

Adaptive Compressed Sensing with Diffusion-Based Posterior Sampling

Noam Elata¹, Tomer Michaeli¹, and Michael Elad²

¹ Dept. of Electric and Computer Engineering ² Dept. of Computer Science
Technion - Israel Institute of Technology, Haifa, Israel
{noamelata@campus, tomer.m@ee, elad@cs}.technion.ac.il

Abstract. Compressed Sensing (CS) facilitates rapid image acquisition by selecting a small subset of measurements sufficient for high-fidelity reconstruction. Adaptive CS seeks to further enhance this process by dynamically choosing future measurements based on information gleaned from data that is already acquired. However, many existing frameworks are often tailored to specific tasks and require intricate training procedures. We propose AdaSense, a novel Adaptive CS approach that leverages zero-shot posterior sampling with pre-trained diffusion models. By sequentially sampling from the posterior distribution, we can quantify the uncertainty of each possible future linear measurement throughout the acquisition process. AdaSense eliminates the need for additional training and boasts seamless adaptation to diverse domains with minimal tuning requirements. Our experiments demonstrate the effectiveness of AdaSense in reconstructing facial images from a small number of measurements. Furthermore, we apply AdaSense for active acquisition of medical images in the domains of magnetic resonance imaging (MRI) and computed tomography (CT), highlighting its potential for tangible real-world acceleration.³

Keywords: Adaptive Compressed Sensing · Posterior Sampling · Pre-trained Diffusion Models

1 Introduction

Compressed sensing (CS) [8, 16] methods have shown promise in image acquisition and reconstruction by leveraging the statistical prior of real-world signals. This allows capturing significant information using fewer measurements, which is crucial when acquisition is time-consuming or costly. To further optimize this process, adaptive CS techniques [2, 29, 43, 50, 56] use reconstructions based on existing measurements to dynamically prioritize subsequent ones. Both adaptive and non-adaptive CS have benefited from the recent surge in deep learning research [27, 65, 73].

Existing methods for measurement selection are typically either heuristic or involve training complex models on simulated subsampled data. The complexity

³ Our code is available at <https://github.com/noamelata/AdaSense>.

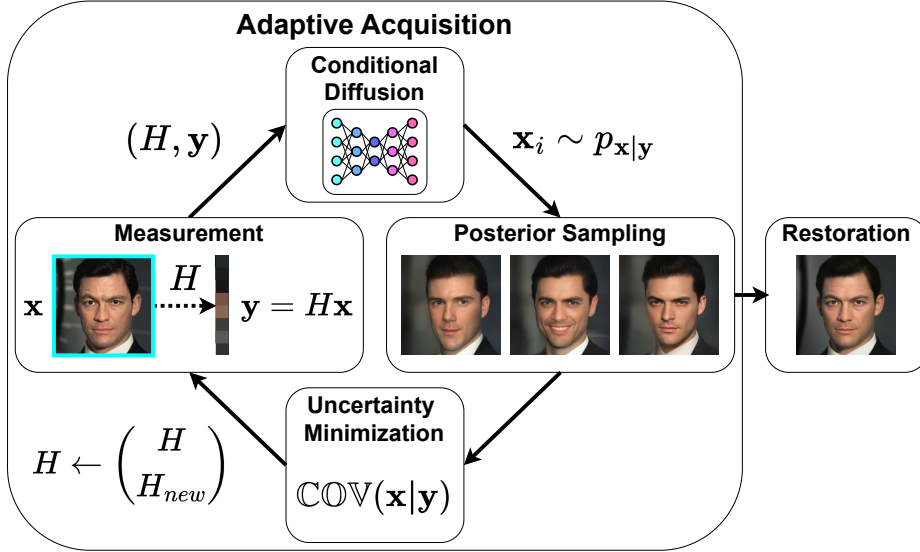


Fig. 1: A diagram of AdaSense. In each acquisition step, the diffusion model generates conditional posterior samples, which are then used to estimate the posterior’s covariance. The ground truth image (highlighted in blue) is measured using our newly-chosen sensing matrix, corresponding to the directions of highest uncertainty in the posterior distribution. This cycle continues until sufficient measurements are acquired. Finally, AdaSense leverages these measurements to restore the final image.

escalates when these methods are adaptive, as they necessitate fine-tuning the algorithm to the specific modality and subsampling schemes applied [3, 51, 64]. Consequently, such approaches are often limited in their modality, and struggle to adapt to changes within their domain.

Driven by recent advancements in zero-shot diffusion-based methods for solving linear inverse problems [15, 24, 52, 59, 62, 63], we propose AdaSense, an image reconstruction technique for active measurement acquisition. These methods leverage pre-trained diffusion models [24, 59, 62] to capture the data distribution, modifying the sampling algorithm to function as posterior distribution samplers. Using candidate posterior samples, AdaSense is able to quantify the reconstruction uncertainty associated with each possible future measurement, guiding the selection of the optimal subsequent measurement through a sequential algorithm. By employing zero-shot diffusion-based methods, AdaSense eliminates the need for additional training or fine-tuning, making it applicable across varied domains and acquisition schemes. Moreover, AdaSense operates under the assumption that access to training data is restricted, which is common in the sensitive medical domain. We offer several key insights to ensure AdaSense is computationally efficient.

To validate our approach, we apply AdaSense in various settings, including general image reconstruction and active acquisition of medical images. For gen-




Fig. 2: Progressive restoration using AdaSense, demonstrating the increasing image reconstruction quality with the accumulation of additional measurements. The number of reconstruction steps is denoted below each image, each step adds 24 measurements.

eral reconstruction, we compare AdaSense to other common subsampling techniques on facial images, and attempt to quantify the performance gain of adaptive restoration compared to non-adaptive alternatives. In the medical image domain, we showcase AdaSense as an active acquisition tool for both k-space magnetic resonance imaging (MRI) subsampling and sparse-view computed tomography (CT). In such scenarios, measurement subsampling may expedite the lengthy acquisition process and reduce patient exposure to radiation [36, 48, 66, 75, 78]. To the best of our knowledge, AdaSense is the first active acquisition algorithm that both utilizes generative priors and is training-free, making it a unique and potentially impactful contribution to the field.

2 Related Work

Diffusion Models. Diffusion models have emerged as the leading approach for image generation and modeling [15, 24, 52, 59, 62, 63]. Diffusion models have been successfully extended to diverse domains, including audio and speech [28, 37, 44], video [39, 41], and medical data [1, 14, 34, 61]. Their capabilities extend beyond generation, excelling at solving inverse problem [12, 53, 54], including many types of image reconstruction tasks. Recent years have seen a surge in zero-shot diffusion-based inverse problems solvers [10, 13, 18, 33, 35, 45, 60, 71] using pre-trained diffusion models. These methods modify the reverse diffusion process by leveraging the available measurements \mathbf{y} through the exploitation of the known degradation form. AdaSense utilizes such a posterior sampler as its core, adapting to any chosen degradation or sensing matrix for measurement acquisition.

Compressed Sensing. Compressed Sensing attempts to find an optimal sensing matrix for reconstruction of a given set of signals, initially employed using a sparsifying basis [8, 16], and later using generative priors [5, 31]. Adaptive CS integrates reconstruction with the continual optimization of the sensing matrix [29]. Similar works [2, 21, 22, 43, 50, 56] have refined the algorithmic underpinnings of adaptive compressed sensing, albeit relying on assumptions regarding data sparsity. In recent years, the integration of deep learning has propelled adaptive CS

methods even further [23, 47], with some methods incorporating the deep priors of generative models [7, 64, 76]. However, many existing methods are confined to specific domains and are thus tuned for particular tasks [58, 78]. Additionally, certain methods are limited in the types of degradation they can handle, focusing solely on subsampling of pixels [7, 64]. In this work, we propose a novel perspective on the application of adaptive CS as a unified framework, using a stochastic posterior sampler. As these samplers grow in quantity and quality following the recent advances of deep learning and diffusion models, we hope to lay the foundation of a practical and straightforward adaptive CS algorithm.

Medical Imaging. In medical imaging, such as MRI or CT, measurement subsampling may expedite the lengthy acquisition process and reduce patient exposure to radiation [4, 48, 78]. In MRI, images are obtained in the Fourier spectrum, named the k-space, using magnetic fields. Common MRI subsampling involves acquiring only a subset of possible k-space frequencies. Extensive research efforts have been directed towards developing reconstruction algorithms for subsampled data [9, 20, 55, 67, 68, 79]. Specifically, diffusion models have demonstrated their prowess in the restoration domain for MR images [26, 61, 72], yet previous work falls short of exploring their potential for active MRI acquisition.

Current approaches for active MRI acquisition often involve deep neural network (DNN) training [7, 30, 64], in many cases using reinforcement learning for policy selection [3, 51, 78]. Despite impressive results, previous MRI active acquisition techniques’ reliance on DNN training is both costly and a limiting factor when adapting to new subsampling schemes. In contrast, AdaSense only requires a pre-trained diffusion model. This allows the same algorithm to be applied seamlessly across diverse subsampling patterns, such as switching from vertical to horizontal frequency sampling in the MRI domain.

CT images are acquired using X-ray projections and reconstructed by means of tomography. Sparse-view CT aims to reduce acquisition time and radiation dose by utilizing fewer projection angles for reconstruction. Recent work has introduced deep learning, including diffusion and score-based models, for sparse-view CT reconstruction [19, 69, 72]. Yet, in these methods the subsampling scheme is assumed to be pre-determined and non-adaptive. Similar to MRI, the active acquisition of sparse-view CT has been explored using reinforcement learning for DNN training [58, 70]. In our work, we show AdaSense offers a training-free approach, demonstrating its potential as an active sparse-view acquisition tool for CT reconstruction.

Blind Inverse Problems. Recent research in blind image restoration utilizes diffusion-based inverse problem solvers to adapt the sensing matrix [11, 38, 49]. While our approach shares similarities with these methods by using diffusion-based approaches for determining the sensing matrix, our goals differ fundamentally. We focus on actively selecting the optimal series of measurements taken from an undamaged image, rather than attempting to revert an unknown degradation in an already captured image.

3 Method

We are interested in reconstructing a signal $\mathbf{x} \in \mathbb{R}^D$ from $d < D$ linear measurements of the form

$$\mathbf{y} = \mathbf{H}\mathbf{x}, \quad (1)$$

where $\mathbf{H} \in \mathbb{R}^{d \times D}$ is the sensing matrix and $\mathbf{y} \in \mathbb{R}^d$ is the vector of measurements. We assume that \mathbf{x} is a random vector whose distribution $p_{\mathbf{x}}$ is either known or has been learned by a generative model based on a large corpus of training samples. Our goal is to determine the matrix \mathbf{H} that allows the best possible recovery of \mathbf{x} in the mean squared error (MSE) sense. Importantly, we assume we may select the measurements (*i.e.* the rows of \mathbf{H}) in a sequential manner, taking advantage of information gained from previous measurements.

Let us first discuss the simplified non-adaptive setting, in which the entire matrix \mathbf{H} is determined a-priori without dependence on the measurements. In this setting, the optimization problem we would like to solve is

$$\operatorname{argmin}_{\mathbf{H}, f} \mathbb{E} [\|f(\mathbf{H}\mathbf{x}) - \mathbf{x}\|^2], \quad (2)$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}^D$ is a function that reconstructs \mathbf{x} from the measurements $\mathbf{H}\mathbf{x}$. In principle, the optimal \mathbf{H} and $f(\cdot)$ may be determined by training an auto-encoder with a linear encoder \mathbf{H} and a nonlinear decoder $f(\cdot)$. However, as will become clear below, adapting this strategy to the adaptive CS setting is computationally impractical. A more scalable approach is to restrict the decoder function to be an affine transformation of the form $f(\mathbf{y}) = \mathbf{U}\mathbf{y} + \mathbf{b}$, relying on the assumption that the optimal \mathbf{H} for this case is also nearly optimal for Eq. (2). In this setting, the optimization problem becomes

$$\operatorname{argmin}_{\mathbf{U}, \mathbf{b}, \mathbf{H}} \mathbb{E} [\|\mathbf{U}\mathbf{H}\mathbf{x} + \mathbf{b} - \mathbf{x}\|^2]. \quad (3)$$

This problem corresponds to principal component analysis (PCA) [17, 25]. Therefore, the optimal sensing matrix \mathbf{H} has the top d eigenvectors of $\operatorname{Cov}[\mathbf{x}]$ as its rows, the optimal decoding matrix is $\mathbf{U} = \mathbf{H}^\top$, and the optimal bias is $\mathbf{b} = (\mathbf{I} - \mathbf{H}^\top \mathbf{H})\mathbb{E}[\mathbf{x}]$. Namely, PCA minimizes the *uncertainty*, defined as the MSE attained by the linear minimum-MSE (MMSE) predictor of \mathbf{x} based on the measurements \mathbf{y} . This is achieved by choosing the measurement matrix \mathbf{H} accordingly to minimize this error.

We next address the *adaptive* case, where the rows of \mathbf{H} are chosen sequentially, such that each row (or chunk of rows) is chosen based on the measurements acquired with the preceding rows. We take a greedy approach, where in each stage we choose the new rows so as to allow minimization of the reconstruction error at that stage, without taking into consideration their effect on the reconstruction error at future stages. We denote by $\mathbf{H}_{i:j}$ the sub-matrix of \mathbf{H} corresponding to rows i to j (including row i and excluding row j) and denote by $\mathbf{y}_{i:j}$ the corresponding measurements, $\mathbf{y}_{i:j} = \mathbf{H}_{i:j}\mathbf{x}$. We aim to construct the whole matrix \mathbf{H} in N steps, selecting r new rows in each step. Therefore, at

step n , we rely on the previous nr measurements, $\mathbf{y}_{0:nr}$. Following the same PCA strategy discussed for the non-adaptive case, we propose choosing the sub-matrix $\mathbf{H}_{nr:nr+r}$ at stage n , as

$$\begin{aligned}\mathbf{H}_{nr:nr+r} &= \underset{\mathbf{V}, \mathbf{U}, \mathbf{b}, \tilde{\mathbf{H}}}{\operatorname{argmin}} \mathbb{E} \left[\left\| (\mathbf{V} \mathbf{U}) \begin{pmatrix} \mathbf{H}_{0:nr} \\ \tilde{\mathbf{H}} \end{pmatrix} \mathbf{x} + \mathbf{b} - \mathbf{x} \right\|^2 \middle| \mathbf{y}_{0:nr} \right] \\ &= \underset{\mathbf{V}, \mathbf{U}, \mathbf{b}, \tilde{\mathbf{H}}}{\operatorname{argmin}} \mathbb{E} \left[\left\| \mathbf{V} \mathbf{y}_{0:nr} + \mathbf{U} \tilde{\mathbf{H}} \mathbf{x} + \mathbf{b} - \mathbf{x} \right\|^2 \middle| \mathbf{y}_{0:nr} \right].\end{aligned}\quad (4)$$

Here, $(\mathbf{V} \mathbf{U})$ serves as the linear reconstruction matrix and we used the fact that $\mathbf{H}_{0:nr} \mathbf{x} = \mathbf{y}_{0:nr}$. Due to the conditioning on $\mathbf{y}_{0:nr}$, the vector $\mathbf{V} \mathbf{y}_{0:nr} + \mathbf{b}$ can be viewed as a deterministic bias term (playing the role of \mathbf{b} in Eq. (3)). Therefore, similarly to Eq. (3), the optimal $\tilde{\mathbf{H}}$ in this case corresponds to the top r eigenvectors $\operatorname{Cov}[\mathbf{x} | \mathbf{y}_{0:nr}]$.

Once $\mathbf{H}_{nr:nr+r}$ is determined, we can use it to obtain r new measurements, $\mathbf{y}_{nr:nr+r} = \mathbf{H}_{nr:nr+r} \mathbf{x}$, and append them to our previous ones as $\mathbf{y}_{0:nr+r} = (\mathbf{y}_{0:nr}^\top \mathbf{y}_{nr:nr+r}^\top)^\top$. We may repeat this process until we have chosen all Nr rows of \mathbf{H} . Finally, we may restore our final measurements \mathbf{y} with a nonlinear reconstruction function $f(\mathbf{y})$.

The iterative approach described above requires knowing the posterior covariance matrix $\operatorname{Cov}[\mathbf{x} | \mathbf{y}_{0:nr}]$ in each stage, or at least its top eigenvectors. To obtain an approximation of these principal components, we propose harnessing zero-shot posterior sampling methods that are based on a pre-trained diffusion model [12, 13, 18, 33, 35, 45, 60, 71]. These methods allow generating samples from the posterior distribution of \mathbf{x} given $\mathbf{y}_{0:nr}$ for problems in which $\mathbf{y}_{0:nr}$ is a linear transformation of \mathbf{x} , as in our setting. Given a set of s such posterior samples, we can apply PCA to determine the top eigenvectors of the (empirical) posterior covariance. Our sampling algorithm, Adasense, is presented in Algorithm 1, including the constrained measurement scenario which we describe in Sec. 3.1 below. A diagram of our proposed method may be found in Fig. 1, as well as a qualitative example of progressively increasing restoration quality with the accumulation of additional measurements in Fig. 2.

3.1 Constrained Measurement

In certain real-world applications, the sensing matrix is constrained due to physical limitations. One notable example is MRI, in which the signal is measured in k -space, so that the rows of \mathbf{H} are restricted to be rows of the Fourier transform matrix. Another important example is CT, in which the signal is measured in Radon space. In this setting, the rows of \mathbf{H} are constrained to be rows of the Radon transform matrix. To adapt our method to such scenarios, we therefore need to choose the best option from a predefined set \mathcal{H} of feasible sensing matrices. In other words, we need to solve Eq. (4) under the constraint $\tilde{\mathbf{H}} \in \mathcal{H}$. As we show in App. A.1, the matrix $\tilde{\mathbf{H}} \in \mathcal{H}$ achieving the optimum of this constrained

Algorithm 1 AdaSense

Require: Generative model for $p_{\mathbf{x}}$, number of steps N , number of measurements per step r , number of posterior samples s .
initialize $\mathbf{y}_{0:0}$ as an empty vector of measurements.
for $n \in \{0 : N - 1\}$ **do**
 $\{\mathbf{x}_i\}_{i=1}^s \sim p_{\mathbf{x}|\mathbf{y}_{0:nr}}$ ▷ generate s posterior samples
 $\{\mathbf{x}_i\}_{i=1}^s \leftarrow \{\mathbf{x}_i - \frac{1}{s} \sum_{j=1}^s \mathbf{x}_j\}_{i=1}^s$
 if not Constrained Measurements **then** ▷ select r principal components
 $\mathbf{H}_{nr:nr+r} \leftarrow$ top r right singular vectors of $(\mathbf{x}_1 \dots \mathbf{x}_s)^\top$
 else ▷ select r best sensing rows
 $\mathbf{H}_{nr:nr+r} \leftarrow \operatorname{argmax}_{\tilde{\mathbf{H}} \in \mathcal{H}} \sum_{i=1}^s \mathbf{x}_i^\top \tilde{\mathbf{H}}^\dagger \tilde{\mathbf{H}} \mathbf{x}_i$
 end if
 $\mathbf{y}_{0:nr+r} \leftarrow \text{Append}(\mathbf{y}_{0:nr}, \mathbf{H}_{nr:nr+r} \mathbf{x})$ ▷ measure the real image \mathbf{x}
end for
return $\mathbf{x}_1 = f(\mathbf{y}_{0:Nr})$ ▷ posterior sampling or alternative restoration

optimization problem is the solution to

$$\operatorname{argmax}_{\tilde{\mathbf{H}} \in \mathcal{H}} \operatorname{Tr} \left\{ \tilde{\mathbf{H}} \operatorname{Cov}[\mathbf{x}|\mathbf{y}_{0:nr}]^2 \tilde{\mathbf{H}}^\top \left(\tilde{\mathbf{H}} \operatorname{Cov}[\mathbf{x}|\mathbf{y}_{0:nr}] \tilde{\mathbf{H}}^\top \right)^{-1} \right\}. \quad (5)$$

Note that as opposed to the unconstrained case, if \mathcal{H} is not the entire space $\mathbb{R}^{r \times D}$, then the solution to this optimization problem is not necessarily the top eigenvectors of $\operatorname{Cov}[\mathbf{x}|\mathbf{y}_{0:nr}]$.

Unfortunately, in many practical cases the number of rows r selected in each step is unavoidably large. For example, in MRI, we would often like to measure a whole column of the 2d Fourier transform of the image at a time. Similarly in CT, we would often like to measure a whole column of the sinogram at a time. In such scenarios, attempting to directly approximate the solution to Eq. (5) becomes prohibitively expensive. This is because for the $r \times r$ matrix $\tilde{\mathbf{H}} \operatorname{Cov}[\mathbf{x}|\mathbf{y}_{0:nr}] \tilde{\mathbf{H}}^\top$ to be invertible, the rank of our approximation of the covariance $\operatorname{Cov}[\mathbf{x}|\mathbf{y}_{0:nr}]$ must be at least r . This requires generating $s \geq r$ posterior samples in each step, which is impractical with today's zero-shot posterior samplers for the typical values of r in *e.g.* MRI and CT. To overcome this computational difficulty, we propose a sub-optimal solution, which nonetheless works quite well in practice. Specifically, rather than optimizing over the reconstruction matrix \mathbf{U} in Eq. (4), we set it to be $\tilde{\mathbf{H}}^\dagger$, which is its optimal values in the unconstrained scenario. As we demonstrate in App. A.2, this simplifies the problem to

$$\operatorname{argmax}_{\tilde{\mathbf{H}} \in \mathcal{H}} \mathbb{E} \left[(\mathbf{x} - \mathbb{E}[\mathbf{x}|\mathbf{y}_{0:i}])^\top \tilde{\mathbf{H}}^\dagger \tilde{\mathbf{H}} (\mathbf{x} - \mathbb{E}[\mathbf{x}|\mathbf{y}_{0:i}]) \middle| \mathbf{y}_{0:i} \right]. \quad (6)$$

The solution to this problem can again be approximated using posterior samples, by replacing expectations by averages and exhaustively scanning all matrices in \mathcal{H} . This search space becomes impractically large to scan if we attempt to select several frequency columns in MRI or several sinogram columns in CT at once. In such cases, we use a heuristic, as detailed in App. D.

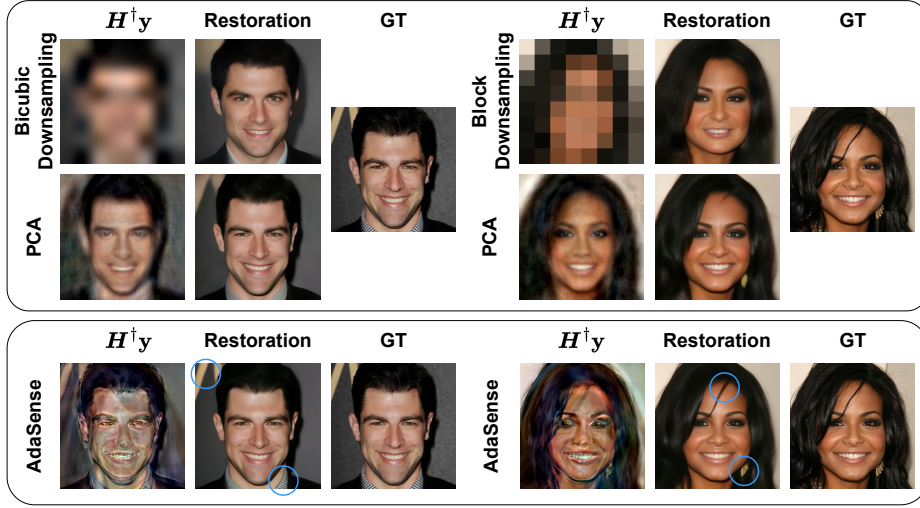


Fig. 3: Face image restoration using AdaSense. We illustrate how our adaptive approach (bottom) is better for image reconstruction than block-downsampling or bicubic-downsampling (top), where measurements have 192 elements. Notably, AdaSense successfully preserves some finer details (circled in blue). Also, we compare AdaSense with the optimal non-adaptive approach, using PCA (Eq. (3)) on the distribution of training images. The images are restored using DDRM.

3.2 Restoration

Once we have obtained the optimal sensing matrix \mathbf{H} using the method outlined above, we seek a function $f(\cdot)$ to restore our final set of measurements $\mathbf{y}_{0:N_T}$. The most straightforward approach would be to use the same zero-shot posterior sampler for the final restoration, generating a single or multiple reconstructions. Another alternative, involves using the average of several posterior sampler outputs. This average approximates the posterior mean, which is the minimum MSE (MMSE) estimator. A third alternative is to take the restoration function $f(\cdot)$ to be a neural network that is specifically tailored to the modality. Section 4.3 experiments with classical and deep learning reconstruction approaches specifically designed for MRI for improved restoration.

3.3 Acceleration of Posterior Sampling

AdaSense, as outlined in Algorithm 1, may be used with any posterior sampler. However, to avoid training for specific degradations we propose using a zero-shot diffusion restoration method. While the mechanics of different methods differ, many approaches solely rely on the computation of the Moore–Penrose pseudo-inverse \mathbf{H}^\dagger of the degradation matrix \mathbf{H} . Also, several zero-shot diffusion restoration methods are consistent with the measurements \mathbf{y} , i.e. for any

Table 1: Comparisons of several DDRM-based restorations on different measurements of rank 192. Also, the table depicts how changing the final restoration algorithm can decrease distortion at the expense of image quality, for the same chosen measurements.

CELEBA-HQ 256	PSNR↑	SSIM↑	LPIPS↓	DeepFace	CosSim↓
Block Downsampling	20.50	0.6128	0.3035		0.6134
Bicubic Downsampling	20.86	0.6193	0.2964		0.5865
Walsh-Hadamard CS	10.75	0.3874	0.5920		0.8836
Missing Pixels	10.85	0.3865	0.6064		0.8782
PCA	24.60	0.7190	0.2307		0.3010
AdaSense	26.20	0.7515	0.1950		0.2674
AdaSense (Mean)	26.87	0.7713	0.2449		0.2614

generated \mathbf{x} the equation $\mathbf{y} = \mathbf{H}\mathbf{x}$ holds⁴. Below, we describe how we can greatly increase the efficiency of our implementation using such approaches.

The computation of \mathbf{H}^\dagger , could generally be computed for any degradation \mathbf{H} using a computationally expensive SVD. Because the degradation matrix \mathbf{H} is chosen at runtime in AdaSense, the SVD computation could potentially lead to slower sampling. Nevertheless, we offer several insights that explain why this repeated computation of SVD can be disregarded. Because a consistent posterior sampler is used, the variance along previously selected sensing matrix is necessarily zero. Therefore, for general measurements and cases of constrained measurements where the measurements are inherently non-overlapping (such as MRI subsampling), our choices for the next measurement $\tilde{\mathbf{H}}$ are necessarily orthogonal to all previous measurements. This leads to the useful property where the low-rank SVD of our matrices $\mathbf{H}_{0:nr}$ is $(U, \Sigma, V^\top) = (I, I, \mathbf{H}_{0:nr})$, eliminating the need for additional computation. Further details on are provided in the supplementary material. In other cases, such as sparse-view CT reconstruction, limiting AdaSense to a small number of total measurements ensures that the matrix $\mathbf{H}_{0:nr}$ remains low-rank, making SVD computation inexpensive.

4 Experiments

4.1 General Images

We begin by testing AdaSense on 256×256 face images taken from the CelebA-HQ [32, 40] validation dataset. The pre-trained model from SDEdit [46] was used with DDRM [33] for diffusion-based sampling. We evaluate AdaSense’s capabilities to select the best possible measurements by comparing reconstructions made using AdaSense’s proposed sensing matrix to reconstructions from different common degradations. All degraded measurements are of dimension 192 and have

⁴ While consistency should hold by the definition of posterior sampling, not all zero-shot diffusion restoration methods guarantee consistency [12, 33, 60].

been restored using the same algorithm (DDRM [33]). AdaSense employed $N = 8$ consecutive iterations, selecting $r = 24$ elements per iteration. Figure 3 showcases a qualitative comparison of AdaSense with other restoration approaches, highlighting its superior ability to preserve fine details and subject identity compared to alternatives. Qualitative result in Tab. 1 reveal that AdaSense outperforms other non-adaptive methods across all metrics, including PSNR, SSIM, LPIPS [77], and DeepFace [57] cosine similarity. Notably, AdaSense surpasses a non-adaptive PCA approach that selects the optimal sensing matrix following Eq. (3) using real training images. This achievement highlights the effectiveness of AdaSense, especially considering it relies solely on approximated generated samples. In addition, we show how using the mean of the posterior sampler as suggested in Sec. 3.2 lowers the distortion of the final reconstruction (measured in PSNR and SSIM) while sacrificing the perceptual quality (measured in LPIPS).

4.2 Importance of Adaptivity

The hyperparameter N governs the number of iterations in the measurement acquisition process. Within a constant total number of measurements $N \cdot r$, AdaSense is deemed ‘less adaptive’ when r is larger at the expense of N , and ‘more adaptive’ otherwise. For $N = 1$, a single iteration is used, essentially eliminating the adaptive component of the algorithm.

To validate the significance of adaptivity in the restoration process, we compare our method’s performance using different numbers of iterations for measurement acquisition, while maintaining the total number of measurements $N \cdot r$. This ensures a fair comparison focused solely on the effect of adaptivity. Also, the value of s remains consistently proportional to r at a ratio of $4/3$, ensuring consistent computational demands (measured by the total number of generated samples) across experiments. We retain the posterior sampler from Sec. 4.1, and adjust the AdaSense hyperparameters N, r, s to change the adaptivity of our algorithm. The results, illustrated in Fig. 4 demonstrates that both PSNR and LPIPS [77] improve as the number of iterations used for measurement acquisition increases. This trend underscores the positive impact of adaptivity on the system’s robustness and effectiveness, highlighting its crucial role within our proposed method.

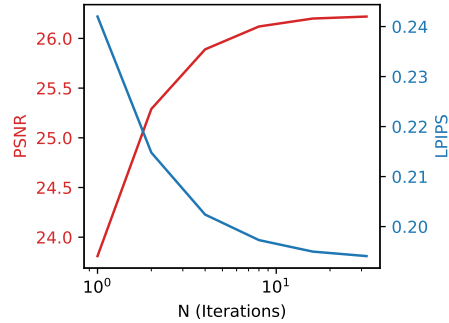


Fig. 4: Graphs of the effect of adaptivity on restoration. AdaSense is used to restore the same set of images using the same total number of measurements $N \cdot r$, but with different numbers of iterations N and measurements per iteration r . The image restoration improves considerably the more adaptive the algorithm.

Table 2: Comparison of AdaSense with non-adaptive MRI restoration algorithms. The table compares subsampling masks generated according to [6, 36, 75] with sensing matrices selected by Adasense. The number following R denotes the acceleration factor. The table displays results for three reconstruction methods: DDRM [33], mean of 16 DDRM samples, and $L1$ wavelet regularized reconstruction [42].

Vertical Subsampling Mask	DDRM		DDRM Mean		Wavelet	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Random R10	24.56	0.4786	25.29	0.5067	22.50	0.4365
Equi-spaced R10	24.73	0.4767	25.44	0.5043	22.56	0.4339
AdaSense R10	27.01	0.5229	27.58	0.5431	25.18	0.4874
General Subsampling Mask	PSNR↑		SSIM↑		PSNR↑	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Poisson Disk R400	23.06	0.3728	23.86	0.4010	18.22	0.3504
AdaSense R400	25.26	0.4124	25.76	0.4314	23.32	0.4010
Poisson Disk R200	24.31	0.3999	24.95	0.4234	22.50	0.3619
AdaSense R200	25.94	0.4331	26.39	0.4510	24.44	0.4219

Table 3: Comparisons of MRI Active Acquisition Algorithms. This table compares the performance of AdaSense with other MRI active acquisition algorithms. The acceleration factors are denoted by R, and the number of pre-set low-frequencies in L, following [51]. We also include results for the non-adaptive baseline strategy ‘Low-to-High’ [51], and the ‘Greedy Oracle’ [51] approach, a theoretical upper-bound on the achievable performance. As shown in the table, AdaSense demonstrates comparable performance to training-based methods.

Method	R8 - 30L		R16 - 2L	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Low-to-High (non-adaptive) [51]	28.32	0.5948	25.66	0.4908
Reducing Uncertainty [78]	28.66	0.6021	26.43	0.5386
SS-DDQN [51]	28.99	0.6129	27.70	0.5691
DS-DDQN [51]	28.98	0.6135	27.60	0.5650
Greedy Oracle [51]	29.22	0.6264	28.23	0.5846
AdaSense (Reconstructor)	28.89	0.6108	27.51	0.5547

4.3 Active MRI Subsampling

In this section, we demonstrate the usefulness of AdaSense in active MRI subsampling. We use a diffusion model for complex 640×368 single-coil knee MRI images which we train on the FastMRI [36, 75] dataset, following established data pre-processing [51] and model architecture and training procedures [24]. We apply AdaSense for image reconstruction using DDRM [33] under various acceleration schemes. Table 2 compares AdaSense to non-adaptive subsampling masks, including general subsampling with acceleration factor of 200 and 400 and vertical subsampling with acceleration factor of 10, created following [6] and [36, 75] respectively. Images are restored using posterior sampling or with




Fig. 5: MRI active restoration using AdaSense. We compare various subsampling masks with AdaSense, restoring vertical subsampling with acceleration 10 and subsampling with acceleration 400. Images are reconstructed using DDRM [33]. Each column of the figure shows, from left to right: the subsampling mask, the zero-filled measurements, the restored image, the ground-truth image, and the absolute error. The central 320×320 region is cropped and the image intensity is displayed.

a classical wavelet approach [42]. We measure PSNR and SSIM on the central 320×320 region of the reconstructed MR images, aligning with the evaluation method used in similar works [51, 75]. As shown, AdaSense outperforms all non-adaptive measurements. Visualizations of the reconstructions are provided in Fig. 5. Furthermore, Tab. 3 compares AdaSense against established MRI active acquisition methods [51, 78]. For a fair comparison, all approaches use the same final reconstruction model employed in the referenced works, termed ‘Reconstructor’. We include results for the non-adaptive baseline strategy ‘Low-to-High’ [51], and the theoretical upper-bound ‘Greedy Oracle’ [51]. Specifically, the ‘Greedy Oracle’ upper-bound is computed by searching all possible future measurements during each acquisition step, and selecting the ones minimizing the reconstruction error relative to the *ground truth*, which is unavailable in real settings. Consistent with [51], frequency selection utilizes only the image intensity of the central 320×320 region, where model evaluation takes place (details in App. D). Visual examples are included in Fig. 6. AdaSense remains competitive




Fig. 6: MRI active restoration using AdaSense, using reconstructor models from [51]. The images depict image intensity for the central 320×320 region, applying vertical acceleration factors 8 and 16, following the strategies 30L and 2L respectively. Each column displays, from left to right: the subsampling mask, the zero-filled measurements, our AdaSense reconstruction, the ground-truth image, and the absolute error.

with similar training-based methods, despite requiring no additional training beyond the pre-trained diffusion model, and never encountering degraded data during the diffusion model’s training. This underscores AdaSense’s advantage in adapting to any subsampling scheme.

4.4 Sparse-View CT Reconstruction

AdaSense can also be used for sparse-view CT reconstruction, by selecting the best projection angles for restoration. For this experiment, we have trained a simple diffusion model following [24] for 256×256 CT images from on the DeepLesion [74] dataset. We simulate restoration of parallel-beam projections and measure PSNR and SSIM only within the valid region of the simulated projections. Similar to previous experiments, we use DDRM [33] as our posterior sampler. Applying AdaSense to various acceleration schemes on samples from the validation set yields promising results, as showcased in Fig. 7 through both qualitative and quantitative data for sparse-view acquisitions.

5 limitations

While AdaSense demonstrates impressive performance without additional training, several limitations warrant discussion. Firstly, AdaSense’s effectiveness is inherently tied to the quality of the posterior sampler, which forms the core of the method. Any limitations in the sampler’s capabilities will be reflected in AdaSense’s performance. Secondly, generating multiple samples requires high computational cost. Fortunately, the rapid advancements in diffusion-based posterior samplers offer promising avenues for addressing both limitations




Fig. 7: CT sparse-view image restoration using AdaSense. Left: Reconstructions obtained using various acquisition strategies. Each reconstruction is accompanied by its corresponding mask, which illustrates the selected acquisition angles. Additionally, an enlarged center crop of the respective reconstruction is provided for detailed visualization. Right: Quantitative evaluations of CT restorations using PSNR and SSIM. Both these qualitative and quantitative findings results demonstrate the potential of AdaSense for active acquisition in diverse domains.

in the near future. Additionally, the current implementation of AdaSense is constrained to linear measurements. Exploring the potential of non-linear measurement schemes presents an interesting avenue for future research.

6 Conclusions

In this work, we introduced AdaSense, a novel method for adaptive image reconstruction that leverages zero-shot sampling with pre-trained diffusion. AdaSense sequentially refines the reconstruction by strategically selecting the most informative measurements based on the current state. A key strength of AdaSense lies in its versatility. The algorithm remains applicable across diverse image domains, including facial images, MRI, and CT, only adapting the domain-specific diffusion model to capture the unique characteristics of each modality.

Future work could investigate methods for accelerating AdaSense, perhaps by using previously generated samples. Additionally, it would be interesting to explore extending AdaSense to handle more complex degradation models, such as burst-photography or three-dimension reconstruction with best-next-view. Despite the promising results in the medical setting, we emphasize that more experiments should be conducted before AdaSense can be used in clinical settings.

Acknowledgements

This research was partially supported by the Council for Higher Education - Planning & Budgeting Committee, Israel. This research was also partially supported by the Israel Science Foundation (grant no. 2318/22), by a gift from Elbit Systems, and by the Ollendorff Minerva Center, ECE faculty, Technion.

References

1. Alcaraz, J.M.L., Strodthoff, N.: Diffusion-based conditional ecg generation with structured state space models. *Computers in Biology and Medicine* p. 107115 (2023)
2. Arias-Castro, E., Candes, E.J., Davenport, M.A.: On the fundamental limits of adaptive sensing. *IEEE Transactions on Information Theory* **59**(1), 472–481 (2012)
3. Bakker, T., van Hoof, H., Welling, M.: Experimental design for mri by greedy policy search. *Advances in Neural Information Processing Systems* **33**, 18954–18966 (2020)
4. Barkan, O., Weill, J., Averbuch, A., Dekel, S.: Adaptive compressed tomography sensing. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2195–2202 (2013)
5. Bora, A., Jalal, A., Price, E., Dimakis, A.G.: Compressed sensing using generative models. In: *International conference on machine learning*. pp. 537–546. PMLR (2017)
6. Bridson, R.: Fast poisson disk sampling in arbitrary dimensions. *SIGGRAPH sketches* **10**(1), 1 (2007)
7. van de Camp, K.C., Joudeh, H., Antunes, D.J., van Sloun, R.J.: Active subsampling using deep generative models by maximizing expected information gain. In: *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2023)
8. Candès, E.J., et al.: Compressive sampling. In: *Proceedings of the international congress of mathematicians*. vol. 3, pp. 1433–1452. Madrid, Spain (2006)
9. Chen, F., Taviani, V., Malkiel, I., Cheng, J.Y., Tamir, J.I., Shaikh, J., Chang, S.T., Hardy, C.J., Pauly, J.M., Vasanawala, S.S.: Variable-density single-shot fast spin-echo mri with deep learning reconstruction by using variational networks. *Radiology* **289**(2), 366–373 (2018)
10. Choi, J., Kim, S., Jeong, Y., Gwon, Y., Yoon, S.: Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938* (2021)
11. Chung, H., Kim, J., Kim, S., Ye, J.C.: Parallel diffusion models of operator and image for blind inverse problems. *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023)
12. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. In: *The Eleventh International Conference on Learning Representations* (2023), <https://openreview.net/forum?id=0nD9zGAGT0k>
13. Chung, H., Sim, B., Ryu, D., Ye, J.C.: Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems* **35**, 25683–25696 (2022)
14. Chung, H., Ye, J.C.: Score-based diffusion models for accelerated MRI. *Medical Image Analysis* **80**, 102479 (2022)

15. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. *Advances in Neural Information Processing Systems* **34**, 8780–8794 (2021)
16. Donoho, D.L.: Compressed sensing. *IEEE Transactions on information theory* **52**(4), 1289–1306 (2006)
17. F.R.S., K.P.: Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**(11), 559–572 (1901). <https://doi.org/10.1080/14786440109462720>
18. Graikos, A., Malkin, N., Jojic, N., Samaras, D.: Diffusion models as plug-and-play priors. *Advances in Neural Information Processing Systems* **35**, 14715–14728 (2022)
19. Guan, B., Yang, C., Zhang, L., Niu, S., Zhang, M., Wang, Y., Wu, W., Liu, Q.: Generative modeling in sinogram domain for sparse-view ct reconstruction. *IEEE Transactions on Radiation and Plasma Medical Sciences* (2023)
20. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine* **79**(6), 3055–3071 (2018)
21. Haupt, J., Baraniuk, R., Castro, R., Nowak, R.: Sequentially designed compressed sensing. In: 2012 IEEE Statistical Signal Processing Workshop (SSP). pp. 401–404. IEEE (2012)
22. Haupt, J., Nowak, R., Castro, R.: Adaptive sensing for sparse signal recovery. In: 2009 IEEE 13th Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop. pp. 702–707. IEEE (2009)
23. Haussmann, M., Hamprecht, F.A., Kandemir, M.: Deep active learning with adaptive acquisition. *arXiv preprint arXiv:1906.11471* (2019)
24. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* **33**, 6840–6851 (2020)
25. Hotelling, H.: Analysis of a complex of statistical variables into principal components. *Journal of educational psychology* **24**(6), 417 (1933)
26. Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A.G., Tamir, J.: Robust compressed sensing mri with deep generative priors. *Advances in Neural Information Processing Systems* **34**, 14938–14954 (2021)
27. Jalal, A., Liu, L., Dimakis, A.G., Caramanis, C.: Robust compressed sensing using generative models. *Advances in Neural Information Processing Systems* **33**, 713–727 (2020)
28. Jeong, M., Kim, H., Cheon, S.J., Choi, B.J., Kim, N.S.: Diff-tts: A denoising diffusion model for text-to-speech. *arXiv preprint arXiv:2104.01409* (2021)
29. Ji, S., Xue, Y., Carin, L.: Bayesian compressive sensing. *IEEE Transactions on signal processing* **56**(6), 2346–2356 (2008)
30. Jin, K.H., Unser, M., Yi, K.M.: Self-supervised deep active accelerated mri. *arXiv preprint arXiv:1901.04547* (2019)
31. Kamath, A., Karmalkar, S., Price, E.: Lower bounds for compressed sensing with generative models. *arXiv preprint arXiv:1912.02938* (2019)
32. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196* (2017)
33. Kavar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. In: *Advances in Neural Information Processing Systems* (2022)
34. Kavar, B., Elata, N., Michaeli, T., Elad, M.: Gsure-based diffusion model training with corrupted data. *arXiv preprint arXiv:2305.13128* (2023)
35. Kavar, B., Vaksman, G., Elad, M.: SNIPS: Solving noisy inverse problems stochastically. *Advances in Neural Information Processing Systems* **34**, 21757–21769 (2021)

36. Knoll, F., Zbontar, J., Sriram, A., Muckley, M.J., Bruno, M., Defazio, A., Par-
ente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdalv, M.,
Romero, A., Rabbat, M., Vincent, P., Pinkerton, J., Wang, D., Yakubova, N.,
Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastMRI: A
publicly available raw k-space and DICOM dataset of knee images for accelerated
MR image reconstruction using machine learning. *Radiology: Artificial Intelligence*
2(1), e190007 (2020)
37. Kong, Z., Ping, W., Huang, J., Zhao, K., Catanzaro, B.: Diffwave: A versatile
diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761* (2020)
38. Laroche, C., Almansa, A., Coupette, E.: Fast diffusion em: a diffusion model for
blind inverse problems with application to deconvolution. In: *Proceedings of the*
IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5271–5281
(2024)
39. Liu, Y., Zhang, K., Li, Y., Yan, Z., Gao, C., Chen, R., Yuan, Z., Huang, Y., Sun,
H., Gao, J., et al.: Sora: A review on background, technology, limitations, and
opportunities of large vision models. *arXiv preprint arXiv:2402.17177* (2024)
40. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild.
In: *Proceedings of the IEEE International Conference on Computer Vision*. pp.
3730–3738 (2015)
41. Luo, Z., Chen, D., Zhang, Y., Huang, Y., Wang, L., Shen, Y., Zhao, D., Zhou,
J., Tan, T.: Videofusion: Decomposed diffusion models for high-quality video gen-
eration. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and*
Pattern Recognition. pp. 10209–10218 (2023)
42. Lustig, M., Donoho, D., Pauly, J.M.: Sparse mri: The application of compressed
sensing for rapid mr imaging. *Magnetic Resonance in Medicine: An Official Journal*
of the International Society for Magnetic Resonance in Medicine **58**(6), 1182–1195
(2007)
43. Malloy, M.L., Nowak, R.D.: Near-optimal adaptive compressed sensing. *IEEE*
Transactions on Information Theory **60**(7), 4001–4012 (2014)
44. Manor, H., Michaeli, T.: Zero-shot unsupervised and text-based audio editing using
DDPM inversion. *arXiv preprint arXiv:2402.10009* (2024)
45. Mardani, M., Song, J., Kautz, J., Vahdat, A.: A variational perspective on solving
inverse problems with diffusion models. *arXiv preprint arXiv:2305.04391* (2023)
46. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: Sdedit: Guided
image synthesis and editing with stochastic differential equations. In: *International*
Conference on Learning Representations (2021)
47. Mi, L., Wang, H., Meirovitch, Y., Schalek, R., Turaga, S.C., Lichtman, J.W.,
Samuel, A.D., Shavit, N.: Learning guided electron microscopy with active acquisi-
tion. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI*
2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings,
Part V 23. pp. 77–87. Springer (2020)
48. Mullainathan, S., Obermeyer, Z.: Solving medicine’s data bottleneck: Nightingale
open science. *Nature Medicine* **28**(5), 897–899 (May 2022)
49. Murata, N., Saito, K., Lai, C.H., Takida, Y., Uesaka, T., Mitsufuji, Y., Ermon, S.:
Gibbsddrm: A partially collapsed gibbs sampler for solving blind inverse problems
with denoising diffusion restoration. *arXiv preprint arXiv:2301.12686* (2023)
50. Nakos, V., Shi, X., Woodruff, D.P., Zhang, H.: Improved algorithms for adaptive
compressed sensing. *arXiv preprint arXiv:1804.09673* (2018)
51. Pineda, L., Basu, S., Romero, A., Calandra, R., Drozdal, M.: Active mr k-space
sampling with reinforcement learning. In: *Medical Image Computing and Computer*

- Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23. pp. 23–33. Springer (2020)
52. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10684–10695 (June 2022)
 53. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–10 (2022)
 54. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. arXiv:2104.07636 (2021)
 55. Schlemper, J., Caballero, J., Hajnal, J.V., Price, A.N., Rueckert, D.: A deep cascade of convolutional neural networks for dynamic mr image reconstruction. *IEEE transactions on Medical Imaging* **37**(2), 491–503 (2017)
 56. Seeger, M.W., Nickisch, H.: Large scale bayesian inference and experimental design for sparse linear models. *SIAM Journal on Imaging Sciences* **4**(1), 166–199 (2011)
 57. Serengil, S.I., Ozpinar, A.: Lightface: A hybrid deep face recognition framework. In: 2020 Innovations in Intelligent Systems and Applications Conference (ASYU). pp. 23–27. IEEE (2020). <https://doi.org/10.1109/ASYU50717.2020.9259802>, <https://doi.org/10.1109/ASYU50717.2020.9259802>
 58. Shen, Z., Wang, Y., Wu, D., Yang, X., Dong, B.: Learning to scan: A deep reinforcement learning approach for personalized scanning in ct imaging. arXiv preprint arXiv:2006.02420 (2020)
 59. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: International Conference on Machine Learning. pp. 2256–2265. PMLR (2015)
 60. Song, J., Vahdat, A., Mardani, M., Kautz, J.: Pseudoinverse-guided diffusion models for inverse problems. In: International Conference on Learning Representations (ICLR) (May 2023)
 61. Song, Y., Shen, L., Xing, L., Ermon, S.: Solving inverse problems in medical imaging with score-based generative models. In: International Conference on Learning Representations (2023)
 62. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2020)
 63. Vahdat, A., Kreis, K., Kautz, J.: Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems* **34**, 11287–11302 (2021)
 64. Van Gorp, H., Huijben, I., Veeling, B.S., Pezzotti, N., Van Sloun, R.J.: Active deep probabilistic subsampling. In: International Conference on Machine Learning. pp. 10509–10518. PMLR (2021)
 65. Van Veen, D., Jalal, A., Soltanolkotabi, M., Price, E., Vishwanath, S., Dimakis, A.G.: Compressed sensing with deep image prior and learned regularization. arXiv preprint arXiv:1806.06438 (2018)
 66. Wang, C., Shang, K., Zhang, H., Zhao, S., Liang, D., Zhou, S.K.: Active ct reconstruction with a learned sampling policy. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 7226–7235 (2023)
 67. Wang, S., Cheng, H., Ying, L., Xiao, T., Ke, Z., Zheng, H., Liang, D.: Deep-complexmri: Exploiting deep residual network for fast parallel mr imaging with complex convolution. *Magnetic Resonance Imaging* **68**, 136–147 (2020)

68. Wang, S., Su, Z., Ying, L., Peng, X., Zhu, S., Liang, F., Feng, D., Liang, D.: Accelerating magnetic resonance imaging via deep learning. In: IEEE 13th International Symposium on Biomedical Imaging (ISBI). pp. 514–517. IEEE (2016)
69. Wang, T., Xia, W., Lu, J., Zhang, Y.: A review of deep learning ct reconstruction from incomplete projection data. *IEEE Transactions on Radiation and Plasma Medical Sciences* (2023)
70. Wang, T., Lucka, F., van Leeuwen, T.: Sequential experimental design for x-ray ct using deep reinforcement learning. *arXiv preprint arXiv:2307.06343* (2023)
71. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. *arXiv preprint arXiv:2212.00490* (2022)
72. Wu, W., Wang, Y., Liu, Q., Wang, G., Zhang, J.: Wavelet-improved score-based generative model for medical imaging. *IEEE transactions on medical imaging* (2023)
73. Wu, Y., Rosca, M., Lillicrap, T.: Deep compressed sensing. In: International Conference on Machine Learning. pp. 6850–6860. PMLR (2019)
74. Yan, K., Wang, X., Lu, L., Summers, R.M.: Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *Journal of medical imaging* **5**(3), 036501–036501 (2018)
75. Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M.J., Defazio, A., Stern, R., Johnson, P., Bruno, M., Parente, M., Geras, K.J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M., Vincent, P., Yakubova, N., Pinkerton, J., Wang, D., Owens, E., Zitnick, C.L., Recht, M.P., Sodickson, D.K., Lui, Y.W.: fastMRI: An open dataset and benchmarks for accelerated MRI (2018)
76. Zhang, P., Wang, F., Xu, W., Li, Y.: Multi-channel generative adversarial network for parallel magnetic resonance image reconstruction in k-space. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I. pp. 180–188. Springer (2018)
77. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
78. Zhang, Z., Romero, A., Muckley, M.J., Vincent, P., Yang, L., Drozdal, M.: Reducing uncertainty in undersampled mri reconstruction with active acquisition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2049–2058 (2019)
79. Zhu, B., Liu, J.Z., Cauley, S.F., Rosen, B.R., Rosen, M.S.: Image reconstruction by domain-transform manifold learning. *Nature* **555**(7697), 487–492 (2018)