, F. Deep learning reconstruction enables highly accelerated biparametric mr imaging of the prostate. *Journal of Magnetic Resonance Imaging*, 56(1):184–195, 2022.
- Kilpeläinen, T., Tammela, T., Määtänen, L., Kujala, P., Stenman, U.-H., Ala-Opas, M., Murtola, T., and Auvinen, A. False-positive screening results in the finnish prostate cancer screening trial. *British journal of cancer*, 102(3): 469–474, 2010.
- Knoll, F., Murrell, T., Sriram, A., Yakubova, N., Zbontar, J., Rabbat, M., Defazio, A., Muckley, M. J., Sodickson, D. K., Zitnick, C. L., et al. Advancing machine learning for mr image reconstruction with an open competition: Overview of the 2019 fastmri challenge. *Magnetic resonance in medicine*, 84(6):3054–3070, 2020.
- Kumar, A., Yazdanbakhsh, A., Hashemi, M., Swersky, K., and Levine, S. Data-driven offline optimization for architecting hardware accelerators. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=GSH-KlVIyy>.
- Lafata, J. E., Simpkins, J., Lamerato, L., Poisson, L., Divine, G., and Johnson, C. C. The economic impact of false-positive cancer screens. *Cancer Epidemiology and Prevention Biomarkers*, 13(12):2126–2132, 2004.
- Liang, Z.-P. and Lauterbur, P. C. *Principles of magnetic resonance imaging*. SPIE Optical Engineering Press Bellingham, 2000.
- Lustig, M., Donoho, D., and 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.
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. Compressed sensing mri. *IEEE signal processing magazine*, 25(2):72–82, 2008.
- Macedo, F., Oliveira, M. R., Pacheco, A., and Valadas, R. Theoretical foundations of forward feature selection methods based on mutual information. *Neurocomputing*, 325:67–89, 2019.
- Mathieu, M., Henaff, M., and LeCun, Y. Fast training of convolutional networks through ffts. *arXiv preprint arXiv:1312.5851*, 2013.
- Pan, L., Manjanna, S., and Hsieh, M. A. Marlas: Multi agent reinforcement learning for cooperated adaptive sampling. *arXiv preprint arXiv:2207.07751*, 2022.
- Pineda, L., Basu, S., Romero, A., Calandra, R., and Drozdal, 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.
- Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.

- Singhal, R., Sudarshan, M., Mahishi, A., Kaushik, S., Ginocchio, L., Tong, A., Chandarana, H., Sodickson, D. K., Ranganath, R., and Chopra, S. On the feasibility of machine learning augmented magnetic resonance for point-of-care identification of disease. *arXiv preprint arXiv:2301.11962*, 2023.
- Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C. L., Yakubova, N., Knoll, F., and 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.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Tolpadi, A. A., Bharadwaj, U., Gao, K. T., Bhattacharjee, R., Gassert, F. G., Luitjens, J., Giesler, P., Morshuis, J. N., Fischer, P., Hein, M., et al. K2s challenge: From undersampled k-space to automatic segmentation. *Bioengineering*, 10(2):267, 2023.
- Tygart, M. and Zbontar, J. Simulating single-coil mri from the responses of multiple coils. *Communications in Applied Mathematics and Computational Science*, 15(2): 115–127, 2020.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Weinreb, J. C., Barentsz, J. O., Choyke, P. L., Cornud, F., Haider, M. A., Macura, K. J., Margolis, D., Schnall, M. D., Shtern, F., Tempany, C. M., et al. Pi-rads prostate imaging–reporting and data system: 2015, version 2. *European urology*, 69(1):16–40, 2016.
- Weiss, T., Senouf, O., Vedula, S., Michailovich, O., Zibulevsky, M., and Bronstein, A. Pilot: Physics-informed learned optimized trajectories for accelerated mri. *arXiv preprint arXiv:1909.05773*, 2019.
- Weiss, T., Vedula, S., Senouf, O., Michailovich, O., Zibulevsky, M., and Bronstein, A. Joint learning of cartesian under sampling andre construction for accelerated mri. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8653–8657. IEEE, 2020.
- Winawer, S. J., Fletcher, R. H., Miller, L., Godlee, F., Stolar, M., Mulrow, C., Woolf, S., Glick, S., Ganiats, T., Bond, J., et al. Colorectal cancer screening: clinical guidelines and rationale. *Gastroenterology*, 112(2):594–642, 1997.
- 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.
- Zhang, Z., Romero, A., Muckley, M. J., Vincent, P., Yang, L., and Drozdal, M. Reducing uncertainty in under-sampled mri reconstruction with active acquisition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2049–2058, 2019.
- Zhao, R., Yaman, B., Zhang, Y., Stewart, R., Dixon, A., Knoll, F., Huang, Z., Lui, Y. W., Hansen, M. S., and Lungren, M. P. fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil mri data. *arXiv preprint arXiv:2109.03812*, 2021.
- Zucker, M., Kuffner, J., and Bagnell, J. A. Adaptive workspace biasing for sampling-based planners. In *2008 IEEE International Conference on Robotics and Automation*, pp. 3757–3762. IEEE, 2008.

Appendix

A. Comparison to Learning-based Methods

Here we present comparisons of ASMR to methods that learn a non-adaptive mask. Table 2 shows the AUC metrics across 3 pathologies in the Knee dataset. We provide the mean and standard deviations across 5 seeds. Table 3 and Table 4 provide the same results for pathologies in the Brain and Prostate datasets respectively.

Pathology - Sampling Rate (%)	DPS	EMRT	LOUPE	ASMR (Ours)
Abnormal - 5	82.0 ± 0.6	81.5 ± 1.0	81.8 ± 1.1	86.8 ± 0.2
Abnormal - 8	82.7 ± 0.8	81.1 ± 0.7	83.9 ± 1.0	87.7 ± 0.4
Abnormal - 10	83.7 ± 1.0	80.3 ± 1.7	83.5 ± 0.9	88.1 ± 0.2
Abnormal - 12.5	82.4 ± 2.9	82.0 ± 1.5	84.6 ± 0.9	88.3 ± 0.2
ACL - 5	91.3 ± 0.5	88.7 ± 0.9	91.0 ± 1.1	93.0 ± 0.4
ACL - 8	91.3 ± 0.7	88.5 ± 2.7	91.4 ± 0.9	93.8 ± 0.4
ACL - 10	91.9 ± 0.5	89.3 ± 1.3	91.8 ± 0.2	93.9 ± 0.3
ACL - 12.5	91.5 ± 1.9	89.4 ± 0.5	92.1 ± 1.3	94.0 ± 0.4
Mensc. Tear - 5	91.2 ± 0.2	90.3 ± 0.9	91.2 ± 0.4	94.2 ± 0.2
Mensc. Tear - 8	91.2 ± 0.2	91.3 ± 1.7	92.1 ± 0.6	94.7 ± 0.2
Mensc. Tear - 10	91.8 ± 0.4	90.9 ± 1.3	92.1 ± 0.4	94.9 ± 0.2
Mensc. Tear - 12.5	91.7 ± 1.0	91.7 ± 1.0	92.6 ± 0.5	95.1 ± 0.2

Table 2. Test AUC for identifying pathologies under different sampling rates on the Knee dataset

Pathology - Sampling Rate (%)	DPS	EMRT	LOUPE	ASMR (Ours)
Edema - 5	76.5 ± 2.5	82.0 ± 1.2	80.1 ± 5.1	84.4 ± 2.2
Edema - 8	74.7 ± 3.3	82.1 ± 0.9	81.1 ± 2.6	86.8 ± 1.2
Edema - 10	77.7 ± 2.3	85.4 ± 1.5	82.5 ± 2.5	87.2 ± 1.6
Edema - 12.5	74.5 ± 1.9	86.6 ± 2.5	85.1 ± 1.4	87.8 ± 1.6
Enlg. Ventricles - 5	76.9 ± 7.3	91.2 ± 1.4	86.9 ± 3.0	89.1 ± 4.1
Enlg. Ventricles - 8	76.5 ± 4.2	89.4 ± 2.1	88.7 ± 1.9	90.6 ± 3.1
Enlg. Ventricles - 10	74.2 ± 2.7	88.0 ± 3.7	87.0 ± 5.7	91.3 ± 2.9
Enlg. Ventricles - 12.5	78.9 ± 6.4	90.4 ± 1.5	90.5 ± 2.2	91.8 ± 2.7
Mass - 5	74.3 ± 1.2	80.5 ± 1.3	77.6 ± 4.4	86.2 ± 3.6
Mass - 8	73.4 ± 2.1	82.4 ± 2.0	83.5 ± 1.7	87.7 ± 3.1
Mass - 10	74.9 ± 2.6	84.9 ± 1.0	82.2 ± 3.1	87.7 ± 2.5
Mass - 12.5	73.8 ± 3.7	85.3 ± 2.5	86.0 ± 3.0	88.0 ± 2.8
Abnormal - 5	64.9 ± 3.2	72.9 ± 0.9	70.0 ± 2.2	72.2 ± 1.0
Abnormal - 8	66.1 ± 1.5	75.0 ± 0.4	70.9 ± 2.5	74.9 ± 1.1
Abnormal - 10	66.8 ± 3.8	75.8 ± 0.9	73.3 ± 2.0	75.6 ± 1.3
Abnormal - 12.5	67.7 ± 1.2	76.8 ± 0.9	74.1 ± 1.4	76.3 ± 1.5

Table 3. Test AUC for identifying pathologies under different sampling rates on the Brain dataset

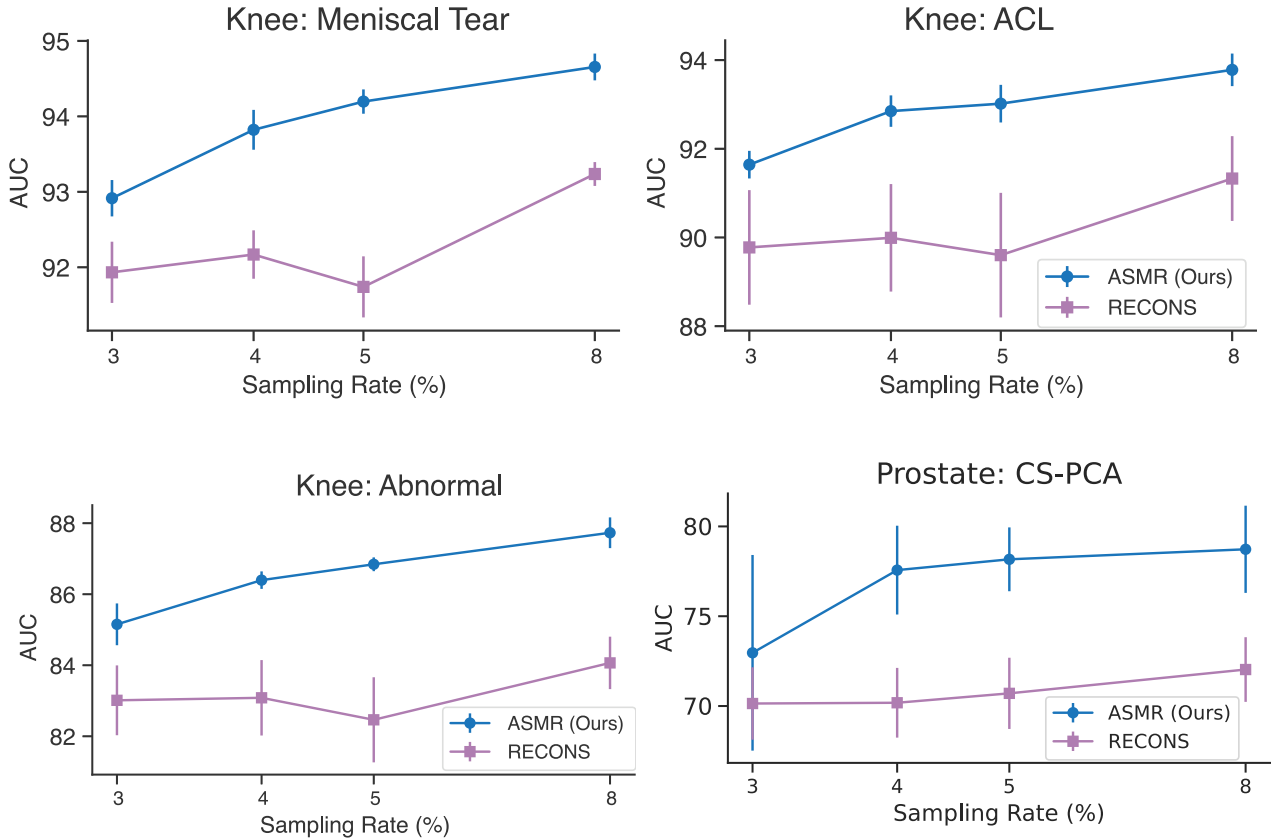
Pathology - Sampling Rate (%)	EMRT	DPS	LOUPE	ASMR (Ours)
CS-PCA - 5	83.8 ± 1.5	66.14 ± 5.8	75.00 ± 4.1	78.46 ± 2.3
CS-PCA - 8	83.5 ± 1.9	63.30 ± 6.6	79.68 ± 2.7	78.73 ± 2.4
CS-PCA - 10	83.4 ± 2.1	66.41 ± 7.4	82.83 ± 0.6	79.06 ± 1.9
CS-PCA - 12.5	82.7 ± 1.4	69.73 ± 4.7	81.70 ± 1.4	80.24 ± 2.2

Table 4. Test AUCs for identifying pathologies under different sampling rates on the Prostate dataset

B. Additional Experiments

B.1. Comparison to Reconstruction Optimized Policies on Knee and Prostate Dataset

To demonstrate the utility of optimizing a policy for classification performance, following Bakker et al. (2020) we also train policies to optimize for reconstruction errors. Using k -space subsets selected by the reconstruction optimized policy, we evaluate these subsets using our pathology classifier. Appendix B.1 shows the results of these experiments and indicate that reconstruction-based objectives are not necessarily suited for classification.



C. Additional Classification Metrics

C.1. Knee Results of ASMR

Pathology - Sampling Rate (%)	Balanced Accuracy	Sensitivity	Specifity	NPV
ACL - 5	86.1 ± 0.5	89.8 ± 0.0	82.4 ± 1.1	99.5 ± 0.0
ACL - 8	86.5 ± 0.7	89.8 ± 0.0	83.2 ± 1.5	99.5 ± 0.0
ACL - 10	86.8 ± 0.6	89.8 ± 0.0	83.9 ± 1.2	99.6 ± 0.0
ACL - 12.5	86.7 ± 0.3	89.8 ± 0.0	83.7 ± 0.7	99.5 ± 0.0
Mensc. Tear - 5	86.8 ± 0.3	89.9 ± 0.1	83.7 ± 0.5	98.4 ± 0.0
Mensc. Tear - 8	87.6 ± 0.5	89.9 ± 0.0	85.3 ± 1.0	98.4 ± 0.0
Mensc. Tear - 10	88.0 ± 0.3	89.9 ± 0.0	86.1 ± 0.6	98.4 ± 0.0
Mensc. Tear - 12.5	88.0 ± 0.5	89.9 ± 0.0	86.2 ± 1.0	98.4 ± 0.0
Abnoraml - 5	76.5 ± 0.9	89.9 ± 0.0	63.2 ± 1.8	96.9 ± 0.1
Abnoraml - 8	77.6 ± 0.8	89.9 ± 0.0	65.3 ± 1.6	97.0 ± 0.1
Abnoraml - 10	77.7 ± 0.4	89.9 ± 0.0	65.4 ± 0.9	97.0 ± 0.0
Abnoraml - 12.5	77.6 ± 0.6	89.9 ± 0.0	65.4 ± 1.2	97.0 ± 0.1

C.2. Brain Results of ASMR

Pathology - Sampling Rate (%)	Balanced Accuracy	Sensitivity	Specifity	NPV
Edema - 5	74.2 ± 2.5	89.8 ± 0.0	82.4 ± 1.1	99.5 ± 0.0
Edema - 8	76.9 ± 1.1	89.8 ± 0.0	83.2 ± 1.5	99.5 ± 0.0
Edema - 10	77.6 ± 2.1	89.8 ± 0.0	83.9 ± 1.2	99.6 ± 0.0
Edema - 12.5	78.7 ± 1.9	89.8 ± 0.0	83.7 ± 0.7	99.5 ± 0.0
Enlg. Ventricles - 5	78.1 ± 7.1	89.9 ± 0.1	83.7 ± 0.5	98.4 ± 0.0
Enlg. Ventricles - 8	78.9 ± 7.2	89.9 ± 0.0	85.3 ± 1.0	98.4 ± 0.0
Enlg. Ventricles - 10	79.8 ± 7.3	89.9 ± 0.0	86.1 ± 0.6	98.4 ± 0.0
Enlg. Ventricles - 12.5	81.2 ± 7.3	89.9 ± 0.0	86.2 ± 1.0	98.4 ± 0.0
Mass - 5	74.6 ± 3.7	89.9 ± 0.0	63.2 ± 1.8	96.9 ± 0.1
Mass - 8	76.5 ± 3.3	89.9 ± 0.0	65.3 ± 1.6	97.0 ± 0.1
Mass - 10	77.1 ± 2.3	89.9 ± 0.0	65.4 ± 0.9	97.0 ± 0.0
Mass - 12.5	77.0 ± 1.8	89.9 ± 0.0	65.4 ± 1.2	97.0 ± 0.1
Abnormal - 5	62.1 ± 2.0	72.9 ± 0.9	70.0 ± 2.2	72.2 ± 1.0
Abnormal - 8	64.1 ± 1.3	75.0 ± 0.4	70.9 ± 2.5	74.9 ± 1.1
Abnormal- 10	64.9 ± 1.8	75.8 ± 0.9	73.3 ± 2.0	75.6 ± 1.3
Abnormal - 12.5	65.4 ± 1.6	76.8 ± 0.9	74.1 ± 1.4	76.3 ± 1.5

C.3. Prostate Results of ASMR

Pathology - Sampling Rate (%)	Balanced Accuracy	Sensitivity	Specifity	NPV
CS-PCA - 5	68.3 ± 2.0	89.7 ± 0.0	47.0 ± 3.9	98.4 ± 0.1
CS-PCA - 8	70.0 ± 3.7	89.7 ± 0.0	50.2 ± 7.5	98.5 ± 0.2
CS-PCA - 10	68.4 ± 3.0	89.7 ± 0.0	47.2 ± 6.0	98.4 ± 0.2
CS-PCA - 12.5	68.0 ± 2.5	89.7 ± 0.0	46.4 ± 5.1	98.4 ± 0.2

D. Training Details

D.1. Training Parameters

The k -space data input to agent is in complex domain with height and width of $d_r \times d_c$. The $d_r \times d_c$ for knee, brain, prostate are (768×400) , (640×400) , and (320×451) , respectively. Hyperparameters used to train the policy are provided in Table 5.

Parameter	Value
Optimizer	Adamw
Learning rate	1e-04
weight decay	1e-04
discount factor	0.99
gae_lambda	0.95
clip ratio	0.2
entropy cost	0.01
grad norm clipping	0.5
value function coefficient	0.5
parallelized rollout	128

Table 5. Hyperparameters of our agent

E. Policy selected Mask Distributions

Below we plot a heatmap of the k -space column selections made by ASMR on the Knee, Brain and Prostate datasets. We evaluate our policy to have a sampling rate of 12.5% and generate heat maps representing the fraction of samples in the test set for which a column was picked.

