Adaptive Selection of Sampling-Reconstruction in Fourier Compressed Sensing

Seongmin Hong[1] , Jaehyeok Bae[1] , Jongho Lee[1,2]* , and
Se Young Chun[1,2,3]*

[1]Dept. of Electrical and Computer Engineering, [2]INMC, [3]IPAI
Seoul National University, Republic of Korea
{smhongok, wogur110, jonghoyi, sychun}@snu.ac.kr
Project page: https://smhongok.github.io/ada-sel.html

Abstract. Compressed sensing (CS) has emerged to overcome the inefficiency of Nyquist sampling. However, traditional optimization-based reconstruction is slow and may not yield a high-quality image in practice. Deep learning-based reconstruction has been a promising alternative to optimization-based reconstruction, outperforming it in accuracy and computation speed. Finding an efficient sampling method with deep learning-based reconstruction, especially for Fourier CS remains a challenge. Existing joint optimization of sampling-reconstruction works ($\mathcal{H}_1$) optimize the sampling mask but yield suboptimal results because it is not adaptive to each data point. Adaptive sampling ($\mathcal{H}_2$) has also disadvantages of difficult optimization and Pareto sub-optimality. Here, we propose a novel adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$) framework that selects the best sampling mask and reconstruction network for each input data. We provide theorems that our method has a lower infimum of the true risk compared to $\mathcal{H}_1$ and effectively solves the Pareto sub-optimality problem in sampling-reconstruction by using separate reconstruction networks for different sampling masks. To select the best sampling mask, we propose to quantify the high-frequency Bayesian uncertainty of the input, using a super-resolution space generation model. Our method outperforms joint optimization of sampling-reconstruction ($\mathcal{H}_1$) and adaptive sampling ($\mathcal{H}_2$) by achieving significant improvements on several Fourier CS problems.
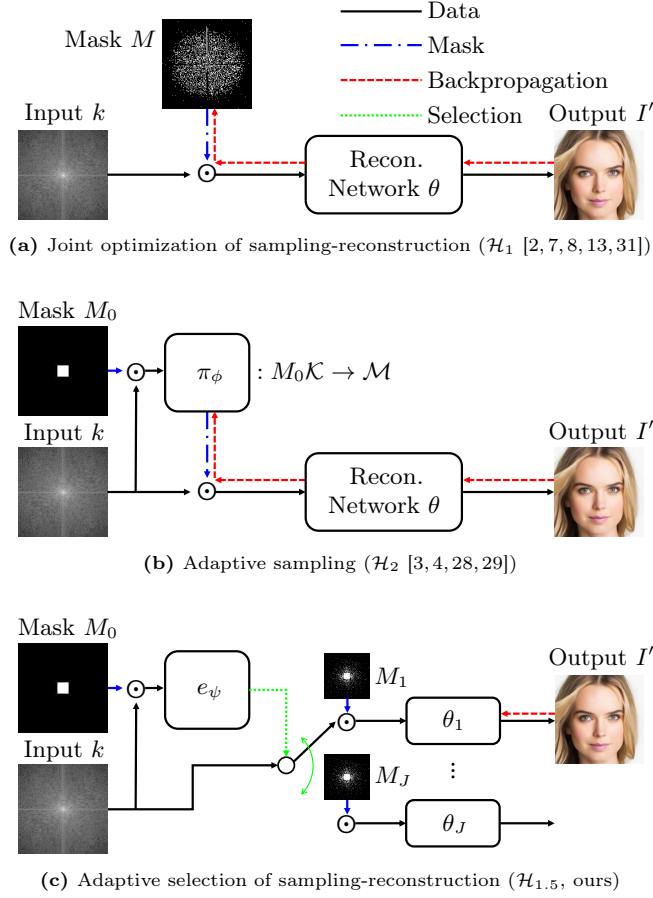
Keywords: Fourier compressed sensing · Sampling-reconstruction · Adaptive selection · Bayesian uncertainty

1 Introduction

Compressed sensing (CS) has revolutionized the field of image acquisition, enabling the reconstruction of high-quality images from a reduced number of measurements. This remarkable feat is achieved by exploiting the sparsity of natural images (or medical images) in certain transform domains. The CS theory [6, 10]

---

* Co-corresponding authors.

(a) Joint optimization of sampling-reconstruction ($\mathcal{H}_1$ [2, 7, 8, 13, 31])

(b) Adaptive sampling ($\mathcal{H}_2$ [3, 4, 28, 29])

(c) Adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$, ours)

Fig. 1: We propose the adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$, c). In Fourier compressed sensing, there were two classes of methods for finding the optimal sampling: joint optimization of sampling-reconstruction ($\mathcal{H}_1$, a) and adaptive sampling ($\mathcal{H}_2$, b). $\mathcal{H}_1$ has low potential as its mask M is not adaptive to each data point. $\mathcal{H}_2$ poses a challenge in optimizing the mask generator (π_ϕ) and then exhibits Pareto suboptimality, where a single θ is not optimal for multiple masks $M \in \mathcal{R}(\pi_\phi)$. In contrast, $\mathcal{H}_{1.5}$ is adaptive for input k (e_ψ in c selects the best M-θ pair), avoids the challenge of backpropagation to discrete space (red lines in a, b, c), and achieves Pareto optimality by dedicating each network θ_j exclusively to M_j. $\odot$ denotes the componentwise multiplication.

guarantees that an image can be accurately recovered from a non-adaptive random sampling pattern, with much fewer samples than the Nyquist-Shannon sampling theorem requires, if the image has a sparse representation in that domain. Before the era of deep learning, CS used to refer to obtaining the final image through l_1-regularized reconstruction, *i.e.*, solving Lasso [6]. Recently, reconstruction has

Table 1: Adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$) alleviates the drawbacks of joint optimization of sampling-reconstruction ($\mathcal{H}_1$) and adaptive sampling ($\mathcal{H}_2$).

Methods	Adaptive to input k	Backprop to a continuous space	Pareto optimal θ
$\mathcal{H}_1$ [2, 7, 8, 13, 26, 31]	✗	✗ (✓ [26])	✓
$\mathcal{H}_2$ [3, 4, 22, 28, 29]	✓	✗ (✓ [22])	✗
$\mathcal{H}_{1.5}$ (ours)	✓	✓	✓

often been performed using deep neural networks trained on the data. In this paper, we focus on deep learning-based reconstructions.

Fourier compressed sensing (Fourier CS) refers to CS where the measurement is in the discrete Fourier transform (DFT) of an image. As electromagnetic waves are inherently wave-like, obtaining spatial information such as pixel values directly is not feasible. Instead, spatial information is acquired through DFT. Obtaining samples for every Fourier-transformed element can be costly. Unlike the CS theory that suggests random sampling is sufficient, the sampling results of Fourier CS, similar to the prior in natural images, concentrate a significant amount of energy in the low-frequency (LF) components [18,27]. Despite deviating from the random sampling principle of CS theory, Fourier CS achieves excellent image quality in many domains by extensively sampling LF components; hence it has been successfully applied to various electromagnetic imaging applications, including magnetic resonance imaging (MRI) [18, 19] or radar [9, 11].

However, finding an optimal sampling method that is both efficient and effective remains a challenge in Fourier CS. One approach is the *joint optimization of sampling-reconstruction* (denoted by $\mathcal{H}_1$) [2,7,8,13,26,31], where the parameters of both the sampling and reconstruction networks are jointly trained using the dataset, as depicted in Fig. 1a. But this approach has two drawbacks, as described in Tab. 1. Firstly, it is not adaptive to each data point, *i.e.*, the optimized sampling and the reconstruction parameter would not be the best pair for a specific input. Moreover, the reconstruction network is usually trained by backpropagation. To employ backpropagation, the parameters are expected to be defined within continuous spaces. This makes training the sampling mask not trivial, as it is defined in a discrete space. Most methods [2, 7, 8, 13, 31] just perform discrete optimization anyway using the straight-through estimator [5, 31].

The other approach is *adaptive sampling* (denoted by $\mathcal{H}_2$) [3,4,20,25,28,29,34], which aims to generate the best sampling mask for each data point (or each image) based on the fact that a predetermined sampling mask may not be optimal for every situation. Most adaptive sampling studies generate the optimal sampling mask based on the information from the initially measured LF components of each data point, which has the potential to achieve excellent results. Then, they usually have a single reconstruction network that is responsible for many optimal masks for all data points. Unfortunately, there are a couple of major issues in

current adaptive sampling, as in Tab. 1. Similar to the joint optimization of sampling-reconstruction models, optimizing the mask generator is challenging due to the broad and discrete mask space, as depicted in Fig. 1b. Secondly, the single reconstruction network for diverse sampling masks may not be Pareto optimal, which we called *Pareto suboptimal reconstruction network*. Note that a similar issue can arise in the task of restoring various degradations (*e.g.*, the performance of a blind denoising network trained on multiple noise levels is usually lower than that of an identical network trained only on the specific noise level used as the actual input [32]).

In this paper, we propose a novel adaptive selection of the sampling-reconstruction framework for Fourier CS that alleviates the drawbacks of joint optimization of sampling-reconstruction and adaptive sampling. It is adaptive to each data point, avoids backpropagation to discrete spaces, and its reconstruction network is Pareto optimal. In the adaptive selection, we first sample LF components quickly and then leverage a super-resolution (SR) space generation model, to quantify the high-frequency (HF) Bayesian uncertainty. This approach ensures that HF components, which contain crucial details, are sampled more effectively, leading to improved reconstruction quality. The main contributions of our paper are as follows:

- Proposing a novel *adaptive selection of sampling-reconstruction* framework for Fourier CS that alleviates the drawbacks of joint optimization of sampling-reconstruction and adaptive sampling with theoretical justification.
- Designing the adaptive selection to efficiently quantify HF Bayesian uncertainty by leveraging an SR space generation model for determining sampling masks.
- Demonstrating that our adaptive selection improves performance in multiple Fourier CS problems such as facial image restoration (up to 0.04 average gain in SSIM) and multi-coil MR reconstruction (up to 0.004 average gain in SSIM).

2   Related Works

2.1   Fourier compressed sensing

Fourier CS can be defined as the following regression problem. Let us define the dataset $\mathcal{D} = \{(k^i, I^i)\}_{i=1}^{N}$ such that $k^1, k^2, \ldots, k^N \in \mathcal{K} \subseteq \mathbb{C}^L$ are fully-sampled k-space data and $I^1, I^2, \ldots, I^N \in \mathcal{I} \subseteq \mathbb{R}^L$ are the corresponding images, respectively. Let us define the mask space by $\mathcal{M} \subseteq \{0, 1\}^{L \times L}$ whose element is a diagonal binary matrix (indicating acquired (1) and unacquired (0) grid points). Let $h(k; M, \theta) : \mathcal{K} \times \mathcal{M} \times \Theta \to \mathcal{I}$ be a reconstruction function of k for a sampling mask $M \in \mathcal{M}$ and a reconstruction network (*e.g.*, U-Net [21] or E2E-VarNet [24]) parameterized by $\theta \in \Theta$. Then, this function h can be used as a joint sampling-reconstruction model that optimizes both the sampling mask M and the reconstruction network θ. Specifically, for the given dataset $\mathcal{D}$, the model is optimized to minimize the following empirical risk

$$\hat{\mathcal{L}}[h(\mathcal{K}; M, \theta)] = \frac{1}{N} \sum_{i=1}^{N} l(I^i, h(k^i; M, \theta)),$$

where l is the loss function $(e.g., l(I, \hat{I}) = 1 - \text{SSIM}(I, \hat{I}))$.

2.2 Joint optimization of sampling-reconstruction

One of the recent approaches to finding a good sampling mask is joint optimization of sampling-reconstruction $[2, 7, 8, 13, 26, 31]$, which reconstructs the image with a non-adaptive mask $M \in \mathcal{M}$. They are defined as follows:

$$\mathcal{H}_1 = \{h(\cdot; M, \theta) | M \in \mathcal{M}, \theta \in \Theta\}. \tag{1}$$

They jointly optimize M and θ for a *dataset*; however, M is not adaptive to each data point. Whether using a tailored M or not, the fundamental limitation of $\mathcal{H}_1$ is that the sampling mask is not optimal for each data point. Moreover, they exhibit highly varying results across different settings, as they require discrete optimization $[2, 7, 8, 13, 31]$ or virtual data $[26]$, as shown in Fig. 1a and Tab. 1.

2.3 Adaptive sampling

Some recent works $[3, 4, 20, 28, 28, 29]$ employ adaptive mask, using a mask generator $\pi_\phi : M_0 \mathcal{K} \to \mathcal{M}$, parameterized by $\phi \in \Phi$, as shown in Fig. 1b. Here, $M_0 \in \mathcal{M}$ denotes a mask that samples only LF components. Adaptive sampling approaches minimize $\hat{\mathcal{L}}[h]$ on

$$\mathcal{H}_2 = \{h(\cdot; \pi_\phi(\cdot), \theta) | \pi_\phi : M_0 \mathcal{K} \to \mathcal{M}, \theta \in \Theta, \phi \in \Phi\}. \tag{2}$$

Obviously, $\mathcal{H}_1 \subseteq \mathcal{H}_2$. That is, $\mathcal{H}_2$ has the greatest potential but is hard to train because of its complexity. Specifically, $\mathcal{H}_2$ faces two main issues: the difficulty of the mask generator (π_ϕ) optimization, and Pareto suboptimality of θ, due to the fact that a single reconstruction network is responsible for multiple masks.

Previous studies on $\mathcal{H}_2$ have used reinforcement learning $[3, 20, 28]$ or backpropagation $[3, 4, 28, 29]$ using straight-through estimator $[5, 31]$ to optimize mask generator π_ϕ, but this is a complicated problem because the action space $\mathcal{M}$ is too broad and discrete. θ is Pareto suboptimal in adaptive sampling studies since there are multiple sampling masks M while only one θ exists at inference time. Due to these difficulties, most adaptive sampling studies in CS-MRI have been conducted in a clinically less relevant simple setting of single-coil $[3, 20, 25, 28, 29, 34]$. There was only one study conducted on realistic multi-coil setting $[4]$, but the final models of $[4]$ turned out to be non-adaptive, which is an unintended consequence.

To avoid optimization in discrete space, $[22]$ proposed adaptive sampling using a conditional Generative Adversarial Network (cGAN). During sampling, cGAN assesses the uncertainty of samples yet to be acquired. Subsequently, the user selects a sample with the highest uncertainty for acquisition $(i.e.,$ greedy algorithm). This process of quantifying and sampling is iteratively repeated. While this method benefits from backpropagation occurring only in continuous space (Tab. 1) it still faces the challenge of a single reconstruction network, having to perform reconstruction for all masks. Consequently, CS-MRI experiments were conducted using a simple single-coil setup.

2.4 Super-resolution space generation

Super-resolution (SR) space generation [16, 17] aims to create diverse high-resolution (HR) images that can be downsampled to the same low-resolution (LR) image (*i.e.*, $q_\psi(I_{\mathrm{HR}}|I_{\mathrm{LR}}) \notin \{\delta_I | I \in \mathcal{I}\}$). For this purpose, a stochastic approach is used rather than a deterministic one. Conditional normalizing flow-based SR space generation methods [12, 15, 23] explicitly obtain $q_\psi(I_{\mathrm{HR}}|I_{\mathrm{LR}})$ using a diffeomorphic mapping $f_\psi : \mathcal{I} \to \mathcal{Z}$ and a simple base distribution q_z (*e.g.*, standard Gaussian), as $q_\psi(I_{\mathrm{HR}}|I_{\mathrm{LR}}) = q_z(f_\psi(I_{\mathrm{HR}}; I_{\mathrm{LR}}))|\det \frac{\partial f_\psi}{\partial I_{\mathrm{HR}}}(I_{\mathrm{HR}}; I_{\mathrm{LR}})|$. Since f_ψ is invertible, q_z and f_ψ^{-1} can be used to directly sample I_{HR} from q_ψ (*i.e.*, $z \sim q_z \implies f_\psi^{-1}(z; I_{\mathrm{LR}}) \sim q_\psi(\cdot|I_{\mathrm{LR}})$). In this work, we trained and exploited a recent robust flow-based SR space generation method [23], with tuned hyperparameters [12] for stability, to generate HR images from the corresponding LR image that is reconstructed from undersampled k-space data with mask M_0.

3   Proposed methods

In Sec. 2.2 and 2.3, we investigated the difficulty of optimizing the mask generator π_ϕ, and Pareto optimal θ (for all masks). This section proposes a novel scheme, *adaptive selection of sampling-reconstruction*, which does not encounter these problems. Using two Theorems 1 and 2, we explain our adaptive selection ($\mathcal{H}_{1.5}$) is better than the joint optimization ($\mathcal{H}_1$) and the adaptive sampling ($\mathcal{H}_2$).

3.1 Adaptive selection of sampling-reconstruction

Our adaptive selection model $\mathcal{H}_{1.5}$ is defined as follows:

$$\mathcal{H}_{1.5} = \left\{ \sum_{j=1}^{J} e_\psi(\cdot)_j h(\cdot; M_j, \theta_j) \middle| e_\psi : M_0\mathcal{K} \to \{e_j\}_{j=1}^{J}, M_j \in \mathcal{M}, \theta_j \in \Theta, \forall j \right\} \quad (3)$$

where e_j is the j-th standard unit vector (*i.e.*, one-hot vector). Each submodel $h(\cdot; M_j, \theta_j)$ contains mask M_j and reconstruction network θ_j as a pair, which is Pareto optimal. At inference time, each data selects an appropriate submodel through the mask selector $e_\psi(\cdot)_j$, which takes input M_0k. This scheme is similar to a segmented regression problem that ensembles multiple submodels using one-hot encoding.

Remark 1. If $\mathcal{H}_1$ is a linear regression, then $\mathcal{H}_{1.5}$ is a segmented linear regression.

We propose the following Theorems 1 and 2. Theorem 1 shows that $\mathcal{H}_{1.5}$ is better than $\mathcal{H}_1$ due to its adaptivity, and Theorem 2 demonstrates that $\mathcal{H}_{1.5}$ is superior to $\mathcal{H}_2$ because $\mathcal{H}_2$ has poor Pareto optimality.

Theorem 1 (Adaptive selection is better than non-adaptive). *For a true risk* $\mathcal{L}$, $\inf_{h \in \mathcal{H}_{1.5}} \mathcal{L}[h] \leq \inf_{h \in \mathcal{H}_1} \mathcal{L}[h]$.

Theorem 2 (Adaptive selection is Pareto optimal). *For a true risk $\mathcal{L}$,* $|\pi_\phi(M_0\mathcal{K})| \leq J \Rightarrow \inf_{h\in\mathcal{H}_{1.5}} \mathcal{L}[h] \leq \inf_{h\in\mathcal{H}_2} \mathcal{L}[h]$.

Please see the supplementary material for the proofs. Theorem 2 requires an assumption that optimizing π_ϕ is difficult (*i.e.*, $|\pi_\phi(M_0\mathcal{K})| \leq J$), which is justified in Section 2.2 (*e.g.*, The final model in [4] converged to $|\pi_\phi(M_0\mathcal{K})| \to 1$, which means non-adaptive). Theorems 1 and 2 suggest that the proposed adaptive selection scheme ($\mathcal{H}_{1.5}$) may outperform both non-adaptive methods ($\mathcal{H}_1$) and adaptive sampling ($\mathcal{H}_2$). In Section 3.2, we describe the implementation of the scheme using the HF Bayesian uncertainty quantified by an SR space generation method [23].

3.2   How to and what to adaptively select?

(How to) proposed mask selector e_ψ: The sample variance of a generative model to produce diverse samples can be utilized to quantify uncertainty for adaptive sampling [22]. Specifically in Fourier CS, inspired by the idea of initially sampling LF components, we employ an SR space generation model [12, 23] as a HF uncertainty quantifier.

The sample variance $v(M_0 k) := (\widehat{\mathrm{Var}}_{q_\psi}[k'_s])_{s=1}^S$ is an estimator of the mean square error in k-space domain, where S is the number of the SR samples and k'_s is the Fourier transform of the s-th sample. We make up the mask selector e_ψ using $v(M_0 k)$. Specifically, at train time, we normalize $v(M_0 k)$ so that $u(v) := v/\|v\|_2$ and then use the k-means++ clustering algorithm [1] to $\{u(v(M_0 k^i))\}_{i=1}^N$ to create centroids $(c_j)_{j=1}^J$. At inference time, we select adaptive mask index j by calculating the distance $u(v(M_0 k))$ and $(c_j)_{j=1}^J$.

We also need to determine the number of the sampling-reconstruction pairs J. Thinking of Remark 1, increasing J doesn't always mean better average performance; while increasing J can help in robustly handling outliers. This trade-off can be organized as Remark 2:

Remark 2 (Trade-off with the number of segments J). As J increases, despite more training resources, the average performance reaches a plateau at some point, but it becomes more robust against outliers.

The choice of J depends on the user's needs; we defaulted to $J = 3$. We delve into and validate Remark 2 in Sec. 5.

(What to) constructed sampling-reconstruction pairs $(M_j, \theta_j)_{j=1}^J$: One might try to create M_j from c_j using just sorting, *i.e.*,

$$M_j = \arg\max_M \|M c_j\|, \tag{4}$$

based on the following proposition.

Proposition 1 (simplified version). *With mild assumptions, the sorted sample variance (4) is the PSNR-maximizing mask.*
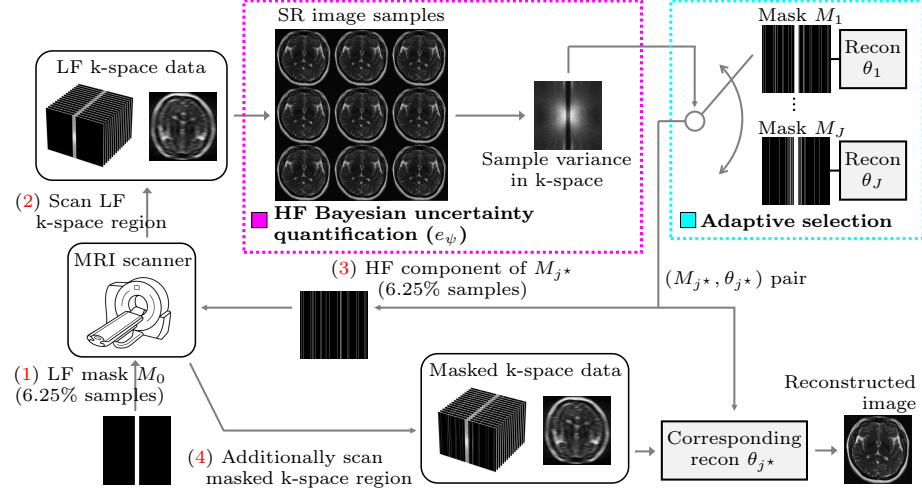
Fig. 2: Using the sample variance of SR space generation [12,23] results, we can quantify HF Bayesian uncertainty (highlighted in magenta dotted box). Then, we can adaptively select a sampling-reconstruction $(M - \theta)$ pair (highlighted in cyan dotted box). Here, we illustrate how our adaptive selection of sampling-reconstruction (Algorithm 2) is employed in CS-MRI. In (1)&(2) in the figure, MRI scanner scans LF k-space region and adaptively selects mask $(M_{j^\star}, \theta_{j^\star})$ pair using HF Bayesian uncertainty quantification. (3)&(4) After additionally scanned masked k-space region from $M_{j^\star}$, reconstructed images are generated from masked k-space data using $\theta_{j^\star}$. See Algorithm 2 for details.

Algorithm 1 Training

Input: Training set $\{k^i\}_{i=1}^N$, initial sampling mask $M_0 \in \mathcal{M}$, trained SR space generation model $f_\psi : \mathcal{I} \to \mathcal{Z}$, the number of segments J, the number of SR generated images S, the number of total sampling points N_M, and the empirical risk $\hat{\mathcal{L}}$.

Output: Masks $(M_j)_{j=1}^J$, reconstruction parameters $(\theta_j)_{j=1}^J$, and centroids of uncertainty $(c_j)_{j=1}^J$

  for $i = 1$ to N **do**
    for $s = 1$ to S **do**
      Sample $z^s \sim \mathcal{N}(0, \sigma_s^2)$
    $m^i \leftarrow \frac{1}{S} \sum_{s'=1}^S f_\psi^{-1}(z^{s'}; M_0 k^i)$
    $v^i \leftarrow \frac{1}{S-1} \sum_{s=1}^S (f_\psi^{-1}(z^s; M_0 k^i) - m^i)^{\circ 2}$
    $u^i \leftarrow v^i/\|v^i\|_2$ ▷ Normalized $\widehat{\mathrm{Var}}$
  $(c_j)_{j=1}^J \leftarrow k\text{-}means\text{++}(\{u^i\}_{i=1}^N, J)$ ▷ k-means++
  for $j = 1$ to J **do**
    $M_j \leftarrow M_0 + RejectionSampling(c_j, N_M - \mathrm{Tr}(M_0))$
    Train θ_j to minimize $\hat{\mathcal{L}}[h(\cdot; M_j, \theta_j)]$

Algorithm 2 Inference

Input: k-space input k, initial sampling mask $M_0 \in \mathcal{M}$, trained SR space generation model $f_\psi : \mathcal{I} \to \mathcal{Z}$, the number of segments J, masks $(M_j)_{j=1}^J$, reconstruction network parameters $(\theta_j)_{j=1}^J$, and the centroids of uncertainty $(c_j)_{j=1}^J$

Output: Reconstructed image I'

  for $s = 1$ to S **do**
    Sample $z^s \sim \mathcal{N}(0, \sigma_s^2)$
  $m \leftarrow \frac{1}{S} \sum_{s'=1}^S f_\psi^{-1}(z^{s'}; M_0 k)$
  $v \leftarrow \frac{1}{S-1} \sum_{s=1}^S (f_\psi^{-1}(z^s; M_0 k) - m)^{\circ 2}$
  $u \leftarrow v/\|v\|_2$ ▷ Normalized $\widehat{\mathrm{Var}}$
  $j^\star \leftarrow \arg\min_j \|u - c_j\|_2$ ▷ Selection
  $I' \leftarrow h(k; M_{j^\star}, \theta_{j^\star})$ ▷ Recon.

However, this method is not the optimal approach for maximizing SSIM. In general, it is known that introducing randomness to the mask is effective in maximizing SSIM [27]. Therefore, we generate M_j using rejection sampling proportional to c_j. Then, we train the dedicated θ_j for the corresponding M_j. Figure 2 shows the overview of adaptive selection, clearly showing why the proposed adaptive selection is *adaptive*. Algorithms 1 and 2 provide detailed descriptions of the training and inference processes of our adaptive selection method, respectively.

4   Experiments

We proposed the *adaptive selection of sampling-reconstruction* scheme in Sec. 3 to address the issues of $\mathcal{H}_1$ and $\mathcal{H}_2$. This section experimentally demonstrates that the proposed method performs well in various settings of Fourier CS.
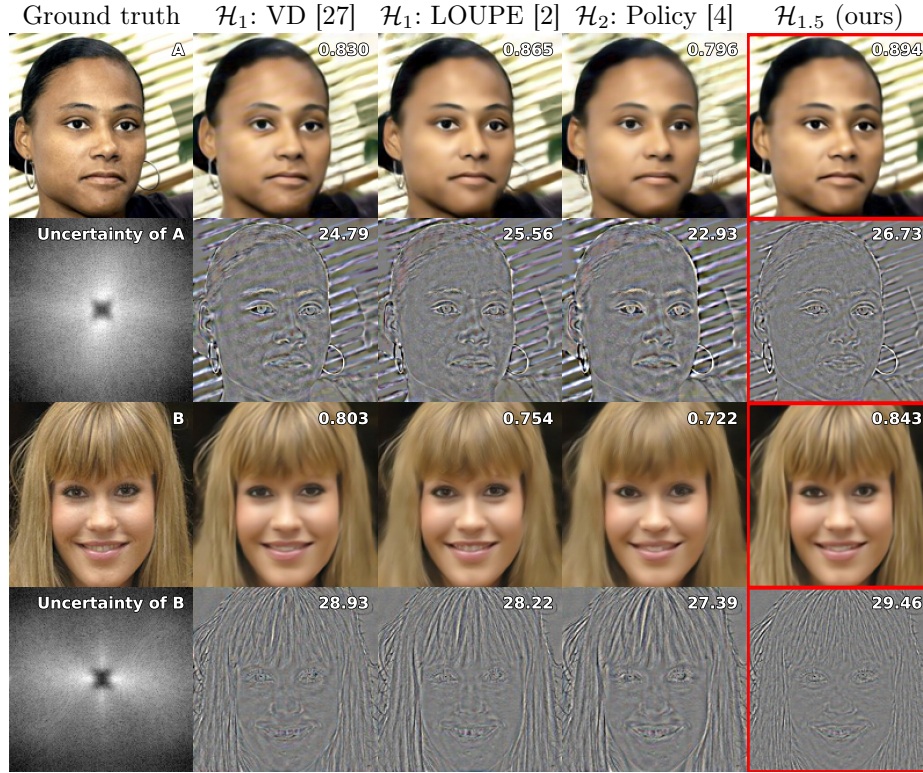
4.1   Fourier CS face reconstruction

We performed Fourier compressed sensing on the CelebA dataset [14], which consists of 160×160 RGB human face images. Similar to LOUPE [2], the reconstruction network used a U-Net [21] architecture, with 6 input channels and 3 output channels, because the input, zero-filling reconstruction, is complex.

Figure 3 shows a qualitative comparison of our method ($\mathcal{H}_{1.5}$) and other $\mathcal{H}_1$, $\mathcal{H}_2$ methods [2, 4, 27] at the acceleration rate $16\times$ with the metrics of SSIM and PSNR. Our final reconstruction results are superior to those of other methods, which supports Theorems 1 and 2. In detail, our method adeptly selects sampling-reconstruction pairs using HF Bayesian uncertainty. Looking at the first column of Fig. 3, in the case of A, the presence of horizontal stripes in the background results in a high uncertainty in the vertical direction, whereas in B, the elongated blonde hair leads to a high uncertainty in the horizontal HF components. In Fig. 4, which shows the sampling masks (with the selection) and the corresponding reconstruction results of $\mathcal{H}_{1.5}$, our e_ψ selected M_2, which has a shape similar to the uncertainty of A in the second column, obtained the highest SSIM for A. M_2 emphasizes in red in the error map, indicating effective suppression of artifacts caused by horizontal high-frequency components in the background, achieved by sampling more in the vertical direction. Similarly, M_3 in the third column, which has a shape similar to the uncertainty of B, achieved the highest SSIM for B. M_3 is highlighted in the error map, revealing reduced errors in the hair region of the subject due to increased sampling in the horizontal direction.

4.2   Multi-coil CS-MRI reconstruction

We also performed Fourier compressed sensing on the fastMRI multi-coil brain dataset [30]. We resized all slices to a size of 320×320. The number of coils was 16. Most implementations of $\mathcal{H}_1$ and $\mathcal{H}_2$ methods were based on the official

Ground truth	$\mathcal{H}_1$: VD [27]	$\mathcal{H}_1$: LOUPE [2]	$\mathcal{H}_2$: Policy [4]	$\mathcal{H}_{1.5}$ (ours)

Fig. 3: Our adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$) shows the strongest reconstruction performance (emphasized in red). Here, we show a qualitative comparison of reconstruction and error map at acceleration rate 16× in the CelebA dataset [14]. For comparison, we also show the results of the variable density (VD) [27], LOUPE [2], and policy-based adaptive sampling [4]. SSIMs and PSNRs are included in the reconstructions and the error maps, respectively.

fastMRI repository[1]. We experimented not only with 2D undersampling patterns, as described in Sec. 5, but also with 1D line subsampling used in actual MRI. For the latter, we modified the SR space generation model [23] to achieve a 16× SR only in the horizontal direction.

Figure 5 shows a qualitative comparison of our method ($\mathcal{H}_{1.5}$) and other $\mathcal{H}_1$, $\mathcal{H}_2$ methods [2, 4, 27] at 4× 1D undersampling with the metrics of SSIM and PSNR. Our final reconstructions outperform other methods, supporting Theorems 1 and 2. In the right half of Fig. 4, our e_ψ selected (M_1, θ_1), which samples more of the low-frequency components, obtained the best reconstruction result for A (second row). Besides, for input B (fourth row), (M_3, θ_3) generated the best reconstruction result (highlighted in red). Since M_3 samples the high-frequency components more, it made the clearest imaging of the longitudinal

---

[1] https://github.com/facebookresearch/fastMRI
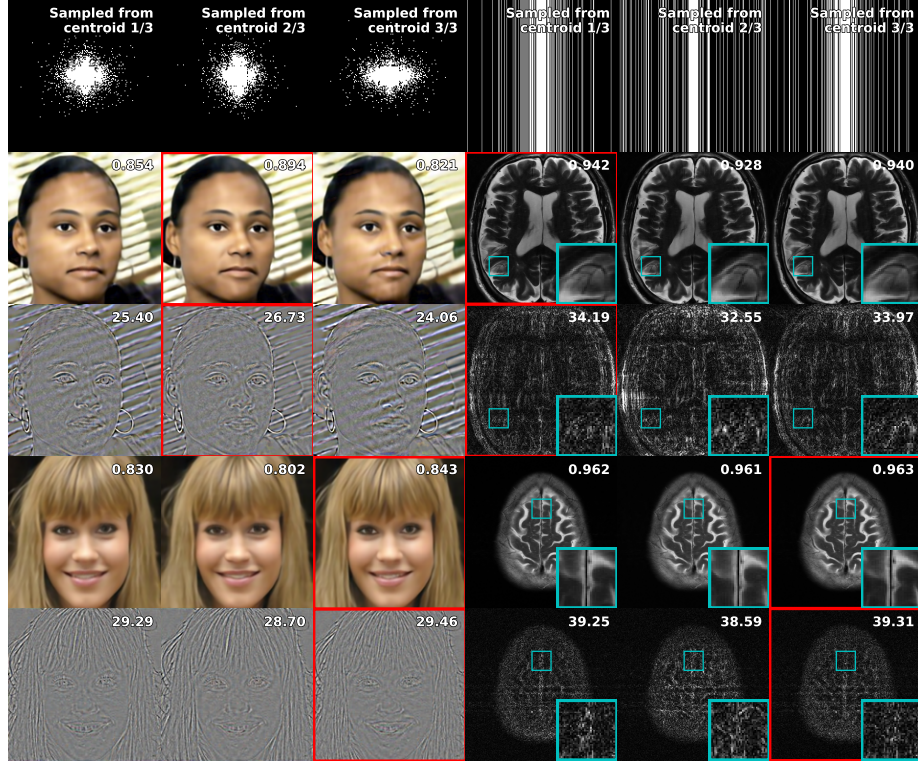
Fig. 4: Our adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$) adaptively selects the best sampling-reconstruction pair based on the HF uncertainty of the input, leading to strong reconstruction performance. Here, we show a qualitative comparison of reconstruction and error map obtained using the mask-reconstruction pairs $((M_j, \theta_j)_{j=1}^3)$ generated from Algorithm 1 at acceleration rate $16\times$ in the CelebA dataset [14] and at acceleration rate $4\times$ in the fastMRI dataset [30]. Our Algorithm 2 estimated the uncertainty of each image as in Figs. 3 and 5, and then selected the appropriate mask M_j (with θ_j) as emphasized in red. For all images in this case, the selected (M_j, θ_j) resulted in the best reconstruction outcomes. SSIMs and PSNRs are included in the reconstructions and the error maps, respectively.

fissure, indicating the effectiveness of the proposed method (*i.e.*, Algorithms 1 and 2).

In Tab. 2, we present the average SSIM of the proposed scheme in various accelerations and datasets. For comparison, LOUPE [2] and policy-based method [4] are evaluated. Two non-adaptive methods, uniformly random mask and sampling from VD [27], are also evaluated. For a 1D line sampling CS-MRI, equispaced masks are additionally evaluated. As shown in Tab. 2, our adaptive selection approach consistently achieves higher SSIM compared to other methods in all scenarios. For example in CelebA dataset at acceleration rate $8\times$, SSIM of our method (0.9405) is about 0.04 higher than the best of $\mathcal{H}_1$ (0.9073) and
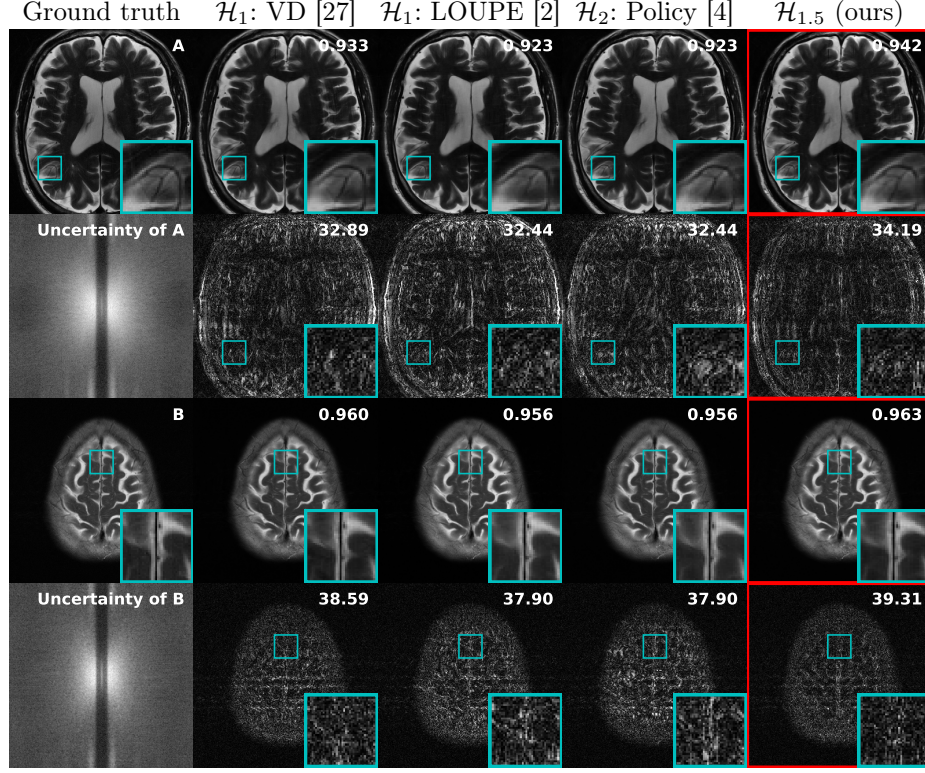
Fig. 5: In a practical multi-coil CS-MRI 1D line sampling scenario, our adaptive selection of sampling-reconstruction shows the highest SSIM (highlighted in red). Here, we show a qualitative comparison of reconstruction and error map at acceleration rate $4\times$ in the fastMRI dataset [30]. For comparison, we also show the results of the variable density (VD) [27], LOUPE [2], and policy-based adaptive sampling [4]. SSIMs and PSNRs are included in the reconstructions and the error maps, respectively.

$\mathcal{H}_2$ (0.8501). In addition in a realistic setting, CS-MRI 1D at acceleration rate $8\times$, SSIM of our method (0.9407) is about 0.004 higher than the best of $\mathcal{H}_1$ (0.9367) and $\mathcal{H}_2$ (0.9240), which is a significant difference in MRI reconstruction problem [33].

5   Discussion

Does the SR space generation model quantify the HF uncertainty well? Since sample variance estimates MSE, evaluating SR space generation can be done by sorting the sample variance, as in Proposition 1. After adaptive sampling and zero-filling for reconstruction, PSNR can be used as a metric for assessment. We qualitatively (in the supplementary material) and quantitatively (in Tab. 3) compare the adaptive sampling results in MRI [30] at acceleration

Table 2: Our adaptive selection of sampling-reconstruction ($\mathcal{H}_{1.5}$, Algorithms 1 and 2) shows the highestSSIM in Fourier CS in various settings (CelebA dataset [14] w/ 2D sampling, fastMRI multi-coil dataset [30] w/ 1D, 2D sampling).

SSIM↑		CelebA		CS-MRI 2D		CS-MRI 1D	
Method \ Accel.		8×	16×	4×	8×	4×	8×
$\mathcal{H}_1$	Random	0.8378	0.8684	0.9663	0.9506	0.9533	0.9255
	VD [27]	0.9073	0.8734	0.9698	0.9578	0.9603	0.9367
	LOUPE [2]	0.8742	0.8673	0.9671	0.9525	0.9541	0.9218
	Equispace (1D)	-	-	-	-	0.9603	0.9258
$\mathcal{H}_2$	Policy [4]	0.8501	0.8394	0.9698	0.9572	0.9569	0.9240
$\mathcal{H}_{1.5}$	Adaptive selection (ours)	0.9405	0.8952	0.9704	0.9585	0.9624	0.9407

Table 3: Quantitative comparison of PSNR and SSIM with *zero-filling*. 'Sorted-Self' achieved the highest PSNR, which supports our Proposition 1. Thus, we can assert that the SR space generation model effectively quantifies the HF uncertainty.

PSNR / SSIM	Sorted-Self	Sorted-Another	VD [27]
4×	**37.15 / 0.939**	36.36 / 0.922	33.33 / 0.854
8×	**34.79 / 0.910**	34.24 / 0.894	32.16 / 0.834

rate 8×. Table 3 presents the average PSNR and SSIM of the proposed methods on the validation dataset. 'Sorted-Self' refers to the zero-filling reconstruction results obtained when sorting its own HF Bayesian uncertainty to create a mask, while 'Sorted-Another' randomly shuffles the masks among the data points. We additionally generate a mask sampled from VD [27] for comparison. As a result, the 'Sorted-Self' approach consistently achieves the highest PSNR and SSIM. These results show that the SR space generation effectively quantifies HF Bayesian uncertainty.

Effect of the number of segments J Here, we validate Remark 2 by conducting an ablation study using the dataset employed in our experiments. Figure 6 shows the average of the lowest 5%, 10%, and 100% of SSIM values ($\overline{\text{SSIM}}_{5\%}^{\text{low}}$, $\overline{\text{SSIM}}_{10\%}^{\text{low}}$, and $\overline{\text{SSIM}}$) for all $J = 1, \ldots, 4$ in the CS-MRI 8× experiments (2D and 1D). The fact that the increase in $\overline{\text{SSIM}}$ is less noticeable when transitioning from $J = 2$ to $J = 3$ or $J = 4$ compared to the transition from $J = 1$ to $J = 2$ supports the first part of Remark 2, "As J increases, despite more training resources, the average performance reaches a plateau at some point." Additionally, the relatively significant increase in $\overline{\text{SSIM}}_{10\%}^{\text{low}}$ and $\overline{\text{SSIM}}_{5\%}^{\text{low}}$ when transitioning from $J = 2$ to $J = 3$ or from $J = 3$ to $J = 4$ supports the latter part of Remark 2, "As J increases, it becomes more robust against outliers." Therefore, users can choose J considering this trade-off.

(a) CS-MRI 2D 8×   **(b)** CS-MRI 1D 8×
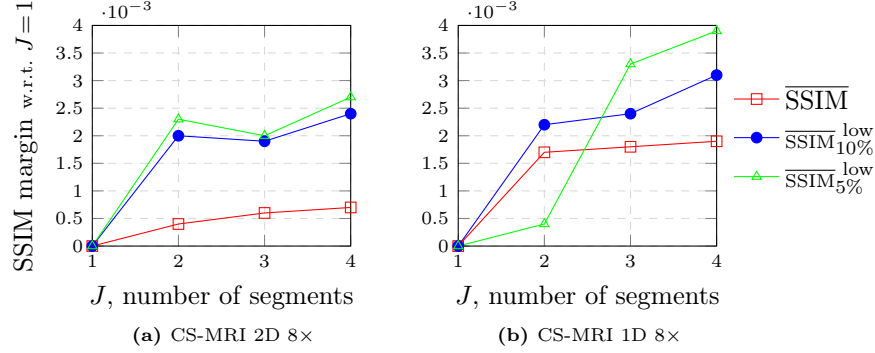
Fig. 6: Trade-off with the number of segments J. The average of the lowest $\{5\%, 10\%, 100\%\}$ of SSIM values according to the number of segments J in CS-MRI (a) 2D 8× and (b) 1D 8× experiment, respectively. All SSIM values were shown as margin with respect to $J = 1$. $\overline{\mathrm{SSIM}}$ reaches a plateau after $J = 2$, but $\overline{\mathrm{SSIM}}_{10\%}^{\mathrm{low}}$ and $\overline{\mathrm{SSIM}}_{5\%}^{\mathrm{low}}$ become more higher in $J = 2, 3$ or 4. These results support our Remark 2.

Table 4: Comparison of rejection sampling and 'kmeans-Sorted'.

SSIM in CS-MRI $J = 3$	2D 4×	2D 8×	1D 4×	1D 8×
Rejection sampling (ours)	**0.9704**	**0.9585**	**0.9624**	**0.9407**
kmeans-Sorted	0.9612	0.9478	0.9493	0.9167

Other sampling methods We employed the rejection sampling to introduce randomness to the mask [27]. To check the effectiveness of the rejection sampling, we compare it with 'kmeans-Sorted' method (*i.e.,* applying 'Self-Sorted' in Tab. 3 to the k-means centroids). In Tab. 4, SSIM in rejection sampling is much higher than 'kmeans-Sorted' in MRI datasets, supporting the effectiveness of rejection sampling.

For more discussions such as runtimes, see the supplementary material.

6   Conclusion

We have presented an adaptive selection of sampling-reconstruction $\mathcal{H}_{1.5}$ framework for Fourier CS. Our method uses an SR space generation model to quantify the high-frequency Bayesian uncertainty of each input; hence is adaptive compared to the existing joint optimization of sampling-reconstruction ($\mathcal{H}_1$) (Theorem 1). Since our method has a dedicated reconstruction network for each sampling mask, unlike adaptive sampling ($\mathcal{H}_2$), our method does not suffer from the Pareto suboptimality (Theorem 2). The proposed method improved SSIM in various Fourier CS experiments, such as CS of facial images and CS-MRI in a practical multi-coil setting.

Acknowledgements

References

1. Arthur, D., Vassilvitskii, S.: K-means++ the advantages of careful seeding. In: ACM-SIAM symposium on Discrete algorithms. pp. 1027–1035 (2007)
2. Bahadir, C.D., Wang, A.Q., Dalca, A.V., Sabuncu, M.R.: Deep-learning-based optimization of the under-sampling pattern in MRI. IEEE Transactions on Computational Imaging **6**, 1139–1152 (2020)
3. Bakker, T., van Hoof, H., Welling, M.: Experimental design for MRI by greedy policy search. NeurIPS **33**, 18954–18966 (2020)
4. Bakker, T., Muckley, M., Romero-Soriano, A., Drozdzal, M., Pineda, L.: On learning adaptive acquisition policies for undersampled multi-coil MRI reconstruction. In: MIDL. pp. 63–85 (2022)
5. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv:1308.3432 (2013)
6. Candes, E., Romberg, J., Tao, T.: Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. IEEE Transactions on Information Theory **52**(2), 489–509 (2006)
7. Chaithya, G., Ramzi, Z., Ciuciu, P.: Learning the sampling density in 2D SPARKLING MRI acquisition for optimized image reconstruction. In: EUSIPCO. pp. 960–964 (2021)
8. Chaithya, G., Weiss, P., Daval-Frérot, G., Massire, A., Vignaud, A., Ciuciu, P.: Optimizing full 3D sparkling trajectories for high-resolution magnetic resonance imaging. IEEE Transactions on Medical Imaging **41**(8), 2105–2117 (2022)
9. De Maio, A., Eldar, Y.C., Haimovich, A.M.: Compressed sensing in radar signal processing. Cambridge University Press (2019)
10. Donoho, D.L.: Compressed sensing. IEEE Transactions on Information Theory **52**(4), 1289–1306 (2006)
11. Hong, S., Kim, S.C., Lee, S.: Advanced direction of arrival estimation using step-learnt iterative soft-thresholding for frequency-modulated continuous wave multiple-input multiple-output radar. IET Radar, Sonar & Navigation **17**(1), 2–14 (2023)
12. Hong, S., Park, I., Chun, S.Y.: On the robustness of normalizing flows for inverse problems in imaging. In: ICCV. pp. 10745–10755 (2023)
13. Lazarus, C., Weiss, P., Chauffert, N., Mauconduit, F., El Gueddari, L., Destrieux, C., Zemmoura, I., Vignaud, A., Ciuciu, P.: SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI. Magnetic Resonance in Medicine **81**(6), 3643–3661 (2019)

14. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: ICCV. pp. 3730–3738 (2015)
15. Lugmayr, A., Danelljan, M., Gool, L.V., Timofte, R.: SRFlow: Learning the super-resolution space with normalizing flow. In: ECCV. pp. 715–732 (2020)
16. Lugmayr, A., Danelljan, M., Timofte, R.: NTIRE 2021 learning the super-resolution space challenge. In: CVPRW. pp. 596–612 (2021)
17. Lugmayr, A., Danelljan, M., Timofte, R., Kim, K.w., Kim, Y., Lee, J.y., Li, Z., Pan, J., Shim, D., Song, K.U., et al.: NTIRE 2022 challenge on learning the super-resolution space. In: CVPRW. pp. 786–797 (2022)
18. Lustig, M., Donoho, D., Pauly, J.M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. Magnetic Resonance in Medicine **58**(6), 1182–1195 (2007)
19. Lustig, M., Donoho, D.L., Santos, J.M., Pauly, J.M.: Compressed sensing MRI. IEEE Signal Processing Magazine **25**(2), 72–82 (2008)
20. Pineda, L., Basu, S., Romero, A., Calandra, R., Drozdzal, M.: Active MR k-space sampling with reinforcement learning. In: MICCAI (2020)
21. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241 (2015)
22. Sanchez, T.: Learning to sample in Cartesian MRI. Phd thesis, EPFL, Lausanne (2022)
23. Song, K.U., Shim, D., Kim, K.w., Lee, J.y., Kim, Y.: FS-NCSR: Increasing diversity of the super-resolution space via frequency separation and noise-conditioned normalizing flow. In: CVPRW. pp. 968–977 (2022)
24. Sriram, A., Zbontar, J., Murrell, T., Defazio, A., Zitnick, C.L., Yakubova, N., Knoll, F., Johnson, P.: End-to-end variational networks for accelerated MRI reconstruction. In: MICCAI (2020)
25. Van Gorp, H., Huijben, I., Veeling, B.S., Pezzotti, N., Van Sloun, R.J.: Active deep probabilistic subsampling. In: ICML. pp. 10509–10518 (2021)
26. Wang, G., Luo, T., Nielsen, J.F., Noll, D.C., Fessler, J.A.: 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)
27. Wang, Z., Arce, G.R.: Variable density compressed image sampling. IEEE Transactions on Image Processing **19**(1), 264–270 (2009)
28. Yang, P., Dong, B.: L2SR: Learning to sample and reconstruct for accelerated MRI. arXiv:2212.02190 (2022)
29. Yin, T., Wu, Z., Sun, H., Dalca, A.V., Yue, Y., Bouman, K.L.: End-to-End sequential sampling and reconstruction for MR imaging. In: MLHC (2021)
30. 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:1811.08839 (2018)
31. Zhang, J., Zhang, H., Wang, A., Zhang, Q., Sabuncu, M., Spincemaille, P., Nguyen, T.D., Wang, Y.: Extending LOUPE for k-space under-sampling pattern optimization in multi-coil MRI. In: MLMIR (2020)
32. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. IEEE Transactions on Image Processing **26**(7), 3142–3155 (2017)
33. Zhang, X., Xu, Z., Jia, N., Yang, W., Feng, Q., Chen, W., Feng, Y.: Denoising of 3d magnetic resonance images by using higher-order singular value decomposition. Medical Image Analysis **19**(1), 75–86 (2015)

34. Zhang, Z., Romero, A., Muckley, M.J., Vincent, P., Yang, L., Drozdzal, M.: Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In: CVPR. pp. 2049–2058 (2019)