

# Constrained Probabilistic Mask Learning for Task-specific Undersampled MRI Reconstruction

Tobias Weber<sup>1,2</sup> Michael Ingrisch<sup>1,2</sup> Bernd Bischl<sup>1,2</sup> David Rügamer<sup>1,2</sup>  
<sup>1</sup> LMU Munich <sup>2</sup> Munich Center for Machine Learning (MCML)  
 tobias.weber@stat.uni-muenchen.de

## Abstract

*Undersampling is a common method in Magnetic Resonance Imaging (MRI) to subsample the number of data points in k-space, reducing acquisition times at the cost of decreased image quality. A popular approach is to employ undersampling patterns following various strategies, e.g., variable density sampling or radial trajectories. In this work, we propose a method that directly learns the undersampling masks from data points, thereby also providing task- and domain-specific patterns. To solve the resulting discrete optimization problem, we propose a general optimization routine called ProM: A fully probabilistic, differentiable, versatile, and model-free framework for mask optimization that enforces acceleration factors through a convex constraint. Analyzing knee, brain, and cardiac MRI datasets with our method, we discover that different anatomic regions reveal distinct optimal undersampling masks, demonstrating the benefits of using custom masks, tailored for a downstream task. For example, ProM can create undersampling masks that maximize performance in downstream tasks like segmentation with networks trained on fully-sampled MRIs. Even with extreme acceleration factors, ProM yields reasonable performance while being more versatile than existing methods, paving the way for data-driven all-purpose mask generation.*<sup>1</sup>

## 1. Introduction

Undersampling plays a crucial role in accelerating the acquisition time of magnetic resonance imaging (MRI) by selectively sampling only specific data points in k-space, a representation of spatial frequency information. Through undersampling, acquisition times in MRI can be shortened, which in turn can be traded into higher spatial resolution, spatial coverage, or shorter scanning sessions. The latter promise to reduce patient discomfort and anxiety, particularly for individuals who may experience claustropho-

bia or have difficulty lying still for extended periods [24]. Likewise, undersampling techniques can be employed to increase temporal resolution in dynamic imaging enabling real-time MRI [23,30], allowing for visualization and monitoring of dynamic processes as they occur. This is particularly relevant in interventional radiology procedures, where real-time guidance is, among others, crucial e.g. for accurate needle placement [16] or catheter navigation [7]. Without optimized reconstruction, however, reducing the number of acquired data points results in decreased image quality and artifacts like infolding, blurring, or distortions, which can hinder accurate diagnosis and interpretation. To tackle this issue, various techniques have been developed to enhance reconstructions and produce high-quality images from undersampled data, e.g., via compressed sensing [21] or parallel imaging [10,14].

Rather than focusing on enhancing the image quality for a predetermined undersampling pattern through image reconstruction, this paper addresses the challenge of identifying the optimal sampling pattern or k-space acquisition mask in terms of reconstruction quality or downstream task performance for a given undersampling ratio. Previously, in the area of deep learning this challenge has been mainly addressed by combined approaches that simultaneously learn a reconstruction network and an undersampling mask [1,31–34]. While image enhancement has proven to work well using such neural network approaches, these black box algorithms give little to no control over the reconstruction process and often have practical as well as legal limitations when being used as the basis of a reliable system in safety-critical applications. Furthermore, the application of large architectures in clinical practice, especially for real-time systems, requires computational resources not affordable by many healthcare institutions, particularly in developing regions where such advanced technology is scarce or cost-prohibitive.

Most closely related to our work, [27,28] propose a direct undersampling mask optimization scheme based on iterative gradient sampling (IGS), which repeatedly determines k-space elements that contribute the most to a loss

<sup>1</sup>Code: <https://github.com/saiboxx/bernoulli-mri>

criterion.

In contrast to previous work, we examine the potential of directly learning the mask itself via projected gradient descent. Our method called *ProM* is a fully differentiable, learning-based, and **probabilistic** framework for **mask** optimization. By framing the search for an optimal mask as a probabilistic optimization problem for a pre-specified acceleration factor, *ProM* is able to find the optimal undersampling distribution using ideas from relaxed categorical optimization in deep learning research. *ProM* optimizes the parameters of the undersampling distribution directly without the need for a model, resulting in runtimes of up to a few seconds. *ProM* does not operate under the assumption of any predetermined undersampling pattern. As such, it adapts to deliver bespoke results tailored to the specific requirements of the downstream task and the anatomical region as depicted in Figure 1. This adaptability makes *ProM* a versatile, data-driven tool, capable of serving as an all-purpose mask generator.

## 2. Related Work

The field of learning-based undersampled MRI reconstruction can largely be divided into two main branches: one focused on enhancing the quality of the reconstructed images, and the other on optimizing the initial sampling pattern itself.

**Image enhancement.** In the field of image quality enhancement, [13] consider the auto-regressive nature of encoder-decoder transformer architectures to derive missing spokes in radial undersampling. Using adversarial methods, [4] is able to enhance images with extreme acceleration factors. [8] apply the conditioning mechanism in diffusion models to retrieve enhanced images from low-quality reconstructions. Instead of conditioning on the actual image, [26] utilize the undersampled mask directly to guide the diffusion process. To mitigate domain shifts, [9] propose a diffusion prior trained via adversarial mapping. Through open-sourcing high-quality and large MRI datasets, the fastMRI [36] challenge actively promotes development in the area of enhancing undersampled MRI.

**Mask optimization and hybrid approaches.** While most approaches that optimize masks fall under the category of hybrid approaches (see below) as they also try to enhance the image quality at the same time, a few learning-based optimization approaches exist that exclusively focus on the sampling pattern. Iterative gradient sampling (IGS; [27,28]) is capable of efficiently deriving undersampling masks in 1D, but its complexity scales quadratically when moving to a 2D setting, whereas *ProM*'s complexity is only dependent on the number of optimization steps. [1] propose

LOUPE which encourages sparsity of the mask through  $\ell_1$ -penalization (without directly enforcing a specific acceleration factor) and approximates mask probabilities using a sigmoid output. As this approach includes a reconstruction network, it can also be seen as a hybrid approach. A different network-assisted idea is to employ a binarized mask on the forward pass and – in contrast to our proposal – transfer its gradients to a continuous real-valued mask on the backward pass [33,35]. An updated binary mask with induced sparsity is then produced on the subsequent iteration by thresholding and pruning the real-valued scores. With a focus on enforcing physical constraints, [19] propose optimized trajectories for compressed sensing. Similarly, [32] solve for k-space trajectories, but additionally use a reconstruction network to enhance image quality. [31] jointly optimize quadratic B-spline kernels together with a reconstruction network for multi-coil data.

**Other approaches.** Further approaches include variational models to jointly synthesize and reconstruct MRI images [6], sharpening networks [12] to counter the absence of high-frequency features in undersampled MRIs, or learnable Fourier interpolation [11]. To account for meta information such as the manufacturer, [20] condition the reconstruction network on side information.

## 3. Methods

In our presentation on learning a distribution for fully probabilistic undersampling masks, we first describe our routine for a single image and later extend this idea to jointly optimize masks across a whole dataset.

### 3.1. Probabilistic Undersampling Masks

In the following, we define  $D$  as the number of elements on the 2-dimensional or 3-dimensional k-space grid and use a vectorized notation for all objects for simplicity. Thus,  $\mathbf{x}_k = (x_k^{(1)}, \dots, x_k^{(D)})^\top \in \mathbb{C}^D$  depicts an image residing in k-space. Partial acquisition is augmented by applying a binary mask  $\mathbf{m} \in \{0, 1\}^D$  to the fully-sampled  $\mathbf{x}_k$  element-wise:  $\mathbf{x}_k \odot \mathbf{m}$ . For a given  $\mathbf{m}$ , the translation of the image from the under-sampled k-space to the complex image domain  $\mathcal{X} \subseteq \mathbb{C}^D$  is done by  $\mathbf{x}_c = \mathcal{F}^{-1}(\mathbf{x}_k \odot \mathbf{m})$ , where  $\mathcal{F}^{-1}$  is the inverse Fourier transform matrix. This procedure amounts to a simple zero-filling reconstruction strategy.

**Empirical Bayes Optimization.** Instead of following a distinct sampling pattern for  $\mathbf{m}$ , we assume that every element  $m^{(i)}, i = 1, \dots, D$ , in  $\mathbf{m}$  is the result of an independent Bernoulli experiment of a random variable  $M^{(i)}$  with distri-




Figure 1. Visualization of the *ProM* optimization routine for the ACDC dataset with an acceleration factor of  $\times 8$  for a 2D mask. The **bottom row** shows our Bernoulli mask distribution  $p_{M|x}$ , where a lighter color implies a higher probability for sampling the respective entry in the cartesian k-space grid. Starting from a randomly initialized distribution, *ProM* gradually optimizes  $p_{M|x}$  to maximize reconstruction quality while simultaneously increasing the sparsity of the masks. The resulting distribution converges to a domain-specific mask with desired acceleration factor that preserves most of the image’s quality (**top row**). The original image is displayed on the **right**.

bution defined by

$$p_M := \mathbb{P}(M = m \mid \theta) = \prod_{i=1}^D \theta_i^{m^{(i)}} (1 - \theta_i)^{1-m^{(i)}}, \quad (1)$$

$M = (M^{(1)}, \dots, M^{(D)})^\top$ , and  $\theta_i \in (0, 1)$  the sampling probability of the  $i$ th element. Eq. 1 can be thought of as a Bayesian prior for the image mask. While this factorized prior allows the method to be adaptable for any task and domain, other choices are possible and discussed in more detail in Section 5.

Given the prior distribution, we strive to optimize the posterior sampling distribution of masks (i.e., after accounting for the specific reconstruction task and data) for arbitrary differentiable loss functions  $\mathcal{L}$ , in particular those used in computer vision. As these loss functions are typically designed to only work in real-valued space, we transform  $x_c$  into a real-valued representation using its magnitude image  $\hat{x} = |x_c| \in \mathbb{R}^D$ . The quality of the reconstruction can then be assessed by  $\mathcal{L}(\hat{x}, x)$ , where  $x$  is the fully-sampled original image and  $\hat{x}$  the undersampled image as a function of  $m$  (or a random variable when  $M$  is not yet observed). Given a data point  $x$ , we take an empirical Bayes approach and directly optimize  $\theta$  by minimizing the expected loss via

$$\arg \min_{\theta} \mathbb{E}_{M \sim p_{M|x}} \mathcal{L}(\hat{x}, x) \approx \arg \min_{\theta} \frac{1}{L} \sum_{l=1}^L \mathcal{L}(\hat{x}^{(l)}, x), \quad (2)$$

where we approximate the expectation w.r.t. the posterior distribution  $p_{M|x}$  using  $L$  Monte-Carlo samples  $\hat{x}^{(l)}$  from  $p_{M|x}$ . Although in general, no analytical representation of the posterior exists, the optimization of Eq. 1 and  $\frac{1}{L} \sum_{l=1}^L \mathcal{L}(\hat{x}^{(l)}, x)$  individually is straightforward and allows us to derive an iterative optimization procedure (Algorithm 1) to optimize  $\theta$ .

**Constrained Optimization.** Without further constraints, the optimal solution of Eq. 2 is  $\theta = \mathbf{1}$ , i.e., the fully-sampled image with minimal loss, as  $\hat{x} \equiv x$ . To incorporate undersampling into our objective, we introduce a constraint similar to [38] by limiting the sum of all probabilities in  $p_{M|x}$  to a pre-specified value  $S$ , i.e.,  $\sum_{i=1}^D \theta_i \leq S$ . Practically speaking,  $S$  will result in the number of sampled k-space elements as  $\sum_{i=1}^D \theta_i$  is the expected value of  $\|m\|_0$ . Given a user-defined acceleration factor  $\alpha$ , i.e. the ratio between the amount of acquired k-space points in the full versus the undersampled image, we can alternatively define  $S = \lfloor \frac{D}{\alpha} \rfloor$  and our final objective as

$$\begin{aligned} & \arg \min_{\theta} \mathbb{E}_{M \sim p_{M|x}} \mathcal{L}(\hat{x}, x) \\ & \text{s.t. } \sum_{i=1}^D \theta_i \leq S \text{ and } S \in \{0, \dots, D\}. \end{aligned} \quad (3)$$

The constraint in Eq. 3, which can equally be expressed as an  $\ell_1$ -norm penalty for  $\theta$ , has an intrinsic affinity for sparse mask distributions (c.f. Figure 2).




Figure 2. Histograms of probabilities in  $\theta$  during different optimization phases for a sample of fastMRI Knee. Starting from a random initialization (leftmost distribution), the sparsification constraint leads to a posterior  $p_{M|x}$  with probabilities that tend to be close to zero or one.

### 3.2. Differentiability through Reparameterization

Except for the stochastic sampling of  $m$ , the proposed approach is fully differentiable including  $\mathcal{F}^{-1}$  and  $|\cdot|$ . In order to use modern autograd frameworks for stochastic masks, we apply the Gumbel-Softmax trick [15] tailored to

the Bernoulli distribution [38]. Let  $\rho := \log\left(\frac{\theta}{1-\theta}\right)$  be the log odds-ratio for  $\theta$ . Then, a “soft mask”  $\mathbf{m}_{\text{soft}}$ , allowing for differentiability, can be obtained by sampling

$$\begin{aligned} \mathbf{m}_{\text{soft}} &\sim \mathbb{1}(\rho + \mathbf{g}_1 - \mathbf{g}_0 \geq \mathbf{0}) \\ &\approx \sigma((\rho + \mathbf{g}_1 - \mathbf{g}_0)\tau^{-1}), \end{aligned} \quad (4)$$

where  $\mathbf{g}_1, \mathbf{g}_0$  are independent and identically distributed samples from a Gumbel(0,1) distribution and the indicator function  $\mathbb{1}$  is relaxed using a sigmoid function  $\sigma$ . A temperature parameter  $\tau$  controls the *softness* of the discrete approximation and is annealed during optimization. Here, stochasticity is rerouted over the Gumbel samples and thus a computational graph is able to propagate gradients to  $\theta$ . As the mask  $\mathbf{m}$  needs to be strictly binary, which is not the case for  $\mathbf{m}_{\text{soft}}$ , we adopt the straight-through Gumbel estimator trick [15], yielding

$$\mathbf{m} = \mathbb{1}(\mathbf{m}_{\text{soft}} \geq 0.5) - \text{sg}[\mathbf{m}_{\text{soft}}] + \mathbf{m}_{\text{soft}}, \quad (5)$$

where  $\mathbb{1}$  is applied element-wise and returns our final binary mask sample. *sg* denotes the *stop gradient* operation, which blocks gradients from backpropagation. In other words, Eq. 5 yields discrete values while we obtain gradients for its soft approximation.

### 3.3. Optimization via Projected Gradient Descent

The optimization problem in Eq. 3 cannot be solved effectively with standard gradient descent. Instead, we follow [38] and use a projected gradient approach by first updating the unconstrained parameter vector  $\tilde{\theta} = [\tilde{\theta}_1, \dots, \tilde{\theta}_D]$  using  $\tilde{\theta} = \theta - \eta \nabla_{\theta} \mathbb{E}_{\mathbf{M} \sim p_{\mathbf{M}}|\mathbf{x}} \mathcal{L}(\hat{\mathbf{x}}, \mathbf{x})$  with  $\eta$  being the learning rate, and then project  $\tilde{\theta}$  into the space of valid elements by solving

$$\sum_{i=1}^D \min(1, \max(0, \tilde{\theta}_i - \lambda)) = S \quad (6)$$

for  $\lambda \in \mathbb{R}$ , yielding

$$\theta = \min(\mathbf{1}, \max(\mathbf{0}, \tilde{\theta} - \max(0, \lambda)\mathbf{1})). \quad (7)$$

See Supplementary Material A.2 for an extended derivation of the objective. A solution to Eq. 6 can be obtained using a convex solver or root-finding method such as bisection search. To foster exploration at the beginning of the training and allow for exploitation at later stages, we anneal  $S$  during optimization. First, exploration iterations with  $S = D$  allow for unrestricted optimization. Then, an annealing phase following the schedule of [39] (see Supplementary Material A.1) decreases  $S$  to meet the desired acceleration factor. Finally, the exploitation phase optimizes  $\theta$  under the nominal constrained  $S$  (c.f. Figure 3).

Our procedure is summarized in Algorithm 1 for a single image  $\mathbf{x}_k$ . While our previous approach is presented




Figure 3. Sparsity versus reconstruction quality for a sample of fastMRI Knee. The amount of active masked elements in  $\mathbf{m}$  is bounded by the continuously annealed dense rate  $\frac{S}{D}$ . During optimization, the average SSIM metric (measuring reconstruction) plummets when  $\theta$  is limited by the constraint (at around 20% progress) but roughly stays constant or even slightly improves while the number of active elements is further restricted.

for only a single sample, it can be equivalently applied to a whole dataset by iteratively taking random batches instead of a single  $\mathbf{x}_k$  for the optimization of  $\theta$ . Note that the only trainable parameters in *ProM* are the  $D$  parameters  $\theta$ . Optimization can thus be done in a matter of seconds (single image) or a few minutes (full dataset).

---

#### Algorithm 1 Optimization routine of *ProM*

---

**Input:** k-space image  $\mathbf{x}_k$ , mask distribution  $p_{\mathbf{M}}$ , number of samples  $L$ , acc. factor  $\alpha$ , iterations  $i$ , learning rate  $\eta$ , criterion  $\mathcal{L}$

**for** 1, ...,  $i$  **do**

▷ Draw samples (cf. Eq. (4) and (5))

$\mathcal{M} \leftarrow \{\mathbf{m}_1, \dots, \mathbf{m}_L\}$

▷ Expand image to match shape of  $\mathcal{M}$

$\mathbf{X}_k \leftarrow \text{expand}(\mathbf{x}_k)$

▷ Compute undersampled magnitude images

$\hat{\mathbf{X}} \leftarrow |\mathcal{F}^{-1}(\mathbf{X}_k \odot \mathcal{M})|$

▷ Compute fully-sampled image and expand

$\mathbf{X} \leftarrow \text{expand}(|\mathcal{F}^{-1}(\mathbf{x}_k)|)$

▷ Compute batched loss and apply gradients

$\tilde{\theta} \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}(\hat{\mathbf{X}}, \mathbf{X})$

▷ Anneal to acc. factor over iterations

$S \leftarrow \text{anneal}(\alpha)$

▷ Project updated weights (cf. Eq. (6) and (7))

$\theta \leftarrow \text{project}(\tilde{\theta}, S)$

**end for**

---

### 3.4. Constraint Posterior Mode

Although realizations of the resulting distribution  $p_{\mathbf{M}|\mathbf{x}}$  of our approach are sparse in practice (cf. Figure 2) and adhere to the constraint, no real data set will result in a posterior with all  $\theta_i$  values being optimized to be either exactly




Figure 4. Optimized undersampling 2D mask distributions for ACDC, BraTS and fastMRI Knee (rows) with varying acceleration factors (columns). Different anatomic regions (right) have a distinct unique optimal distribution. The color indicates the value of each  $\theta_i$  representing a point in k-space. Bright color implies a high probability of acquiring the respective entry.

0 or 1. Whereas this result simply reflects the uncertainty in the data, some practical applications require masks to be deterministic. Therefore, we propose a deterministic variant of our approach by generating the mask  $\mathbf{m}$  using the mode of (the optimized) posterior  $p_{\mathbf{M}|\mathbf{x}}$  conditional on the constraint  $\|\mathbf{m}\|_0 = S$ . In other words, we find this mask  $\mathbf{m}^*$  as

$$\mathbf{m}^* = \arg \max_{\mathbf{m}, \|\mathbf{m}\|_0=S} p_{\mathbf{M}|\mathbf{x}}(\mathbf{m}). \quad (8)$$

This approach is equivalent to setting all mask entries  $m^{*(i)}$  corresponding to the  $S$  largest probabilities  $\theta_{(D-S+1)}, \dots, \theta_{(D)}$  to the value one and all others to zero.

## 4. Experiments

We investigate the performance of *ProM* using slices from ACDC [5], BraTS [2, 3, 22] and fastMRI Knee [36]:

**ACDC** are cardiac MRIs with 100 train and 50 test subjects. We extract the end-diastolic frame in 256px resolution including segmentation labels of the left and right ventricular cavity as well as the left ventricular myocardium, yielding 548 train and 338 test slices. k-space data is emulated via Fourier transform.

**BraTS** contains brain MRIs with T2-FLAIR, T1-, T1Gd- and T2-weighted modalities. The goal of the segmentation is to determine the classes of whole, core, and enhancing tumors. The dataset is split into 387 train and 97 test subjects. We extract 19,350 train and 4,850 test slices using slice indices between 60 and 110 with 256px resolution. k-space data is emulated as in ACDC.

**fastMRI Knee** includes raw k-space data of single coil knee MRI. To focus on pathologies, we extract the annotated subset of fastMRI+ [37] amounting to 8,057 train and 1524 test slices with a center-crop to 320px resolution.

For *ProM* we use 2500 iterations with a learning rate of 0.01 in the Adam [17] optimizer. From the full amount of iterations, we took 250 steps for exploration and exploitation each. This configuration provided an ideal balance between runtime and convergence. We use batches of size 32 and draw  $L = 4$  Bernoulli samples for each sample, yielding a total batch size of 128. Similar to [18], we found that a low number of Monte Carlo samples is sufficient if the batch size is large. The temperature  $\tau$  follows a linearly decreasing schedule from 1 to 0.03 in the last step, which aligns with [38]. For reconstruction, we use the mean squared error for  $\mathcal{L}$ . The mask distribution  $p_{\mathbf{M}}$  is initialized by drawing  $\theta_i$  from  $U(0, 1)$ . The final masks are obtained by applying Bayesian model averaging on 10 different *ProM* solutions.

Optimization is done in PyTorch v.1.13 [25] on an NVIDIA A100 GPU. We compare our approach against an equispaced mask with fully-sampled central region of 4% [36] and random offset, a 2D variable density Gaussian mask, and the learning-based IGS [28]) method. The reported metrics represent the mean across 10 randomly initialized runs per method.

### 4.1. Domain-specific Masks

Each anatomic region (as a composition of different elemental shapes) yields a distinct k-space representation. Our

first experiment not only demonstrates that this is the case in practice but also shows that a data-driven optimization routine for masks such as *ProM* is indeed necessary to facilitate optimal reconstruction (c.f. Figure 4 for the results of *ProM* trained on the three different datasets). For instance, the cardiac ACDC dataset consists dominantly of elliptic primitives, which results in a completely different optimal mask than fastMRI Knee with a lot of vertical lines and some horizontal elements in image space.

## 4.2. Convergence and Stability

*ProM* is a stochastic method, which involves randomness in its optimization process through the initialization of  $p_M$  and Monte Carlo samples  $\hat{x}^{(l)}$  on each forward pass or batch collection of data samples. Figure 5 displays the point-wise variance per mask parameter  $\theta_i$  over multiple independent runs. We observe the presence of a distinct central area that exhibits very small variance / high stability across all runs for most acceleration factors. Only in the case of an extreme acceleration factor ( $\times 32$ ), more variance is found within this central region. This suggests that particular segments within the k-space exhibit varying degrees of relevance, with the central region providing greater importance than the respective peripheral regions. Consequently, an optimal sampling strategy would entail precisely targeting this central area. If this region cannot be fully covered for larger acceleration factors ( $\alpha = 32$ ), the main source of uncertainty is located within this important area with no distinct best solution.

Apart from the convergence of the mask distribution,




Figure 5. Pointwise variance in cartesian k-space per mask parameter over 10 randomly initialized runs on the fastMRI Knee dataset shown for various  $\alpha$ . Brighter color indicates a larger variance.




Figure 6. Comparison of an original scan (**left**) with reconstructions based on the constrained mode (**middle**) and random samples (**right**) of the learned mask distribution for fastMRI Knee.




Figure 7. Pixel-variance on fastMRI Knee reconstructions based on random mask samples. Brighter color indicates larger variance.

sampling masks from  $p_{M|\alpha}$  induces stochasticity in the inference process. As shown in Figure 6, drawing random samples from  $p_{M|\alpha}$  produces coherent images, which only differ slightly in their high-frequency noise patterns. Interestingly, when investigating the variation of reconstructed pixels across multiple mask samples in Figure 7, a periodic pattern along the horizontal axis appears in the case of fastMRI Knee. The wavelength of the observed pattern is getting longer with higher  $\alpha$ . The stochasticity can be eliminated and reconstruction quality can be enhanced by using the suggested approach in Section 3.4 and generating the mask based on the conditional posterior mode.

## 4.3. Reconstruction Quality

To assess the reconstruction quality of *ProM* we evaluate the peak signal-to-noise ratio (PSNR), structural similarity index measure (SSIM), and normalized mean squared error (NMSE) for acceleration factors  $\alpha \in \{4, 8, 16, 32\}$  (abbreviated as  $\times 4$  to  $\times 32$ ) on the tests sets of fastMRI Knee, ACDC and BraTS. Results (Table 1) show that the IGS method works well, the additional flexibility of *ProM* to operate in 2D, however, allows to obtain superior results. For fastMRI Knee this advantage increases for higher  $\alpha$ . As an instance, in the case of factor  $\times 4$ , the SSIM / NMSE of *ProM* is 5.00% / 19.23% better compared to IGS, with the improvement increasing to 15.64% / 134.78% for factor  $\times 32$ . When compared to more conventional masking approaches, the gain in quality of *ProM* is even larger. In addition, Figure 8 shows that the *ProM* mask does not introduce noise nor artifacts, unlike the other evaluated methods. However, even with a custom-tailored undersampling strategy, high-frequency details are omitted to a certain degree.

When evaluating on the BraTS dataset, *ProM* is superior for acceleration factors  $\times 4$  and  $\times 8$ , but reveals performance limitations on higher  $\alpha$  values. Interestingly, for  $\alpha = 16$  *ProM* exhibits the highest PSNR among the compared methods but has a rather low SSIM. This can be reasoned by PSNR and SSIM capturing different aspects of image quality, where PSNR is primarily concerned with the amount of noise and SSIM focuses on structural information. Thus, this implies the presence of a low distortion level but differences in perceptual quality, as it is also visible in Figure 9.




Figure 8. Reconstructions of a fastMRI Knee sample with  $\alpha = 16$  (top) and  $\alpha = 32$  (bottom). Each cell contains an evaluated method, where the reconstructed scan is displayed on the left and the applied mask on the right side. The original scan is shown in the right column. Equidistant spacing and IGS display strong infolding artifacts, which are less pronounced with Gaussian sampling. The reconstruction with the data-driven mask of *ProM* is much closer to the fully-sampled image and reduces image noise as well as artifacts very clearly.

		x4			x8			x16			x32		
		PSNR $\uparrow$	SSIM $\uparrow$	NMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	NMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	NMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	NMSE $\downarrow$
fastMRI	Equi.	24.481	0.601	0.066	23.803	0.524	0.077	23.365	0.474	0.085	23.158	0.448	0.090
	Gauss.	29.570	0.664	0.027	28.118	0.560	0.035	23.342	0.440	0.112	17.373	0.299	0.396
	IGS	28.553	0.640	0.031	26.428	0.532	0.047	24.376	0.458	0.070	22.313	0.409	0.108
	<i>ProM</i>	<b>29.787</b>	<b>0.672</b>	<b>0.026</b>	<b>28.472</b>	<b>0.570</b>	<b>0.034</b>	<b>27.575</b>	<b>0.511</b>	<b>0.040</b>	<b>26.739</b>	<b>0.473</b>	<b>0.046</b>
ACDC	Equi.	23.566	0.716	0.069	22.552	0.674	0.0857	21.785	0.653	0.103	21.772	0.650	0.104
	Gauss.	39.169	0.959	<b>0.001</b>	24.681	0.705	0.073	18.038	0.512	0.265	15.786	0.428	0.422
	IGS	33.856	0.925	0.007	27.804	0.813	0.026	23.804	0.712	0.065	<b>20.705</b>	<b>0.621</b>	<b>0.132</b>
	<i>ProM</i>	<b>40.906</b>	<b>0.977</b>	<b>0.001</b>	<b>33.852</b>	<b>0.896</b>	<b>0.007</b>	<b>26.344</b>	<b>0.735</b>	<b>0.035</b>	20.401	0.580	0.140
BraTS	Equi.	26.732	0.716	0.033	25.599	0.687	0.043	25.660	0.693	0.042	<b>25.149</b>	<b>0.692</b>	<b>0.048</b>
	Gauss.	38.372	0.779	<b>0.002</b>	22.452	0.301	0.112	17.543	0.225	0.307	14.357	0.152	0.574
	IGS	35.550	0.911	0.004	30.607	0.821	0.014	26.995	<b>0.739</b>	0.032	24.054	0.679	0.061
	<i>ProM</i>	<b>39.623</b>	<b>0.943</b>	<b>0.002</b>	<b>35.660</b>	<b>0.874</b>	<b>0.004</b>	<b>27.002</b>	0.375	<b>0.031</b>	15.495	0.174	0.425

Table 1. Quality of reconstruction for the fastMRI Knee, ACDC, and BraTS.




Figure 9. *ProM* reconstructions of T2-weighted BraTS samples with  $\alpha = 16$  with high PSNR and low SSIM.

#### 4.4. Zero-shot Undersampled Segmentation

Visual reconstruction quality does not necessarily correlate with performance in downstream tasks such as segmentation [27]. We now investigate the quality of *ProM* in a transfer learning scenario, where we apply a frozen segmentation network trained on fully-sampled MRIs and learn a mask to maximize its performance with undersampled MRIs. For this, we choose a standard U-net [29] implementation and training routine with channel multipliers of (16, 32, 64, 128, 256). We substitute  $\mathcal{L}$  and use the trained segmentation network *net* as a proxy paired with a com-

bination of Dice and cross-entropy loss  $\mathcal{L}_{\text{seg}}$ , i.e.  $\mathcal{L}(\hat{x}, x)$  becomes  $\mathcal{L}_{\text{seg}}(\text{net}(\hat{x}), x_{\text{seg}})$  where  $x_{\text{seg}}$  is the segmentation label of  $x$ . Additionally, for the task of segmentation, we introduce a 1D variant of *ProM* to better understand its gain in performance compared to other methods.

The segmentation performance is evaluated in Table 2 via Dice score and Intersection-over-Union (IoU). Examples are visualized in Figure 10. The fully-sampled MRI achieves a macro-averaged Dice score of 0.855 and an IoU of 0.763 for ACDC as well as 0.772 and 0.710 for BraTS, respectively. While 1D *ProM* does not surpass the performance of IGS, our contention is that 1D *ProM* would benefit from optimized parameters specifically tailored for 1D masking. 2D *ProM* achieves competitive segmentation performance for  $\alpha = 8$  and is overall notably better for extreme acceleration factors such as  $\alpha = 32$  with a Dice score / IoU of 0.727 / 0.606 in ACDC and 0.706 / 0.608 in BraTS. The Gaussian mask appears to be an efficient and straightforward method to achieve competitive performance for smaller  $\alpha$ . Note that the *ProM* mask determined by  $\mathcal{L}_{\text{seg}}$  differs substantially from the one obtained to maximize reconstruction quality (Figure 4), emphasizing the fact that an optimal mask does not only vary with the dataset but also with the task.




Figure 10. Segmentation with a pre-trained U-net using  $\mathcal{L}_{\text{seg}}$  in  $\times 16$  accelerated *ProM*. The contours in ACDC correspond to the left (red) and right (green) ventricular cavity as well as the left ventricular myocardium (yellow). BraTS markings in the TIGd sample imply whole (green), core (red) and enhancing (yellow) tumor.

			Equi.	Gauss.	IGS	<i>ProM</i> (1D)	<i>ProM</i> (2D)	
ACDC	$\times 8$	Dice	↑	0.671	<b>0.847</b>	0.828	0.762	0.839
		IoU	↑	0.546	<b>0.752</b>	0.726	0.650	0.742
	$\times 16$	Dice	↑	0.645	0.644	0.745	0.717	<b>0.789</b>
		IoU	↑	0.517	0.534	0.627	0.599	<b>0.679</b>
	$\times 32$	Dice	↑	0.644	0.399	0.592	0.587	<b>0.727</b>
		IoU	↑	0.517	0.301	0.466	0.460	<b>0.606</b>
BraTS	$\times 8$	Dice	↑	0.596	0.733	0.716	0.646	<b>0.739</b>
		IoU	↑	0.489	0.638	0.619	0.542	<b>0.646</b>
	$\times 16$	Dice	↑	0.589	0.597	0.651	0.537	<b>0.735</b>
		IoU	↑	0.481	0.494	0.546	0.426	<b>0.640</b>
	$\times 32$	Dice	↑	0.580	0.315	0.537	0.483	<b>0.706</b>
		IoU	↑	0.472	0.226	0.428	0.374	<b>0.608</b>

Table 2. Segmentation metrics using a pre-trained U-net on ACDC and BraTS, showing that *ProM* (2D) excels especially for extreme acceleration factors.

## 5. Conclusion

This paper proposes *ProM* as a general framework and building block for data-driven and probabilistic undersampling mask learning, with theoretically guaranteed acceleration factor enforcement. Our evaluation protocol consisted in investigating image reconstruction and zero-shot segmentation of cardiac, brain, and knee MRIs. *ProM* shows promising performance, with increased benefits compared to other methods, especially in the downstream task of segmentation, enabling extreme acceleration factors, which can be traded into higher spatial resolution or coverage. Further, *ProM* again demonstrated the benefit of not only having domain-specific but also task-specific undersampling masks, as the efficacy for a task does not necessarily overlap with visually correct reconstructions.

**Clinical Practice.** MRI is an extremely flexible imaging modality, which involves optimization and multiple trade-offs, e.g., between total acquisition time (or temporal resolution in case of time-resolved imaging), spatial resolution and coverage, signal-to-noise ratio, and k-space data to be acquired. Generally, any approach to reduce the amount of

k-space data to be acquired while maintaining the quality of reconstructed images relaxes constraints on a multitude of protocol parameters and opens up possibilities for protocol optimization. Learning-based reconstruction and optimized k-space sampling for task-specific undersampling patterns as provided by *ProM* provides such an approach for better and faster MR data acquisition. This is, e.g., particularly useful for high-speed interventions or pathology localization. The minimalistic and model-free strategy of *ProM* allows us to derive these optimal masks within seconds from a single data sample. As this also does not require extensive computational infrastructure, our approach further reduces barriers to actual deployment.

Currently, our evaluations have primarily focused on a 2D Cartesian pattern, forming scanner trajectories along the depth axis. However, more scanner parameters and physical limitations need to be taken into account when considering deployment. An extension of *ProM* to work with radial or spiral acquisition trajectories could open a new area of investigation. Furthermore, the implications of deriving custom masks in a multi-coil setting need to be considered. Additionally, custom masks for specific scanner parameters, acquisition protocols, or manufacturers could be developed.

**Future Work.** Despite our take on not using a reconstruction network, *ProM* as an encapsulated layer with end-to-end differentiability can be integrated into any architecture of choice without major adaptations. Another possible direction to improve our approach is to integrate prior domain knowledge into  $p_M$ . While our factorized prior distribution assumption proved to work well in most experiments, it could be beneficial to incorporate further prior knowledge if data is scarce, i.e., for very large acceleration factors.

## Acknowledgments

This work has been partially funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of BERD@NFDI - grant number 460037581. The authors gratefully acknowledge LMU Klinikum for providing computing resources on their Clinical Open Research Engine (CORE).



## References

- [1] Cagla Deniz Bahadir, Adrian V Dalca, and Mert R Sabuncu. Learning-based optimization of the under-sampling pattern in MRI. In *Information Processing in Medical Imaging: 26th International Conference, IPMI*, pages 780–792, 2019. 1, 2
- [2] Spyridon Bakas, Hamed Akbari, Aristeidis Sotiras, Michel Bilello, Martin Rozycki, et al. Advancing the cancer genome atlas glioma mri collections with expert segmentation labels and radiomic features. *Scientific data*, 4(1):1–13, 2017. 5
- [3] Spyridon Bakas, Mauricio Reyes, Andras Jakab, Stefan Bauer, Markus Rempfler, et al. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv:1811.02629 [cs, stat]*, 2018. 5
- [4] Aleksandr Belov, Joël Stadelmann, Sergey Kastruyulin, and Dmitry V Dylov. Towards ultrafast mri via extreme k-space undersampling and superresolution. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 254–264, 2021. 2
- [5] Olivier Bernard, Alain Lalonde, Clement Zotti, Frederick Cervenansky, Xin Yang, Pheng-Ann Heng, Irem Cetin, Karim Lekadir, Oscar Camara, Miguel Angel Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE Transactions on Medical Imaging*, 37(11):2514–2525, 2018. 5
- [6] Wanyu Bian, Qingchao Zhang, Xiaojing Ye, and Yunmei Chen. A learnable variational model for joint multimodal MRI reconstruction and synthesis. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 354–364, 2022. 2
- [7] Adrienne E Campbell-Washburn, Mohammad A Tavallaei, Mihaela Pop, Elena K Grant, Henry Chubb, Kawal Rhode, and Graham A Wright. Real-time mri guidance of cardiac interventions. *Journal of Magnetic Resonance Imaging*, 46(4):935–950, 2017. 1
- [8] Hyungjin Chung and Jong Chul Ye. Score-based diffusion models for accelerated mri. *Medical Image Analysis*, 80:102479, 2022. 2
- [9] Salman UH Dar, Şaban Öztürk, Yilmaz Korkmaz, Gokberk Elmas, Muzaffer Özbey, et al. Adaptive diffusion priors for accelerated mri reconstruction. *arXiv:2207.05876 [cs, eess]*, 2022. 2
- [10] Anagha Deshmane, Vikas Gulani, Mark A Griswold, and Nicole Seiberlich. Parallel mr imaging. *Journal of Magnetic Resonance Imaging*, 36(1):55–72, 2012. 1
- [11] Qiaoqiao Ding and Xiaoqun Zhang. Mri reconstruction by completing under-sampled k-space data with learnable fourier interpolation. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 676–685, 2022. 2
- [12] Siyuan Dong, Eric Z Chen, Lin Zhao, Xiao Chen, Yikang Liu, Terrence Chen, and Shanhui Sun. Invertible sharpening network for mri reconstruction enhancement. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 582–592, 2022. 2
- [13] Chang Gao, Shu-Fu Shih, J Paul Finn, and Xiaodong Zhong. A projection-based k-space transformer network for under-sampled radial mri reconstruction with limited training subjects. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 726–736, 2022. 2
- [14] Robin M Heidemann, Özkan Özsarlak, Paul M Parizel, Johan Michiels, Berthold Kiefer, Vladimir Jellus, Mathias Müller, Felix Breuer, Martin Blaimer, Mark A Griswold, et al. A brief review of parallel magnetic resonance imaging. *European radiology*, 13:2323–2337, 2003. 1
- [15] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *5th International Conference on Learning Representations, ICLR*, 2017. 3, 4
- [16] Irving Kaplan, Nicklas E Oldenburg, Paul Meskell, Michael Blake, Paul Church, and Edward J Holupka. Real time mri-ultrasound image guided stereotactic prostate biopsy. *Magnetic resonance imaging*, 20(3):295–299, 2002. 1
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR*, 2015. 5
- [18] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR*, 2014. 5
- [19] Carole Lazarus, Pierre Weiss, Nicolas Chauffert, Franck Mauconduit, Loubna El Gueddari, Christophe Destrieux, Ilyess Zemmoura, Alexandre Vignaud, and Philippe Ciuciu. Sparkling: variable-density k-space filling curves for accelerated t2\*-weighted mri. *Magnetic resonance in medicine*, 81(6):3643–3661, 2019. 2
- [20] Xinwen Liu, Jing Wang, Cheng Peng, Shekhar S Chandra, Feng Liu, and S Kevin Zhou. Undersampled mri reconstruction with side information-guided normalisation. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 323–333, 2022. 2
- [21] Michael Lustig, David Donoho, and John M Pauly. Sparse MRI: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. 1
- [22] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Transactions on Medical Imaging*, 34(10):1993–2024, 2014. 5
- [23] Krishna S Nayak, Yongwan Lim, Adrienne E Campbell-Washburn, and Jennifer Steeden. Real-time magnetic resonance imaging. *Journal of Magnetic Resonance Imaging*, 55(1):81–99, 2022. 1
- [24] Murat Alp Oztek, Christina L Brunnuell, Michael N Hoff, Daniel J Boulter, Mahmud Mossa-Basha, Luke H Beauchamp, David L Haynor, and Xuan V Nguyen. Practical considerations for radiologists in implementing a patient-friendly mri experience. *Topics in Magnetic Resonance Imaging*, 29(4):181–186, 2020. 1
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An im-

- perative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. [5](#)
- [26] Cheng Peng, Pengfei Guo, S Kevin Zhou, Vishal M Patel, and Rama Chellappa. Towards performant and reliable undersampled mr reconstruction via diffusion model sampling. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 623–633, 2022. [2](#)
- [27] Artem Razumov, Oleg Y Rogov, and Dmitry V Dylov. Optimal MRI undersampling patterns for ultimate benefit of medical vision tasks. *arXiv:2108.04914 [cs, eess]*, 2021. [1](#), [2](#), [7](#)
- [28] Artem Razumov, Oleg Y Rogov, and Dmitry V Dylov. Optimal mri undersampling patterns for pathology localization. In *Medical Image Computing and Computer Assisted Intervention, MICCAI*, pages 768–779, 2022. [1](#), [2](#), [5](#)
- [29] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention, MICCAI*, pages 234–241, 2015. [7](#)
- [30] Martin Uecker, Shuo Zhang, Dirk Voit, Alexander Karaus, Klaus-Dietmar Merboldt, and Jens Frahm. Real-time mri at a resolution of 20 ms. *NMR in Biomedicine*, 23(8):986–994, 2010. [1](#)
- [31] Guanhua Wang, Tianrui Luo, Jon-Fredrik Nielsen, Douglas C Noll, and Jeffrey A Fessler. B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri. *IEEE Transactions on Medical Imaging*, 41(9):2318–2330, 2022. [1](#), [2](#)
- [32] Tomer Weiss, Ortal Senouf, Sanketh Vedula, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Pilot: Physics-informed learned optimized trajectories for accelerated mri. *arXiv:1909.05773*, 2019. [1](#), [2](#)
- [33] Tomer Weiss, Sanketh Vedula, Ortal Senouf, Oleg Michailovich, Michael Zibulevsky, and Alex Bronstein. Joint learning of cartesian under sampling andre construction for accelerated mri. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8653–8657. IEEE, 2020. [1](#), [2](#)
- [34] Jingfen Xie, Jian Zhang, Yongbing Zhang, and Xiangyang Ji. Puert: Probabilistic under-sampling and explicable reconstruction network for cs-mri. *IEEE Journal of Selected Topics in Signal Processing*, 16(4):737–749, 2022. [1](#)
- [35] Shengke Xue, Zhaowei Cheng, Guangxu Han, Chaoliang Sun, Ke Fang, et al. 2d probabilistic undersampling pattern optimization for mr image reconstruction. *Medical Image Analysis*, 77:102346, 2022. [2](#)
- [36] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al. fastmri: An open dataset and benchmarks for accelerated mri. *arXiv:1811.08839 [physics, stat]*, 2019. [2](#), [5](#)
- [37] Ruiyang Zhao, Burhaneddin Yaman, Yuxin Zhang, Russell Stewart, Austin Dixon, Florian Knoll, Zhengnan Huang, Yvonne W Lui, Michael S Hansen, and Matthew P Lungren. fastmri+: Clinical pathology annotations for knee and brain fully sampled multi-coil mri data. *arXiv:2109.03812 [physics]*, 2021. [5](#)
- [38] Xiao Zhou, Weizhong Zhang, Hang Xu, and Tong Zhang. Effective sparsification of neural networks with global sparsity constraint. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3599–3608, 2021. [3](#), [4](#), [5](#), [11](#)
- [39] Michael Zhu and Suyog Gupta. To prune, or not to prune: Exploring the efficacy of pruning for model compression. In *6th International Conference on Learning Representations, ICLR, Workshop Track Proceedings*, 2018. [4](#)