

AMBIENT DIFFUSION POSTERIOR SAMPLING: SOLVING INVERSE PROBLEMS WITH DIFFUSION MODELS TRAINED ON CORRUPTED DATA

Anonymous authors

Paper under double-blind review

ABSTRACT

We provide a framework for solving inverse problems with diffusion models learned from linearly corrupted data. Firstly, we extend the Ambient Diffusion framework to enable training directly from measurements corrupted in the Fourier domain. Subsequently, we train diffusion models for MRI with access only to Fourier subsampled multi-coil measurements at acceleration factors $R=2, 4, 6, 8$. Secondly, we propose *Ambient Diffusion Posterior Sampling* (A-DPS), a reconstruction algorithm that leverages generative models pre-trained on one type of corruption (e.g. image inpainting) to perform posterior sampling on measurements from a different forward process (e.g. image blurring). For MRI reconstruction in high acceleration regimes, we observe that A-DPS models trained on subsampled data are better suited to solving inverse problems than models trained on fully sampled data. We also test the efficacy of A-DPS on natural image datasets (CelebA, FFHQ, and AFHQ) and show that A-DPS can sometimes outperform models trained on clean data for several image restoration tasks in both speed and performance.

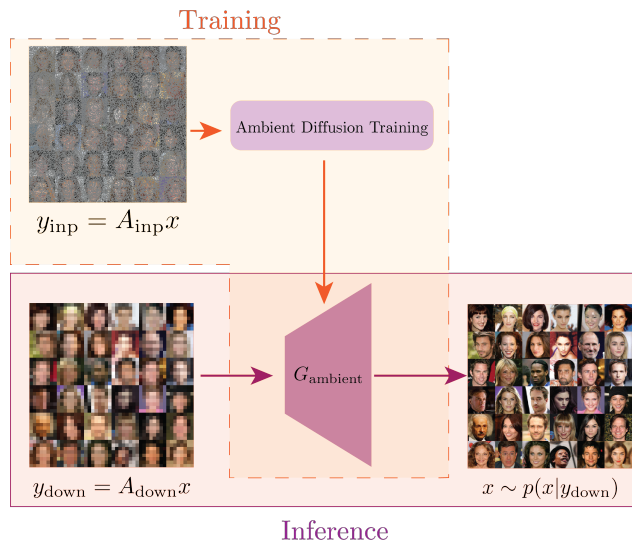


Figure 1: Ambient Diffusion Posterior Sampling (Ambient DPS). During training, we only have access to linearly corrupted data from a forward operator A_{train} . We use the Ambient Diffusion framework to learn a generative model, G_{ambient} , for the uncorrupted distribution, $p(x_0)$. At inference time, we sample from the posterior distribution $p(x_0|y_{A_{\text{inf}}})$, for measurements y_{inf} coming from a different forward operator, A_{inf} .

1 INTRODUCTION

For some applications, it is expensive or impossible to acquire fully observed or uncorrupted data (Collaboration et al., 2019; Gao et al., 2023; Yaman et al., 2020) but possible to acquire partially observed

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

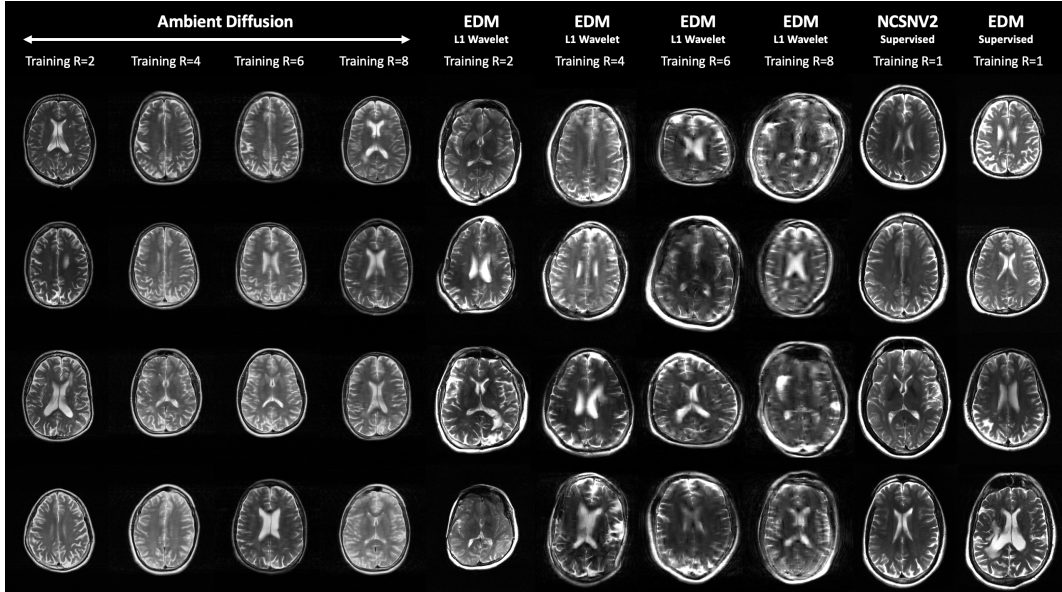


Figure 2: Prior samples from diffusion models trained on MRI scans. Columns 1 – 4: Diffusion models trained on subsampled MRI scans at acceleration factors $R = 2, 4, 6, 8$, using the Ambient Diffusion framework extended for Fourier subsampled training. Columns 5 – 8: EDM models trained with L1-wavelet reconstructions of subsampled scans at $R = 2, 4, 6, 8$. Column 9: NCSNV2 trained with fully sampled scans. Column 10: EDM trained with fully sampled scans. We observe that Ambient Diffusion models consistently produce high-quality and realistic MRI scans even in high acceleration regimes.

samples. Further, in some cases, it may be desirable to train generative models with noisy or corrupted data since that reduces memorization of the training set (Daras et al., 2023b; Carlini et al., 2023; Somepalli et al., 2022). Prior works have shown how to train Generative Adversarial Networks (GANs) (Bora et al., 2018; Cole et al., 2021), flow models (Kelkar et al., 2023) and restoration models (Lehtinen et al., 2018; Krull et al., 2019; Yaman et al., 2020; Millard & Chiew, 2023) with corrupted training data. More recently, there has been a shift towards training *diffusion generative models* given corrupted data (Daras et al., 2024; 2023a;b; Aali et al., 2023; Kawar et al., 2023; Cui et al., 2022; Kim & Ye, 2021). What remains unexplored is how to use generative models trained on a certain type of corruption (e.g. inpainted data) to solve inverse problems that arise from a different forward process (e.g. blurring).

We propose the first framework to solve inverse problems with Ambient Diffusion models (Daras et al., 2023b). These models are trained using only access to linear measurements and they estimate the *ambient score*, i.e. how to best reconstruct given an input with a corrupted linear forward operator *corrupted noisy input*. We show how to use these models for solving linear inverse problems outside their training distribution. Our experiments on datasets of natural images and multi-coil MRI show something surprising: Ambient Models can outperform (in the high corruption regime) models trained on clean data while being substantially faster. Our algorithm extends Diffusion Posterior Sampling (Chung et al., 2023) to Ambient Diffusion models. **Our contributions:**

- We propose *Ambient Diffusion Posterior Sampling* (Ambient DPS), an algorithm that uses diffusion models trained on linearly corrupted data as priors for solving inverse problems with arbitrary linear measurement models.
- We extend the Ambient Diffusion training framework to train models using Fourier subsampled measurements. Then, we train Ambient Diffusion models on subsampled multi-coil MRI scans at various retrospective acceleration factors ($R=2, 4, 6, 8$); we observe that models trained on subsampled data are better priors for solving inverse problems in the high acceleration regime.
- We use pre-trained Ambient Diffusion models to solve inverse problems (compressed sensing, super-resolution) on natural image datasets (CelebA, FFHQ, AFHQ) and show that they can even outperform models trained on clean data in the high corruption regime.

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128

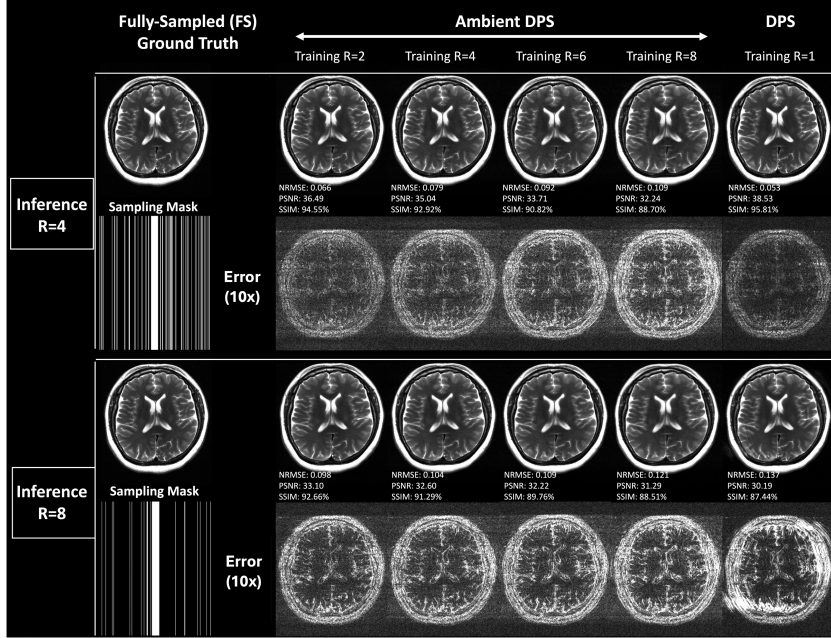


Figure 3: Posterior sampling reconstructions for MRI scans using models trained on Fourier subsampled data at various acceleration factors (Ambient DPS, columns 2 – 5) and a model trained on clean data (FS-DPS (Chung et al., 2023), column 6). Rows 1 and 3 show reconstructions at $R = 4$ and $R = 8$, respectively, while rows 2 and 4 display the difference to the ground truth on a $10\times$ scale. At high inference acceleration ($R=8$), *Ambient DPS*, outperforms FS-DPS, despite that the underlying models were trained solely on corrupted data.

2 METHOD

2.1 BACKGROUND AND NOTATION

Diffusion Posterior Sampling (DPS). Diffusion models are typically trained to reconstruct a clean image $\mathbf{x}_0 \sim p_0(\mathbf{x}_0)$ from a noisy observation $\mathbf{x}_t = \mathbf{x}_0 + \sigma_t \boldsymbol{\eta}$, $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, I)$. Despite the simplicity of the training objective, diffusion models can approximately sample from $p(\mathbf{x})$ by running a discretized version of the Stochastic Differential Equation (Song et al., 2020; Anderson, 1982):

$$d\mathbf{x} = -2\dot{\sigma}_t(\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] - \mathbf{x}_t)dt + g(t)d\mathbf{w}, \quad (2.1)$$

where \mathbf{w} is the standard Wiener process and $\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]$ is estimated by the network. Given measurement $\mathbf{y}_{\text{inf}} = A_{\text{inf}}\mathbf{x}_0$, one can sample from the posterior $p(\mathbf{x}_0|\mathbf{y}_{\text{inf}})$ by running the process:

$$d\mathbf{x} = -2\dot{\sigma}_t\sigma_t \left(\frac{\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] - \mathbf{x}_t}{\sigma_t} + \underbrace{\nabla \log p(\mathbf{y}_{\text{inf}}|\mathbf{x}_t)}_{\text{likelihood term}} \right) dt + g(t)d\mathbf{w}. \quad (2.2)$$

For most forward operators it is intractable to write the likelihood in closed form. Hence, several approximations have been proposed to use diffusion models for inverse problems (Chung et al., 2023; Kawar et al.; Jalal et al., 2021; Song et al., 2021; Chung et al., 2022; Feng et al., 2023; Graikos et al., 2022). One of the simplest and most effective approximations is Diffusion Posterior Sampling (DPS) (Chung et al., 2023). DPS estimates \mathbf{x}_0 using \mathbf{x}_t and uses the conditional likelihood $p(\mathbf{y}_{\text{inf}}|\hat{\mathbf{x}}_0)$ instead of the intractable term, i.e. DPS approximates $p(\mathbf{y}_{\text{inf}}|\mathbf{x}_t)$ with $p(\mathbf{y}_{\text{inf}}|\mathbf{x}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$, where γ_t is a tunable guidance parameter. The update rule becomes:

$$d\mathbf{x} = -2\dot{\sigma}_t\sigma_t \left(\frac{\mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] - \mathbf{x}_t}{\sigma_t} + \gamma_t \nabla_{\mathbf{x}_t} \log p(\mathbf{y}_{\text{inf}}|\mathbf{x}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t]) \right) dt + g(t)d\mathbf{w}, \quad (2.3)$$

Ambient Diffusion. In some settings we do not have a large training set of clean data but we have access to a large set of lossy measurements that we would like to leverage to train a diffusion model for the clean distribution.

159
160
161

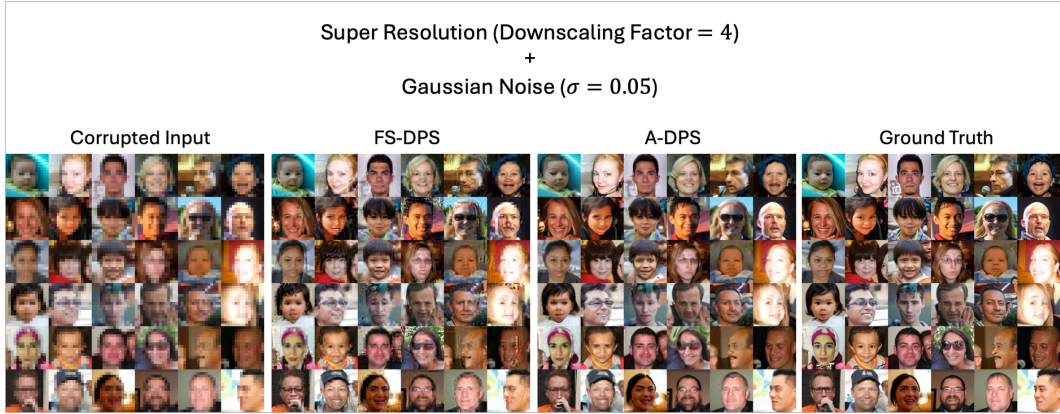


Figure 4: Posterior sampling reconstructions for FFHQ, showing: (1) Corrupted Input: Examples from FFHQ dataset corrupted using Resolution Downscaling ($4\times$) and Gaussian Noise ($\sigma = 0.05$), (2) FS-DPS: Reconstruction using Diffusion Posterior Sampling (DPS), with a diffusion model trained with clean fully-sampled data ($p = 0.0$), (3) A-DPS: Reconstruction using Ambient-DPS, with a diffusion model trained on randomly inpainted data with erasure probability ($p = 0.6$), (4) Ground Truth: Original uncorrupted examples from FFHQ dataset. We observe that *Ambient DPS*, provides better reconstructions (qualitatively) even though the underlying models were trained on corrupted data ($p = 0.6$).

The authors of Daras et al. (2023b) consider the setting of having access to linearly corrupted data $\{\mathbf{y}_0 = A_{\text{train}}\mathbf{x}_0, A_{\text{train}}\}$, where the distribution of A_{train} , denoted as $p(A_{\text{train}})$, is assumed to be known. The ultimate goal is to learn the best restoration model given a linearly corrupted noisy input $\mathbf{y}_{t,\text{train}} = A_{\text{train}}(\mathbf{x}_0 + \sigma_t\boldsymbol{\eta})$, at all noise levels t . The authors of Daras et al. (2023b) form further corrupted iterates $\tilde{\mathbf{y}}_{t,\text{train}} = \tilde{A}_{\text{train}}(\mathbf{x}_0 + \sigma_t\boldsymbol{\eta})$ using a matrix \tilde{A}_{train} (that is a perturbation of the given matrix A_{train}) and train with the following objective:

$$J(\theta) = \mathbb{E}_{\mathbf{x}_0, \mathbf{x}_t, A_{\text{train}}, \tilde{A}_{\text{train}}} \left[\left\| A_{\text{train}} \mathbf{h}_\theta(\tilde{\mathbf{y}}_{0,\text{train}}, \tilde{A}_{\text{train}}) - \mathbf{y}_{t,\text{train}} \right\|^2 \right], \quad (2.4)$$

that provably learns $\mathbb{E}[\mathbf{x}_0 | \tilde{A}_{\text{train}}, \tilde{\mathbf{y}}_{t,\text{train}}]$ as long as the matrix $\mathbb{E}[A_{\text{train}}^T A_{\text{train}} | \tilde{A}_{\text{train}}]$ is full-rank. In certain cases, it is possible to introduce *minimal* additional corruption and satisfy this condition. For example, if A_{train} is a random inpainting matrix (i.e. $A_{ij} \sim \text{Be}(1 - p)$), then \tilde{A}_{train} can be formed by taking A_{train} and erasing additional pixels with any non-zero probability $\delta > 0$.

Multi-coil Magnetic Resonance Imaging. MRI is a prototypical use case for a framework that can learn generative models from linearly corrupted data, as in many cases it is not feasible to collect a large training set of fully sampled data (Tibrewala et al., 2023; Desai et al., 2022; Tariq et al.). In settings such as 3D+time dynamic contrast-enhanced MRI (Zhang et al., 2015) it is impossible to collect fully sampled data due to the time-varying dynamics of the contrast agent (Zhang et al.).

In the multi-coil MRI setting, the acquisition involves collecting measurements of an image directly in the spatial frequency, known as k -space, from a set of spatially localized coils. Mathematically, there are N_c coils, each of which gives measurements:

$$\mathbf{z}_{\mathbf{x},i} = P\mathcal{F}S_i\mathbf{x} + \mathbf{w}_i, \quad i \in [N_c], \quad (2.5)$$

where \mathbf{x} is the (complex-valued) image of interest, S_i represents the coil-sensitivity profile of the i^{th} coil, \mathcal{F} is the Fourier transform, P represents the Fourier subsampling operator and \mathbf{w}_i is complex-valued Gaussian i.i.d noise. We assume the noiseless case and we point the reader to Daras et al. (2023a); Aali et al. (2023); Kawar et al. (2023) for approaches that handle the noisy case.

For the discrete approximation of the continuous signal as an image $\mathbf{x} \in \mathbb{C}^m$, the composition of P, \mathcal{F}, S_i can be written as a matrix $A_i \in \mathbb{C}^{m \times n}$, where the number of measurements, m , depends on the subsampling operator P . It is common to denote with R the ratio $\frac{n}{m}$, which is known as the acceleration factor. At inference time (i.e., for a new patient), we typically want to acquire data with a high acceleration factor because this reduces scan time and patient discomfort.

Due to the time required to collect k -space measurements, it is often not possible to acquire fully sampled k -space. Hence, training cannot rely on a fully sampled image to guide reconstruction

quality as is done in the fully supervised setting (Aggarwal et al., 2019; Hammernik et al., 2018). If two independent measurements are available for the same underlying signal, restoration models can be trained without having access to uncorrupted data (Lehtinen et al., 2018; Gan et al., 2023). If only one measurement is available (as in our setting), some end-to-end techniques use a loss on the measurement domain by partitioning the training measurements for each sample into (1) measurements for reconstruction, and (2) measurements for applying the loss function (Yaman et al., 2020). Other approaches leverage structure in the MRI acquisition to learn from limited-resolution data (Wang et al., 2023). More recently, works have begun leveraging signal set properties such as group invariance to assist in learning from subsampled measurements for a variety of inverse problems (Tachella et al., 2022b; Scanvic et al., 2023; Chen et al., 2022a; Tachella et al., 2022a). The idea of additional corruption (as we do in our work), is more related to the Noisier2Noise framework (Moran et al., 2020) which has been initially developed for denoising. Millard & Chiew (2023) extended this idea to the MRI setting to learn models for MRI restoration without access to clean data. All these approaches learn a restoration model without access to reference data. However, as they are inherently end-to-end methods, their performance on out-of-distribution tasks (e.g., due to different acquisition trajectories) is known to degrade (Jalal et al., 2021; Zach et al., 2022).

2.2 AMBIENT DIFFUSION FOR MRI

The MRI acquisition process results in linearly corrupted measurements and thus it should be possible to use the Ambient Diffusion framework to learn from corrupted observations. Yet, we identified three important changes that differ from the setting studied by Daras et al. (2023b): i) the inpainting happens in the Fourier Domain, ii) the inpainting has structure, i.e. whole vertical lines in the spatial Fourier are either observed or not observed (instead of masking random pixels) and, iii) the image is measured from an array of spatially varying receive coils. In what follows, we show how to account for these factors and apply the Ambient Diffusion framework to MRI data.

First, as in Ambient Diffusion, we will corrupt further the given measurements, but in our case, the additional corruption will happen in the Fourier space. We create further corrupted measurements for each coil by decreasing the Fourier subsampling ratio, i.e. we create iterates:

$$\tilde{z}_{x,i} = \tilde{P} \mathcal{F} S_i x. \quad (2.6)$$

This is possible to do because we can just subsample the available data $z_{x,i}$. To further corrupt the given measurements, we keep the inpainting structure that there is in the measurements, i.e. we delete (at random) 2-D lines in Fourier space (instead of pixel masking as in Daras et al. (2023b)). Due to the multiplicity of coils, we combine their measurements to form a crude estimator of x by taking the adjoint of \tilde{A}_{train} : $\tilde{y}_{\text{train}} = \sum_i S_i^H \mathcal{F}^{-1}(\tilde{z}_{x,i})$. Notice that for a discrete signal $x \in \mathbb{C}^n$, all these operations are linear and hence \tilde{y} can be written as:

$$\tilde{y}_{\text{train}} = \underbrace{\left(\sum_i S_i^H \mathcal{F}^{-1} \tilde{P} \mathcal{F} S_i \right)}_{\tilde{A}_{\text{train}}} x. \quad (2.7)$$

As in Ambient Diffusion, we will train the network to predict the corresponding signal before the additional corruption, i.e. our target will be:

$$y_{\text{train}} = \underbrace{\left(\sum_i S_i^H \mathcal{F}^{-1} P \mathcal{F} S_i \right)}_{A_{\text{train}}} x. \quad (2.8)$$

We are now ready to state our main Theorem.

Theorem 2.1 ((Informal)). *Let $y_{t,\text{train}}, \tilde{y}_{t,\text{train}}$ represent the noisy versions of $y_{\text{train}}, \tilde{y}_{\text{train}}$ respectively, i.e.:*

$$\begin{cases} y_{t,\text{train}} = A_{\text{train}}(x + \sigma_t \eta) \\ \tilde{y}_{t,\text{train}} = \tilde{A}_{\text{train}}(x + \sigma_t \eta) \end{cases}. \quad (2.9)$$

Then, the minimizer of the objective:

$$J(\theta) = \mathbb{E}_{y_{0,\text{train}}, \tilde{y}_{t,\text{train}}, A_{\text{train}}, \tilde{P}} \left[\left\| A_{\text{train}} \mathbf{h}_{\theta}(\tilde{y}_{t,\text{train}}, \tilde{P}) - y_{0,\text{train}} \right\|^2 \right], \quad (2.10)$$

is: $\mathbf{h}_{\theta}(\tilde{y}_{t,\text{train}}, \tilde{P}) = \mathbb{E}[x_0 | \tilde{y}_{t,\text{train}}, \tilde{P}]$.

Proof overview. The formal proof of this Theorem is given in the Appendix. We provide a sketch of the proof here. The techniques are based on Theorem 4.2 of Ambient Diffusion (Daras et al., 2023b). In Theorem 4.2, the condition that needs to be satisfied is that $\mathbb{E}[A_{\text{train}}^T A_{\text{train}} \tilde{A}_{\text{train}}]$ is full-rank. Our objective is slightly different because the network doesn't see the whole forward operator \tilde{A}_{train} , but only part of it, i.e. the matrix \tilde{P} . By following the steps of the Ambient Diffusion proof, we arrive at the necessary condition for our case, which is to prove that $\mathbb{E}[A_{\text{train}} | \tilde{P}]$ is full-rank.

We start the argument by noting that $\mathbb{E}[P | \tilde{P}]$ is full-rank. Intuitively, this is because there is a non-zero probability of observing any Fourier coefficient and hence any deleted Fourier co-efficient could be due to the extra corruption introduced by \tilde{P} . Since \mathcal{F} is an invertible matrix, then also $\mathbb{E}[\mathcal{F}^{-1} P \mathcal{F} | \tilde{P}]$ is full-rank. Finally, we use the properties of the sensitivity masks to show that $\mathbb{E}[\sum_i S_i^H \mathcal{F}^{-1} P \mathcal{F} S_i | \tilde{P}]$ is also full-rank and this completes the proof.

2.3 AMBIENT DIFFUSION POSTERIOR SAMPLING

DPS requires access to $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ to approximately sample from $p(\mathbf{x}_0 | \mathbf{y}_{\text{inf}})$. Since Ambient Diffusion models can only work with corrupted inputs, we propose the following update rule instead:

$$d\mathbf{x} = -2\dot{\sigma}_t \sigma_t \left(\underbrace{\frac{\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{t,\text{train}}, \tilde{A}_{\text{train}}] - \mathbf{x}_t}{\sigma_t}}_{\text{Ambient Score}} + \gamma_t \nabla_{\mathbf{x}_t} \log p(\mathbf{y}_{\text{inf}} | \mathbf{x}_0 = \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{t,\text{train}}, \tilde{A}_{\text{train}}]) \right) dt + g(t) d\mathbf{w}, \quad (2.11)$$

for a fixed $\tilde{A}_{\text{train}} \sim p(\tilde{A}_{\text{train}})$ ¹. Comparing this to the DPS update rule (E.q. 2.3), all the $\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]$ terms have been replaced with their ambient counterparts, i.e. with $\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{t,\text{train}}, \tilde{A}_{\text{train}}]$. The latter is approximated by a neural network that is trained with Ambient Diffusion.

We term our approximate sampling algorithm for solving inverse problems with diffusion models learned from corrupted data **Ambient DPS (A-DPS)**. We remark that similar to DPS, the proposed algorithm is an approximation to sampling from the true posterior distribution. Similarly to Theorem 1 of the DPS paper, for a given measurement \mathbf{y}_{inf} at noise level $\sigma_{\mathbf{y}}$ and a given matrix A_{inf} , one can upper-bound the Jensen gap of A-DPS as follows:

$$\mathcal{J} \leq \frac{n}{\sqrt{2\pi\sigma_{\mathbf{y}}^2}} \exp(-1/2\sigma_{\mathbf{y}}^2) \|A_{\text{inf}}\| \int_{\mathbf{x}_0} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\| p(\mathbf{x}_0 | \mathbf{x}_t) d\mathbf{x}_0, \quad (2.12)$$

where $\hat{\mathbf{x}}_0 = \mathbb{E}[\mathbf{x}_0 | \tilde{A}_{\text{train}} \mathbf{x}_t, \tilde{A}_{\text{train}}]$, for a fixed \tilde{A}_{train} sampled from $p(\tilde{A}_{\text{train}})$. We omit the proof of this proposition since it follows the exact same steps as the proof of Theorem 1 in the DPS paper.

2.4 AMBIENT DIFFUSION FOR IN-DOMAIN RECONSTRUCTIONS.

Ambient DPS can be used to solve *any* inverse problem for which the forward operator is known, using a diffusion model trained on linearly corrupted data of some form. It is important to underline though that if the inverse problem that we want to solve at inference time has the same forward operator as the one that was used for the training measurements, then we can use the Ambient Diffusion as a supervised restoration model. This is because Ambient Diffusion models are trained to estimate $\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{\text{train}}, \tilde{A}_{\text{train}}]$, and hence if A_{inf} comes from the same distribution as A_{train} , then, the one-step prediction of the model, Ambient One Step (A-OS), is the Mean Squared Error (MSE) minimizer. Similarly, if we are interested in unconditional generation, we can simply run Ambient DPS without the likelihood term.

¹The Ambient MRI models take as input \tilde{P} instead of \tilde{A} , as in Equation 2.10.

Table 1: Unconditional and Conditional Sampling FID for models trained on MRI data at different acceleration factors.

Training Data	Training Method	Unconditional FID ↓
$R = 1$	EDM NCSNV2	10.41 13.20
$R = 2$	L1-EDM Ambient Diffusion	18.55 30.34
$R = 4$	L1-EDM Ambient Diffusion	27.64 32.31
$R = 6$	L1-EDM Ambient Diffusion	51.43 31.50
$R = 8$	L1-EDM Ambient Diffusion	102.98 48.15

Training Data	Inference Data	Reconstruction Method	Conditional FID ↓
$R = 1$	$R = 8$	FS-DPS	4.16
$R = 2$	$R = 8$	A-DPS (<i>our method</i>)	0.71
$R = 4$			0.87
$R = 6$			1.29
$R = 8$			2.00

3 EXPERIMENTS

3.1 AMBIENT MRI DIFFUSION MODELS

In this section, we extend Ambient Diffusion to the multi-coil Fourier subsampled MRI setting (as detailed in Section 2.2) and we train our own MRI models from scratch. We give details about our dataset preparation in Section B in the Appendix.

Sampling masks. We retrospectively subsample the k -space training data by applying randomly subsampled masks that correspond to acceleration factors $R \in \{2, 4, 6, 8\}$. The sampling masks include fully sampled vertical (readout) lines corresponding to the observed Fourier coefficients (phase encodes). We always sample the central 20 lines for autocalibration. To form further corrupted measurements, we randomly remove additional lines such that we create measurements at acceleration factor $R + 1$. For example, we take data at $R = 2$, we create further corrupted measurements at $R = 3$ and we train the model to predict the clean image by measuring its error with the available data at $R = 2$ given its prediction for the $R = 3$ corrupted input.

Comparison Models. We train one model for each acceleration factor $R \in \{1, 2, 4, 6, 8\}$. The $R = 1$ model is trained on clean data (no extra corruption) based on the EDM approach (Karras et al., 2022). The $R > 1$ models are trained with our modified Ambient Diffusion framework. While we focus on EDM for fully sampled training, we also train an NCSNV2 model for $R = 1$ following the approach in Jalal et al. (2021). As baselines, we also train EDM models after L1-Wavelet compressed sensing reconstruction (L1-EDM) of the training set at each acceleration factor (Lustig et al., 2007).

Unconditional generation evaluation. We evaluate the unconditional generation performance of each model and show prior samples for Ambient Diffusion and EDM models at $R = 1$ in Figure 2. Similar to the results in Daras et al. (2023b), we observe that the Ambient Diffusion samples are qualitatively similar to EDM models at low acceleration, and become slightly blurry at higher accelerations. Notably, the samples generated by Ambient Diffusion at $R = 8$ have no residual aliasing artifacts, in contrast to the samples generated from L1-Wavelet reconstructions.

To quantify unconditional sample quality, we follow the approach in Bendel et al. (2022) and calculate FID scores from 100 samples using a pre-trained VGG network. Table 1, shows FID scores for the diffusion models. While there is an increase in FID for the Ambient models, we see that those trained at higher acceleration factors outperform L1-EDM models trained on L1-Wavelet reconstructions.

Training Method	Training Data	Reconstruction Method	$R = 2$			$R = 4$			$R = 6$			$R = 8$		
			NRMSE	SSIM	PSNR	NRMSE	SSIM	PSNR	NRMSE	SSIM	PSNR	NRMSE	SSIM	PSNR
Supervised	$R = 1$	FS-DPS	0.034	97.83	43.70	0.065	95.20	38.17	0.121	90.60	32.99	0.179	85.21	29.46
		FS-ALD	0.059	94.55	40.17	0.092	90.71	36.13	0.122	88.53	33.62	0.157	86.11	31.22
		L1-CS	0.092	91.10	36.13	0.143	82.08	31.58	0.217	77.07	27.74	0.259	75.04	26.09
		MoDL ($R = 2$)	0.036	97.56	43.03	0.076	94.34	36.65	0.132	87.56	31.84	0.182	81.13	29.01
		MoDL ($R = 4$)	0.041	96.67	41.95	0.066	95.19	37.85	0.095	93.09	34.62	0.124	90.81	32.31
		MoDL ($R = 6$)	0.045	95.37	41.04	0.067	94.83	37.65	0.090	93.41	35.04	0.111	91.87	33.21
		MoDL ($R = 8$)	0.050	94.13	40.24	0.079	90.30	36.20	0.096	90.09	34.46	0.113	91.08	33.08
Self-Supervised	$R = 2$	A-DPS (<i>our method</i>)	0.058	94.88	38.94	0.074	93.99	36.82	0.093	92.90	34.84	0.112	91.63	33.19
		A-ALD (<i>our method</i>)	0.058	94.84	38.89	0.074	93.99	36.82	0.093	92.83	34.79	0.114	91.52	33.07
		A-OS	0.076	94.93	36.42	0.111	91.81	33.19	0.159	88.04	30.04	0.196	84.85	28.26
		L1-DPS	0.036	97.54	43.20	0.088	93.35	35.64	0.151	87.57	30.94	0.205	81.84	28.09
		SSDU	0.057	94.80	39.25	0.134	88.55	31.76	0.199	82.09	28.23	0.247	77.53	26.35
		1D-Partitioned SSDU	0.076	94.94	36.38	0.083	90.47	35.61	0.107	86.14	33.45	0.128	82.82	31.9
	$R = 4$	A-DPS (<i>our method</i>)	0.072	93.31	37.0	0.084	92.82	35.67	0.099	91.98	34.24	0.113	91.06	33.05
		A-ALD (<i>our method</i>)	0.071	93.32	37.03	0.084	92.77	35.64	0.099	91.98	34.24	0.113	91.13	33.08
		A-OS	0.108	91.25	31.95	0.105	91.83	32.81	0.136	89.90	31.60	0.162	87.74	30.43
		L1-DPS	0.040	96.59	42.45	0.112	89.98	33.58	0.167	84.68	29.98	0.211	80.26	27.77
		SSDU	0.072	87.84	37.20	0.088	92.82	35.36	0.131	89.74	31.81	0.167	87.01	29.72
		1D-Partitioned SSDU	0.099	93.45	34.05	0.103	89.79	33.71	0.114	85.96	32.84	0.127	82.96	31.91
	$R = 6$	A-DPS (<i>our method</i>)	0.084	91.85	35.63	0.094	91.47	34.62	0.106	90.87	33.60	0.118	90.20	32.69
		A-ALD (<i>our method</i>)	0.084	91.85	35.65	0.095	91.42	34.56	0.108	90.74	33.48	0.119	90.12	32.60
		A-OS	0.127	89.83	31.95	0.116	90.25	32.81	0.134	89.45	31.60	0.153	88.16	30.43
		L1-DPS	0.043	95.94	41.73	0.112	88.67	32.72	0.175	83.68	29.49	0.215	79.75	27.59
		SSDU	0.075	89.72	36.84	0.105	91.77	33.82	0.138	89.42	31.34	0.166	87.29	29.76
		1D-Partitioned SSDU	0.081	94.44	35.81	0.103	89.60	33.73	0.121	85.52	32.37	0.134	82.47	31.48
	$R = 8$	A-DPS (<i>our method</i>)	0.099	90.51	34.15	0.107	90.47	33.46	0.117	89.94	32.71	0.126	89.50	32.10
		A-ALD (<i>our method</i>)	0.100	90.47	34.11	0.107	90.44	33.46	0.118	89.81	32.63	0.127	89.29	32.00
		A-OS	0.135	86.46	31.46	0.136	87.08	31.46	0.152	86.95	30.53	0.164	86.54	29.86
		L1-DPS	0.044	95.94	41.53	0.126	88.47	32.41	0.183	82.93	29.07	0.223	78.94	27.24
		SSDU	0.118	77.19	32.95	0.122	88.15	32.42	0.151	87.93	30.55	0.177	86.16	29.17
		1D-Partitioned SSDU	0.083	94.07	35.59	0.114	88.99	32.85	0.134	84.75	31.41	0.148	81.61	30.55

Table 2: Accelerated MRI reconstruction performance evaluation using Normalized Root Mean Squared Error (NRMSE \downarrow), Structural Similarity Index Measure (SSIM \uparrow), and Peak Signal-to-Noise Ratio (PSNR \uparrow) averaged across 100 test samples.

3.2 ACCELERATED MRI RECONSTRUCTION

We now evaluate the accelerated reconstruction performance of Ambient Diffusion models trained on subsampled MRI data. Our method, *A-DPS* was implemented by following Eq. 2.11 for 500 steps with likelihood weighting $\gamma_t = \frac{1}{\|\mathbf{y} - \mathbb{A}\mathbb{E}[\mathbf{x}_0 | \mathbf{y}_t, \mathbf{A}_{\text{train}}]\|_2}$. We consider the following baselines:

- **Fully-Sampled Diffusion Posterior Sampling (FS-DPS)**. For models trained on clean data, we can use the Diffusion Posterior Sampling (DPS) (Chung et al., 2023) algorithm. Specifically, we implemented the DPS algorithm with a Diffusion model trained with fully sampled ground truth images (FS-DPS). Inference was run using the update rule of Eq. 2.3, for a total of 500 steps and with $\gamma_t = \frac{1}{\|\mathbf{y} - \mathbb{A}\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]\|_2}$.
- **Fully-Sampled Annealed Langevin Dynamics (FS-ALD)**. A popular method for solving inverse problems via score-based priors is ALD (Jalal et al., 2021). We train a score-based model using the same fully sampled 10,000 samples as above and run inference for 1,300 steps. This method acts as another fully sampled comparison to our technique.
- **L1-Wavelet Compressed Sensing (L1-CS)**. We use the same L1-Wavelet reconstruction described previously as a standalone non-deep learning comparison.
- **Model Based Deep Learning Supervised (MoDL-Sup)**. We use a supervised end-to-end model using the MoDL architecture in a supervised fashion (Aggarwal et al., 2019) for each acceleration factor. For inference, we pass the undersampled data through the trained MoDL network. Further details of this baseline are mentioned in Section C in the appendix.
- **Ambient Annealed Langevin Dynamics (A-ALD)**. We utilize our Ambient Diffusion MRI models with the inverse problem solver ALD (Jalal et al., 2021) for 1,300 steps.
- **Ambient One Step (A-OS)**. Our method also admits a one-step solution. This is outlined in section 2.4 by noting that we train our model to estimate $\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{\text{train}}, \tilde{\mathbf{A}}_{\text{train}}]$. Thus we can use our models in a one-step fashion. The performance of the one-step prediction should only be expected to be good in-distribution, i.e. when the model is evaluated at the same acceleration factor as the one during training.
- **L1-Diffusion Posterior Sampling (L1-DPS)**. We utilize the diffusion models trained on the L1-Wavelet reconstructions of data from measurement sets at $R = 2, 4, 6, 8$. For reconstruction, we use DPS for a total of 500 steps with $\gamma_t = \frac{1}{\|\mathbf{y} - \mathbb{A}\mathbb{E}[\mathbf{x}_0 | \mathbf{x}_t]\|_2}$.
- **Self-Supervised learning via Data Under-sampling (SSDU)**. We train SSDU (Yaman et al., 2020) models using the end-to-end restoration network by splitting the available

432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485

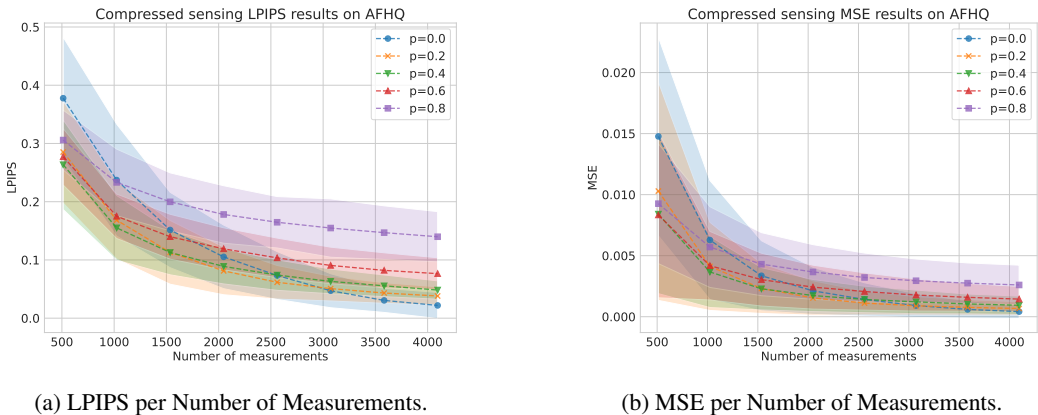


Figure 5: Compressed Sensing results, AFHQ: performance metric and standard deviation. As shown, the model trained with clean data ($p = 0.0$) only outperforms the models trained with corrupted data for more than 1000 measurements, in both LPIPS and MSE.

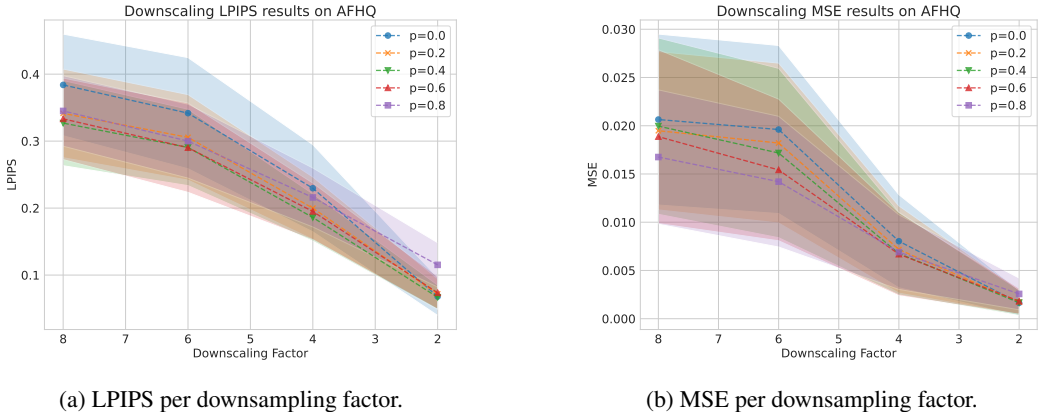


Figure 6: Super-resolution results, AFHQ: Performance metric and standard deviation. The model trained with clean data ($p = 0.0$) performs worse, except at downsampling factor 2.

measurements (Ω) of each sample y_i^Ω into training sets (Θ) and loss sets (Λ) or measurements y_i^Θ and y_i^Λ . Further SSDU training details are given in Section C.

- **1D-Partitioned SSDU.** We train 1D-Partitioned SSDU (Millard & Chiew, 2023) which stems from the SSDU framework and introduces partitioning of the sampling set to ensure each subset has a distribution similar to the original mask, substantially improving SSDU’s restoration performance. The loss function, network architecture, and training hyperparameters were set according to Millard & Chiew (2023).

Results Quantitative metrics for the listed methods are included in Table 2 under the NRMSE, SSIM and PSNR metrics respectively. Furthermore, we report feature-based metrics including LPIPS and DISTS in Table 3. We observe that at low acceleration levels ($R = 2, 4$) FS-DPS outperforms most methods, including A-DPS. This is expected FS-DPS uses models trained on clean data. However, as the validation acceleration increases to higher ratios R , A-DPS begins to outperform other models, including those trained on fully sampled data.

To gain further insight into the findings, we visualize reconstructions for a subset of the methods listed at various accelerations in Fig. 3. Here, we see that at lower accelerations ($R = 4$), just as the metrics suggest, FS-DPS outperforms A-DPS. However, at higher acceleration ($R = 8$) we can see that FS-DPS introduces significant artifacts into the reconstruction while our A-DPS reconstructions maintain a high degree of visual fidelity. We also compute the distribution metrics (Conditional FID)

between samples from the posterior and the ground-truth distribution. These FID scores for A-DPS and FS-DPS methods are shown in Table 1. Similar to our findings from posterior reconstruction metrics, we find that in the higher acceleration ($R = 8$) regime, A-DPS outperforms FS-DPS. We also include graphical trends for NRMSE, SSIM, and PSNR for A-DPS, A-OS, FS-DPS, L1-DPS, FS-ALD, L1-CS, MoDL-Sup, and SSDU in the Appendix (Figs 7, 8, 9, 10, 11).

3.3 AMBIENT DIFFUSION WITH PRE-TRAINED MODELS ON DATASETS OF NATURAL IMAGES

We use the models from the Ambient Diffusion (Daras et al., 2023b) that are trained on randomly inpainted data with different erasure probabilities. Specifically, for AFHQ we use the Ambient Models with erasure probability $p \in \{0.2, 0.4, 0.6, 0.8\}$ and for Celeb-A we use the pre-trained models with $p \in \{0.6, 0.8, 0.9\}$. We underline that all the Ambient Models have worse performance for unconditional generation than those trained with clean data (i.e. the models trained with $p = 0.0$). This work aims to explore the conditional generation performance of Ambient Models, where the conditioning is in the measurements y_{inf} , and compare it with models trained on uncorrupted data. To ensure that Ambient Models do not have an unfair advantage, we test only on restoration tasks that differ from those encountered in their training. Specifically, we use models trained on random inpainting and we evaluate Gaussian Compressed Sensing (Baraniuk, 2007) and Super Resolution.

Results. In Figure 4 we show visual examples from our experiment on the noisy super-resolution task for natural images using a model: (1) trained with fully-sampled data (FS-DPS), (2) trained with randomly inpainted data at $p = 0.6$ (A-DPS). Figure 5 presents Gaussian Compressed Sensing reconstruction results (i.e. reconstructing a signal from Gaussian random projections). We show MSE and LPIPS performance metrics for the AFHQ dataset across varying number of measurements. The results are given for models that are *trained with inpainted* images at different levels of corruption, indicated by the erasure probability p . The model trained with clean data outperforms the models trained with corrupted data when the number of measurements is high. However, as we reduce the number of measurements, Ambient Models outperform the models trained with clean data in the very low measurements regime. To the best of our knowledge, there is no known theoretical argument that explains this performance cross-over, and understanding this further is an interesting research direction. Similar results are presented in Figure 6 for the task of super-resolution at AFHQ. The model trained on clean data ($p = 0$) slightly outperforms the Ambient Models in both LPIPS and MSE for reconstructing a $2\times$ downsampled image, as expected. Yet, as the resolution decreases, there is again a cross-over in performance and models trained on corrupted data start to outperform the models trained on uncorrupted data. We include results for LPIPS and MSE for Compressed Sensing and Downsampling in FFHQ and Celeb-A in the Appendix (Figs 13, 14, 15, 16).

We compute MSE in other inverse problem settings for FFHQ: Box Inpainting, Additive Gaussian Noise, and Super Resolution (with Gaussian Noise $\sigma = 0.05$). We find that trends in this experiment are consistent with previous findings. The results across multiple downsampling factors and noise levels in the Appendix (Figs 17, 18, 19). Finally, we examine how the number of sampling steps affects the performance. The MSE results for Compressed Sensing with 4000 measurements on AFHQ are shown in Figure 12. As shown, the higher the erasure probability p during training, the better the Compressed Sensing performance of the model for low Number of Function Evaluations (NFEs). Models trained with higher corruption are faster since they require fewer steps for similar performance. For increased NFEs, the models trained on clean(er) data outperform. This result is consistent across different datasets (AFHQ, FFHQ, CelebA), reconstruction tasks (Compressed Sensing, Downsampling), and metrics (MSE, LPIPS) (Figs 20, 21, 22, 23, 24 in the Appendix).

4 CONCLUSION

We present Ambient Diffusion Posterior Sampling (A-DPS), a simple framework based on DPS for solving inverse problems with Ambient Diffusion models. We show that Ambient Diffusion models trained on corrupted data can be better suited for handling ill-posed inverse problems under severe corruption. By being exposed to corrupted training data, A-DPS exhibits robust priors that generalize better to high-acceleration regimes, unlike FS-DPS, which may overfit clean data and fail to adapt effectively to severe undersampling. Our framework fully unlocks the potential of Ambient Diffusion models that are critical in applications where access to full data is impossible or undesirable.

REFERENCES

- 540
541
542 Asad Aali, Marius Arvinte, Sidharth Kumar, and Jonathan I Tamir. Solving inverse problems with
543 score-based generative priors learned from noisy data. *arXiv preprint arXiv:2305.01166*, 2023.
- 544 Hemant K. Aggarwal, Merry P. Mani, and Mathews Jacob. Modl: Model-based deep learning
545 architecture for inverse problems. *IEEE Transactions on Medical Imaging*, 38(2):394–405, 2019.
546 doi: 10.1109/TMI.2018.2865356.
- 547 Brian D.O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their
548 Applications*, 12(3):313–326, 1982.
- 550 Richard G Baraniuk. Compressive sensing [lecture notes]. *IEEE signal processing magazine*, 24(4):
551 118–121, 2007.
- 552 Matthew Bendel, Rizwan Ahmad, and Philip Schniter. A Regularized Conditional GAN for Posterior
553 Sampling in Image Recovery Problems. *arXiv e-prints*, art. arXiv:2210.13389, October 2022. doi:
554 10.48550/arXiv.2210.13389.
- 556 Moritz Blumenthal, Christian Holme, Volkert Roeloffs, Sebastian Rosenzweig, Philip Schaten, Nick
557 Scholand, Jon Tamir, Xiaoqing Wang, and Martin Uecker. mrrecon/bart: version 0.8.00, September
558 2022. URL <https://doi.org/10.5281/zenodo.7110562>.
- 559 Ashish Bora, Eric Price, and Alexandros G Dimakis. Ambientgan: Generative models from lossy
560 measurements. In *International conference on learning representations*, 2018.
- 562 Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramer, Borja
563 Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *32nd
564 USENIX Security Symposium (USENIX Security 23)*, pp. 5253–5270, 2023.
- 565 Dongdong Chen, Julián Tachella, and Mike E. Davies. Robust equivariant imaging: A fully unsuper-
566 vised framework for learning to image from noisy and partial measurements. In *Proceedings of
567 the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5647–5656,
568 June 2022a.
- 570 Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru R Zhang. Sampling is as easy
571 as learning the score: theory for diffusion models with minimal data assumptions. *arXiv preprint
572 arXiv:2209.11215*, 2022b.
- 573 Sitan Chen, Giannis Daras, and Alex Dimakis. Restoration-degradation beyond linear diffusions:
574 A non-asymptotic analysis for ddim-type samplers. In *International Conference on Machine
575 Learning*, pp. 4462–4484. PMLR, 2023.
- 576 Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for
577 inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*,
578 35:25683–25696, 2022.
- 580 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul Ye.
581 Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International
582 Conference on Learning Representations*, 2023. URL [https://openreview.net/forum?
583 id=OnD9zGAGT0k](https://openreview.net/forum?id=OnD9zGAGT0k).
- 584 Elizabeth K. Cole, Frank Ong, Shreyas S. Vasanawala, and John M. Pauly. Fast unsupervised mri
585 reconstruction without fully-sampled ground truth data using generative adversarial networks. In
586 *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pp. 3971–
587 3980, 2021. doi: 10.1109/ICCVW54120.2021.00444.
- 588 The Event Horizon Telescope Collaboration, Kazunori Akiyama, Antxon Alberdi, Walter Alef,
589 Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John
590 Barrett, Dan Bintley, Lindy Blackburn, Wilfred Boland, Katherine L. Bouman, Geoffrey C. Bower,
591 Michael Bremer, Christiaan D. Brinkerink, Roger Brissenden, Silke Britzen, Avery E. Broderick,
592 Dominique Brogiere, Thomas Bronzwaer, Do-Young Byun, John E. Carlstrom, Andrew Chael,
593 Chi kwan Chan, Shami Chatterjee, Koushik Chatterjee, Ming-Tang Chen, Yongjun Chen, Ilje

- 594 Cho, Pierre Christian, John E. Conway, James M. Cordes, Geoffrey B. Crew, Yuzhu Cui, Jordy
595 Davelaar, Mariafelicia De Laurentis, Roger Deane, Jessica Dempsey, Gregory Desvignes, Jason
596 Dexter, Sheperd S. Doeleman, Ralph P. Eatough, Heino Falcke, Vincent L. Fish, Ed Fomalont,
597 Raquel Fraga-Encinas, William T. Freeman, Per Friberg, Christian M. Fromm, José L. Gómez,
598 Peter Galison, Charles F. Gammie, Roberto García, Olivier Gentaz, Boris Georgiev, Ciriaco Goddi,
599 Roman Gold, Minfeng Gu, Mark Gurwell, Kazuhiro Hada, Michael H. Hecht, Ronald Hesper,
600 Luis C. Ho, Paul Ho, Mareki Honma, Chih-Wei L. Huang, Lei Huang, David H. Hughes, Shiro
601 Ikeda, Makoto Inoue, Sara Issaoun, David J. James, Buell T. Jannuzi, Michael Janssen, Britton
602 Jeter, Wu Jiang, Michael D. Johnson, Svetlana Jorstad, Taehyun Jung, Mansour Karami, Ramesh
603 Karuppusamy, Tomohisa Kawashima, Garrett K. Keating, Mark Kettenis, Jae-Young Kim, Junhan
604 Kim, Jongsoo Kim, Motoki Kino, Jun Yi Koay, Patrick M. Koch, Shoko Koyama, Michael Kramer,
605 Carsten Kramer, Thomas P. Krichbaum, Cheng-Yu Kuo, Tod R. Lauer, Sang-Sung Lee, Yan-Rong
606 Li, Zhiyuan Li, Michael Lindqvist, Kuo Liu, Elisabetta Liuzzo, Wen-Ping Lo, Andrei P. Lobanov,
607 Laurent Loinard, Colin Lonsdale, Ru-Sen Lu, Nicholas R. MacDonald, Jirong Mao, Sera Markoff,
608 Daniel P. Marrone, Alan P. Marscher, Iván Martí-Vidal, Satoki Matsushita, Lynn D. Matthews, Lia
609 Medeiros, Karl M. Menten, Yosuke Mizuno, Izumi Mizuno, James M. Moran, Kotaro Moriyama,
610 Monika Moscibrodzka, Cornelia Müller, Hiroshi Nagai, Neil M. Nagar, Masanori Nakamura,
611 Ramesh Narayan, Gopal Narayanan, Iniyan Natarajan, Roberto Neri, Chunhong Ni, Aristeidis
612 Noutsos, Hiroki Okino, Héctor Olivares, Tomoaki Oyama, Feryal Özel, Daniel C. M. Palumbo,
613 Nimesh Patel, Ue-Li Pen, Dominic W. Pesce, Vincent Piétu, Richard Plambeck, Aleksandar
614 PopStefanija, Oliver Porth, Ben Prather, Jorge A. Preciado-López, Dimitrios Psaltis, Hung-Yi
615 Pu, Venkatesh Ramakrishnan, Ramprasad Rao, Mark G. Rawlings, Alexander W. Raymond,
616 Luciano Rezzolla, Bart Ripperda, Freek Roelofs, Alan Rogers, Eduardo Ros, Mel Rose, Arash
617 Roshanineshat, Helge Rottmann, Alan L. Roy, Chet Ruszczyk, Benjamin R. Ryan, Kazi L. J.
618 Rygl, Salvador Sánchez, David Sánchez-Arguelles, Mahito Sasada, Tuomas Savolainen, F. Peter
619 Schloerb, Karl-Friedrich Schuster, Lijing Shao, Zhiqiang Shen, Des Small, Bong Won Sohn,
620 Jason SooHoo, Fumie Tazaki, Paul Tiede, Remo P. J. Tilanus, Michael Titus, Kenji Toma, Pablo
621 Torne, Tyler Trent, Sascha Trippe, Shuichiro Tsuda, Ilse van Bemmelen, Huib Jan van Langevelde,
622 Daniel R. van Rossum, Jan Wagner, John Wardle, Jonathan Weintraub, Norbert Wex, Robert
623 Wharton, Maciek Wielgus, George N. Wong, Qingwen Wu, André Young, Ken Young, Ziri
624 Younsi, Feng Yuan, Ye-Fei Yuan, J. Anton Zensus, Guangyao Zhao, Shan-Shan Zhao, Ziyang
625 Zhu, Joseph R. Farah, Zheng Meyer-Zhao, Daniel Michalik, Andrew Nadolski, Hiroaki Nishioka,
626 Nicolas Pradel, Rurik A. Primiani, Kamal Souccar, Laura Vertatschitsch, and Paul Yamaguchi.
627 First m87 event horizon telescope results. iv. imaging the central supermassive black hole. *The
628 Astrophysical Journal Letters*, 875(1):L4, apr 2019. doi: 10.3847/2041-8213/ab0e85. URL
629 <https://dx.doi.org/10.3847/2041-8213/ab0e85>.
- 628 Zhuo-Xu Cui, Chentao Cao, Shaonan Liu, Qingyong Zhu, Jing Cheng, Haifeng Wang, Yanjie Zhu,
629 and Dong Liang. Self-score: Self-supervised learning on score-based models for mri reconstruction.
630 *arXiv preprint arXiv:2209.00835*, 2022.
- 631
- 632 Giannis Daras, Yuval Dagan, Alexandros G Dimakis, and Constantinos Daskalakis. Consistent
633 diffusion models: Mitigating sampling drift by learning to be consistent. *arXiv preprint
634 arXiv:2302.09057*, 2023a.
- 635 Giannis Daras, Kulin Shah, Yuval Dagan, Aravind Gollakota, Alexandros G Dimakis, and Adam
636 Klivans. Ambient diffusion: Learning clean distributions from corrupted data. *arXiv preprint
637 arXiv:2305.19256*, 2023b.
- 638
- 639 Giannis Daras, Alexandros G Dimakis, and Constantinos Daskalakis. Consistent diffusion
640 meets tweedie: Training exact ambient diffusion models with noisy data. *arXiv preprint
641 arXiv:2404.10177*, 2024.
- 642
- 643 Mauricio Delbracio and Peyman Milanfar. Inversion by direct iteration: An alternative to denoising
644 diffusion for image restoration. *arXiv preprint arXiv:2303.11435*, 2023.
- 645 Arjun D Desai, Andrew M Schmidt, Elka B Rubin, Christopher M Sandino, Marianne S Black,
646 Valentina Mazzoli, Kathryn J Stevens, Robert Boutin, Christopher Ré, Garry E Gold, Brian A
647 Hargreaves, and Akshay S Chaudhari. Skm-tea: A dataset for accelerated mri reconstruction with
dense image labels for quantitative clinical evaluation, 2022.

- 648 Berthy T Feng, Jamie Smith, Michael Rubinstein, Huiwen Chang, Katherine L Bouman, and
649 William T Freeman. Score-based diffusion models as principled priors for inverse imaging.
650 *arXiv preprint arXiv:2304.11751*, 2023.
- 651 Weijie Gan, Chunwei Ying, Parna Eshraghi Boroojeni, Tongyao Wang, Cihat Eldeniz, Yuyang
652 Hu, Jiaming Liu, Yasheng Chen, Hongyu An, and Ulugbek S Kamilov. Self-supervised deep
653 equilibrium models with theoretical guarantees and applications to mri reconstruction. *IEEE*
654 *Transactions on Computational Imaging*, 2023.
- 656 Angela F Gao, Oscar Leong, He Sun, and Katherine L Bouman. Image reconstruction without explicit
657 priors. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal*
658 *Processing (ICASSP)*, pp. 1–5. IEEE, 2023.
- 659 Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as
660 plug-and-play priors. *Advances in Neural Information Processing Systems*, 35:14715–14728, 2022.
- 662 Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas
663 Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data.
664 *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018. doi: <https://doi.org/10.1002/mrm.26977>.
665 URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.26977>.
- 666 Reinhard Heckel and Paul Hand. Deep decoder: Concise image representations from untrained
667 non-convolutional networks, 2019.
- 669 Yuyang Hu, Mauricio Delbracio, Peyman Milanfar, and Ulugbek S. Kamilov. A restoration network
670 as an implicit prior. 2023. arXiv:2310.01391.
- 671 Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alexandros G Dimakis, and Jon Tamir. Robust
672 compressed sensing mri with deep generative priors. *Advances in Neural Information Processing*
673 *Systems*, 34:14938–14954, 2021.
- 675 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
676 based generative models. *arXiv preprint arXiv:2206.00364*, 2022.
- 677 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration
678 models. In *Advances in Neural Information Processing Systems*.
- 680 Bahjat Kawar, Noam Elata, Tomer Michaeli, and Michael Elad. Gsure-based diffusion model training
681 with corrupted data. *arXiv preprint arXiv:2305.13128*, 2023.
- 682 Varun A Kelkar, Rucha Deshpande, Arindam Banerjee, and Mark A Anastasio. Ambientflow: Invert-
683 ible generative models from incomplete, noisy measurements. *arXiv preprint arXiv:2309.04856*,
684 2023.
- 686 Kwanyoung Kim and Jong Chul Ye. Noise2score: tweedie’s approach to self-supervised image
687 denoising without clean images. *Advances in Neural Information Processing Systems*, 34:864–874,
688 2021.
- 689 Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single
690 noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*
691 *recognition*, pp. 2129–2137, 2019.
- 693 Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala,
694 and Timo Aila. Noise2noise: Learning image restoration without clean data. *arXiv preprint*
695 *arXiv:1803.04189*, 2018.
- 696 Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed
697 sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195, 2007. doi:
698 <https://doi.org/10.1002/mrm.21391>. URL [https://onlinelibrary.wiley.com/doi/](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21391)
699 [abs/10.1002/mrm.21391](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.21391).
- 700 Christopher A Metzler, Arian Maleki, and Richard G Baraniuk. From denoising to compressed
701 sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016.

- 702 Charles Millard and Mark Chiew. A theoretical framework for self-supervised mr image reconstruc-
703 tion using sub-sampling via variable density noisier2noise. *IEEE Transactions on Computational*
704 *Imaging*, 9:707–720, 2023. doi: 10.1109/TCI.2023.3299212.
- 705 Nick Moran, Dan Schmidt, Yu Zhong, and Patrick Coady. Noisier2noise: Learning to denoise from
706 unpaired noisy data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
707 *Recognition*, pp. 12064–12072, 2020.
- 708 Jérémy Scanvic, Mike Davies, Patrice Abry, and Julián Tachella. Self-supervised learning for image
709 super-resolution and deblurring, 2023.
- 710 Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffu-
711 sion art or digital forgery? investigating data replication in diffusion models. *arXiv preprint*
712 *arXiv:2212.03860*, 2022.
- 713 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
714 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
715 *arXiv:2011.13456*, 2020.
- 716 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging
717 with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021.
- 718 Julián Tachella, Dongdong Chen, and Mike Davies. Unsupervised learning from incomplete mea-
719 surements for inverse problems. *arXiv preprint arXiv:2201.12151*, 2022a.
- 720 Julián Tachella, Dongdong Chen, and Mike Davies. Sensing theorems for unsupervised learning in
721 linear inverse problems, 2022b.
- 722 U Tariq, P Lai, M Lustig, M Alley, M Zhang, G Gold, and Vasanaawala Shreyas S. MRI Data:
723 Undersampled Knees. URL <http://old.mridata.org/undersampled/knees>.
- 724 Radhika Tibrewala, Tarun Dutt, Angela Tong, Luke Ginocchio, Mahesh B Keerthivasan, Steven H
725 Baete, Sumit Chopra, Yvonne W Lui, Daniel K Sodickson, Hersh Chandarana, and Patricia M
726 Johnson. Fastmri prostate: A publicly available, biparametric mri dataset to advance machine
727 learning for prostate cancer imaging, 2023.
- 728 Martin Uecker, Peng Lai, Mark J. Murphy, Patrick Virtue, Michael Elad, John M. Pauly, Shreyas S.
729 Vasanaawala, and Michael Lustig. Espirit—an eigenvalue approach to autocalibrating parallel
730 mri: Where sense meets grappa. *Magnetic Resonance in Medicine*, 71(3):990–1001, 2014. doi:
731 <https://doi.org/10.1002/mrm.24751>. URL [https://onlinelibrary.wiley.com/doi/](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24751)
732 [abs/10.1002/mrm.24751](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.24751).
- 733 Martin Uecker, Jonathan I Tamir, Frank Ong, and Michael Lustig. The bart toolbox for computational
734 magnetic resonance imaging. In *Proc Intl Soc Magn Reson Med*, volume 24, 2016.
- 735 Frederic Wang, Han Qi, Alfredo De Goyeneche, Reinhard Heckel, Michael Lustig, and Efrat Shimron.
736 K-band: Self-supervised MRI Reconstruction via Stochastic Gradient Descent over K-space
737 Subsets. *arXiv e-prints*, art. arXiv:2308.02958, August 2023. doi: 10.48550/arXiv.2308.02958.
- 738 Burhaneddin Yaman, Seyed Amir Hossein Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil,
739 and Mehmet Akçakaya. Self-supervised learning of physics-guided reconstruction neural networks
740 without fully sampled reference data. *Magnetic Resonance in Medicine*, 84(6):3172–3191, 2020.
741 doi: <https://doi.org/10.1002/mrm.28378>. URL [https://onlinelibrary.wiley.com/](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28378)
742 [doi/abs/10.1002/mrm.28378](https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28378).
- 743 Martin Zach, Florian Knoll, and Thomas Pock. Stable Deep MRI Reconstruction using Generative
744 Priors. *arXiv e-prints*, art. arXiv:2210.13834, October 2022. doi: 10.48550/arXiv.2210.13834.
- 745 Tao Zhang, John Pauly, Shreyas Vasanaawala, and Michael Lustig. MRI Data: Undersampled
746 Abdomens. URL <http://old.mridata.org/undersampled/abdomens>.
- 747 Tao Zhang, John M Pauly, Shreyas S Vasanaawala, and Michael Lustig. Coil compression for
748 accelerated imaging with cartesian sampling. *Magnetic resonance in medicine*, 69(2):571–582,
749 2013.

756 Tao Zhang, Joseph Y. Cheng, Aaron G. Potnick, Richard A. Barth, Marcus T. Alley, Martin Uecker,
757 Michael Lustig, John M. Pauly, and Shreyas S. Vasanawala. Fast pediatric 3d free-breathing
758 abdominal dynamic contrast enhanced mri with high spatiotemporal resolution. *Journal of Mag-*
759 *netic Resonance Imaging*, 41(2):460–473, 2015. doi: <https://doi.org/10.1002/jmri.24551>. URL
760 <https://onlinelibrary.wiley.com/doi/abs/10.1002/jmri.24551>.
761
762 Hongkai Zheng, Weili Nie, Arash Vahdat, and Anima Anandkumar. Fast training of diffusion models
763 with masked transformers. *arXiv preprint arXiv:2306.09305*, 2023.
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

810 A THEORETICAL RESULTS

811 A.1 THEORY PROOF

812 In this section, we provide the proof of Theorem 2.1 that was stated in the main paper. We begin by
813 giving the formal statement of the theorem.

814 **Theorem A.1** (Formal Statement of Theorem 2.1). *Let $\mathbf{x} \in \mathbb{C}^n$ be an unknown signal from a
815 distribution that admits density $p(\mathbf{x})$. Let $\{S_i\}$ be a random set of diagonal matrices $\in \mathbb{C}^n$ that
816 satisfies: $\sum_i S_i^H S_i = I_{n \times n}$. Assume sample access to linearly corrupted measurements of \mathbf{x} given
817 by:*

$$818 \mathbf{y}_{\text{train}} = \underbrace{\left(\sum_i S_i^H \mathcal{F}^{-1} P \mathcal{F} S_i \right)}_{A_{\text{train}}} \mathbf{x}, \quad (\text{A.1})$$

819 where \mathcal{F} is the discrete Fourier matrix and P is a random inpainting matrix, i.e. a diagonal matrix
820 with either zeros or ones in the diagonal, such that $\Pr[P_{ii} = 1] > 0$, for all entries i of the diagonal.
821 Define:

$$822 \tilde{\mathbf{y}}_{\text{train}} = \underbrace{\left(\sum_i S_i^H \mathcal{F}^{-1} \tilde{P} \mathcal{F} S_i \right)}_{\tilde{A}_{\text{train}}} \mathbf{x}, \quad (\text{A.2})$$

823 where \tilde{P} is a further corrupted version of P in the sense that with some non-zero probability p , a
824 diagonal element that was 1 in P becomes 0 in \tilde{P} . Define also $\mathbf{x}_t, \mathbf{y}_{t,\text{train}}, \tilde{\mathbf{y}}_{t,\text{train}}$ the noisy versions
825 of $\mathbf{x}, \mathbf{y}_{\text{train}}, \tilde{\mathbf{y}}_{\text{train}}$ respectively, as in:

$$826 \mathbf{x}_t = \mathbf{x} + \sigma_t \boldsymbol{\eta}, \quad \mathbf{y}_{t,\text{train}} = A_{\text{train}} (\mathbf{x} + \sigma_t \boldsymbol{\eta}), \quad \tilde{\mathbf{y}}_{t,\text{train}} = \tilde{A}_{\text{train}} (\mathbf{x} + \sigma_t \boldsymbol{\eta}). \quad (\text{A.3})$$

827 Then, the minimizer of the objective:

$$828 J(\theta) = \mathbb{E}_{\mathbf{y}_{0,\text{train}}, \tilde{\mathbf{y}}_{t,\text{train}}, A_{\text{train}}, \tilde{P}} \left[\left\| A_{\text{train}} \mathbf{h}_\theta(\tilde{\mathbf{y}}_{t,\text{train}}, \tilde{P}) - \mathbf{y}_{0,\text{train}} \right\|^2 \right], \quad (\text{A.4})$$

829 is: $\mathbf{h}_\theta(\tilde{\mathbf{y}}_{t,\text{train}}, \tilde{P}) = \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_{t,\text{train}}, \tilde{P}]$.

830 *Proof.* We adapt the proof Theorem 4.1 in Ambient Diffusion (Daras et al., 2023b). To avoid notation
831 clutter, we will be omitting the subscript “train”, when necessary.

832 Let $\mathbf{h}_{\theta^*}(\tilde{\mathbf{y}}_t, \tilde{P}) = \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] + \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P})$ be the optimal solution. The value of the objective for the
833 optimal solution becomes:

$$834 J(\theta^*) = \mathbb{E}_{\mathbf{y}_0, \tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\left\| A \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] + A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) - \mathbf{y}_0 \right\|^2 \right] \quad (\text{A.5})$$

$$835 = \mathbb{E}_{\mathbf{y}_0, \tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\underbrace{\left\| A \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] - \mathbf{y}_0 \right\|^2}_{\text{irreducible error}} + \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P})^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) - 2 \left(\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] - \mathbf{x}_0 \right)^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right]. \quad (\text{A.6})$$

836 We will now work with the last term.

$$837 \mathbb{E}_{\mathbf{x}_0, \tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\left(\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] - \mathbf{x}_0 \right)^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right] \quad (\text{A.7})$$

$$838 = \mathbb{E}_{\tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\mathbb{E}_{\mathbf{x}_0 | \tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\left(\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] - \mathbf{x}_0 \right)^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right] \right] \quad (\text{A.8})$$

$$839 = \mathbb{E}_{\tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\left(\mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, \tilde{P}] - \mathbb{E}[\mathbf{x}_0 | \tilde{\mathbf{y}}_t, A, \tilde{P}] \right)^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right] \quad (\text{A.9})$$

$$840 = 0, \quad (\text{A.10})$$

864 since $\mathbb{E}[\mathbf{x}_0|\tilde{\mathbf{y}}_t, \tilde{P}] = \mathbb{E}[\mathbf{x}_0|\tilde{\mathbf{y}}_t, A, \tilde{P}]$, i.e. the value of \mathbf{x}_0 does not depend on \tilde{A} given \tilde{P} and $\tilde{\mathbf{y}}_t$.

865 We will now work with the middle term. We have that:

$$866 \mathbb{E}_{\tilde{\mathbf{y}}_t, A, \tilde{P}} \left[\mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P})^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right] \quad (\text{A.11})$$

$$867 = \mathbb{E}_{\tilde{\mathbf{y}}_t, \tilde{P}} \left[\mathbb{E}_{A|\tilde{\mathbf{y}}_t, \tilde{P}} \left[\mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P})^T A^T A \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right] \right] \quad (\text{A.12})$$

$$868 = \mathbb{E}_{\tilde{\mathbf{y}}_t, \tilde{P}} \left[\mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P})^T \mathbb{E}[A^T A|\tilde{P}] \mathbf{f}(\tilde{\mathbf{y}}_t, \tilde{P}) \right]. \quad (\text{A.13})$$

871 From the last equation, it is evident that if $\mathbb{E}[A^T A|\tilde{P}]$ is full-rank, and hence the minimizer is
 872 $\mathbb{E}[\mathbf{x}_0|\tilde{\mathbf{y}}_t, \tilde{P}]$ as needed. Since our A matrix is square, it suffices to show that $\mathbb{E}[A|\tilde{P}]$ is full-rank.
 873 By Corollary A.1. in Ambient Diffusion, we have that: $\mathbb{E}[P|\tilde{P}]$ is full-rank. We will now show
 874 that $\mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F}$ is also full-rank. This can be easily proved with contradiction. Assume that
 875 $\mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F}$ is not full-rank. Then, there exists a vector $\mathbf{w} \neq 0$ such that:

$$876 \mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F}\mathbf{w} = 0 \iff \quad (\text{A.14})$$

$$877 \mathbb{E}[P|\tilde{P}] \underbrace{\mathcal{F}\mathbf{w}}_{z \neq 0} = 0 \iff \quad (\text{A.15})$$

$$878 \mathbb{E}[P|\tilde{P}]z = 0, \quad (\text{A.16})$$

879 which is a contradiction, since $z \neq \mathbf{0}$ and $\mathbb{E}[P|\tilde{P}]$ is full-rank.

880 So far, we have established that $\mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F}$ is full-rank. By linearity of expectation, we further
 881 have that:

$$882 \mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F} = \mathbb{E}[\mathcal{F}^{-1}P\mathcal{F}|\tilde{P}], \quad (\text{A.17})$$

883 and hence, the latter is also full-rank. Finally, since $\sum_i S_i^H S_i = I$, for any full-rank matrix C , we
 884 have that $\sum_i S_i^H C S_i$ is also full-rank. This is true for any given set S_i and hence it is also true for
 885 the expectation over the sets $\{S_i\}$. Putting everything together, we have that:

$$886 \mathbb{E}_{\{S_i\}, P} \left[\sum_i S_i^H \mathcal{F}^{-1}\mathbb{E}[P|\tilde{P}]\mathcal{F} S_i \middle| \tilde{P} \right] = \mathbb{E}[A|\tilde{P}], \quad (\text{A.18})$$

887 is full-rank and the proof is complete.

888 \square

890 **Note on alternative Ambient Diffusion designs.** Our Ambient Diffusion MRI training approach
 891 can be extended to use other linear operators to aggregate the measurements from different coils (e.g.
 892 we could have used the pseudoinverse). Alternatively, we can also further condition on the sensitivity
 893 maps, or directly feed all the coil measurements as input to the network without aggregation. We
 894 opted for a simple and fast method (aggregation through the adjoint). However, these other approaches
 895 can lead to improved performance because they fully leverage coil information.

896 B DATASET PREPARATION

897 We follow standard practices in the MRI literature. We download the FastMRI dataset, randomly
 898 select 2,000 T2-Weighted scans, and pick 5 center slices from each scan, giving us us a dataset of
 899 10,000 T2-weighted brain scans (k -space measurements). We pass each multi-channel k -space sample
 900 through a noise pre-whitening filter to transform the noise into standard white Gaussian. Then, we
 901 normalize each k -space sample by the absolute value of 99-th percentile of the root sum-of-squares
 902 reconstruction of the autocalibration region ($24x \times 24$ pixels from the center). Given the fully sampled
 903 k -space samples, we estimate sensitivity maps using the ESPiRiT calibration method (Uecker et al.,
 904 2014). These measurements serve as a fully sampled reference dataset (i.e., $R = 1$).

C DETAILS FOR BASELINES.

L1 Methods. We used the BART implementation (Uecker et al., 2016; Blumenthal et al., 2022) and searched for the best regularization parameter over the training set. Specifically, we used a regularization weighting of 0.001 with 100 iterations.

Self-Supervised learning via Data Under-sampling (SSDU) (Yaman et al., 2020) In a supervised setting, end-to-end image restoration networks can be trained by passing measurements through a network like MoDL (Aggarwal et al., 2019) or VarNet (Hammernik et al., 2018) and taking a loss directly on the output of the network with the ground truth image. In the self-supervised setting, however, this is not possible. SSDU, trains the end-to-end restoration network by splitting the available measurement set (Ω) of each sample y_i^Ω into training sets (Θ) and loss sets (Λ) or measurements y_i^Θ and y_i^Λ respectively where $\Omega = \Theta \cup \Lambda$. Where the reconstruction network is given access to measurements in the training set to obtain an estimated image $\hat{x}_i^\Theta = \mathbf{h}_\theta(y_i^\Theta)$ and the loss is then defined over the loss measurements as

$$L(\mathbf{y}_i^\Lambda, \hat{\mathbf{x}}_i^\Theta) = \frac{\|\mathbf{y}_i^\Lambda - A_\Lambda \hat{\mathbf{x}}_i^\Theta\|_1}{\|\mathbf{y}_i^\Lambda\|_1} + \frac{\|\mathbf{y}_i^\Lambda - A_\Lambda \hat{\mathbf{x}}_i^\Theta\|_2}{\|\mathbf{y}_i^\Lambda\|_2} \quad (\text{C.1})$$

where A_Λ is the forward operator with ones in the selection mask P at only the measurement locations in the loss set and zeros elsewhere. In this work, we selected $\mathbf{h}_\theta(\cdot)$ to have the MoDL architecture (Aggarwal et al., 2019). We trained individual models at four different acceleration levels ($R = 2, 4, 6, 8$). Each model was trained for ten epochs using the same training set (10,000 slices) as previously described. We used uniform random sampling to separate the training and loss measurement groups for each sample and found that a split of $\rho = \frac{|\Lambda|}{|\Omega|} = 0.2$ provided the best performance and thus was used for reporting all subsequent metrics.

MoDL Supervised (MoDL-Sup) To provide an upper bound on performance for supervised end-to-end methods we also trained the MoDL architecture in a supervised fashion (Aggarwal et al., 2019). Again, we trained individual models at four different acceleration levels ($R = 2, 4, 6, 8$). Each model was trained using the same 10,000 slices as above. However, for these models, we used the normalized root mean squared error (NRMSE) as the loss function:

$$L(\mathbf{x}_i, \hat{\mathbf{x}}_i) = \frac{\|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2}{\|\mathbf{x}_i\|_2} \quad (\text{C.2})$$

where \mathbf{x}_i are the ground truth images from our dataset and $\hat{\mathbf{x}}_i = \mathbf{h}_\theta(\mathbf{y}_i, A_i)$ are the reconstructions provided by our network based on under-sampled measurements \mathbf{y}_i .

D HYPERPARAMETERS FOR MRI MODELS

Training Hyperparameters. The EDM model trained on the fully-sampled ($R = 1$) dataset, as well as the four EDM models trained on L1-Wavelet reconstructions at each acceleration factor $R \in \{2, 4, 6, 8\}$ were all trained with 65 million parameters. The Ambient Diffusion models on the other hand were trained with 36 million parameters, for faster training. All the Ambient Diffusion models were trained for 250 epochs. For the Ambient Diffusion training experiments, the further corrupted sampling masks are given as an input to the model by concatenating with the measurements along the channel dimension, as in the original Ambient Diffusion (Daras et al., 2023b) paper.

Sampling Hyperparameters. The only tunable parameters for DPS (Eq. 2.3) and Ambient DPS (Eq. 2.11) are in the scheduling of the magnitude of the measurements likelihood term. In all the experiments in the DPS paper, this term is kept constant throughout the diffusion sampling trajectory and the authors recommend selecting a value in the range between $[0.1, 10]$. We follow this recommendation and we keep this term constant. The value of the step size for each model is selected with a hyperparameter search in the recommended range. For all our experiments, we follow exactly the DPS implementation provided in the official code repository of the paper. The other parameter

that impacts performance is the number of steps we are going to run each algorithm for, i.e. the discretization level of the SDEs of Equations 2.3, 2.11. Typically, the higher the number of steps the better the performance since the discretization error decreases (Chen et al., 2022b; 2023). For the performance results, we run each method for several steps $\in \{50, 100, 150, 200, 250, 300\}$.

E ADDITIONAL MRI RESULTS

Table 3: Accelerated MRI reconstruction performance evaluation using LPIPS \downarrow and DISTs \downarrow , averaged across 100 test samples.

Training Method	Training Data	Reconstruction Method	$R = 2$		$R = 4$		$R = 6$		$R = 8$	
			LPIPS	DISTS	LPIPS	DISTS	LPIPS	DISTS	LPIPS	DISTS
Supervised	$R = 1$	FS-DPS	0.061	0.028	0.106	0.059	0.166	0.094	0.231	0.126
Self-Supervised	$R = 2$	A-DPS	0.118	0.062	0.127	0.068	0.138	0.076	0.150	0.082
		A-ALD	0.119	0.063	0.127	0.068	0.139	0.077	0.151	0.083
		L1-DPS	0.067	0.032	0.131	0.075	0.201	0.115	0.262	0.143
		1D-Partitioned SSDU	0.141	0.084	0.144	0.081	0.187	0.108	0.223	0.129
	$R = 4$	A-DPS	0.155	0.090	0.155	0.092	0.161	0.095	0.168	0.099
		A-ALD	0.156	0.090	0.155	0.092	0.161	0.095	0.168	0.099
		L1-DPS	0.096	0.051	0.193	0.112	0.258	0.147	0.301	0.167
		1D-Partitioned SSDU	0.169	0.105	0.169	0.092	0.187	0.104	0.214	0.123
	$R = 6$	A-DPS	0.180	0.110	0.178	0.110	0.181	0.111	0.186	0.114
		A-ALD	0.179	0.109	0.179	0.111	0.181	0.111	0.186	0.114
		L1-DPS	0.112	0.061	0.221	0.128	0.282	0.161	0.318	0.179
		1D-Partitioned SSDU	0.143	0.092	0.169	0.102	0.196	0.117	0.222	0.133
	$R = 8$	A-DPS	0.212	0.129	0.204	0.127	0.204	0.127	0.204	0.127
		A-ALD	0.213	0.130	0.204	0.127	0.204	0.127	0.204	0.127
		L1-DPS	0.114	0.064	0.227	0.134	0.291	0.168	0.327	0.187
		1D-Partitioned SSDU	0.142	0.094	0.176	0.109	0.206	0.127	0.231	0.140




Figure 7: Ambient diffusion posterior sampling (A-DPS) multi-step performance metrics at $R = 2, 4, 6, 8$ for models trained at $R = 2, 4, 6, 8$.

1026
 1027
 1028
 1029
 1030
 1031
 1032
 1033
 1034
 1035
 1036
 1037
 1038
 1039
 1040
 1041
 1042
 1043
 1044
 1045
 1046
 1047
 1048
 1049
 1050
 1051
 1052
 1053
 1054
 1055
 1056
 1057
 1058
 1059
 1060
 1061
 1062
 1063
 1064
 1065
 1066
 1067
 1068
 1069
 1070
 1071
 1072
 1073
 1074
 1075
 1076
 1077
 1078
 1079




Figure 8: Ambient diffusion one step (A-OS) performance metrics at $R = 2, 4, 6, 8$ for models trained at $R = 2, 4, 6, 8$.




Figure 9: Performance metrics for FS-DPS, L1-DPS, FS-ALD, and L1-CS at $R = 2, 4, 6, 8$




Figure 10: Fully supervised MoDL performance metrics at $R = 2, 4, 6, 8$ for models trained at $R = 2, 4, 6, 8$.

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133


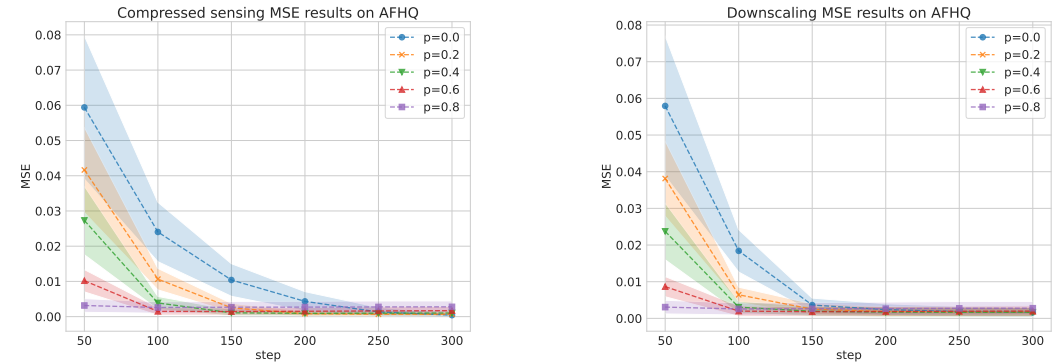


Figure 11: SSDU performance metrics at $R = 2, 4, 6, 8$ for models trained at $R = 2, 4, 6, 8$.

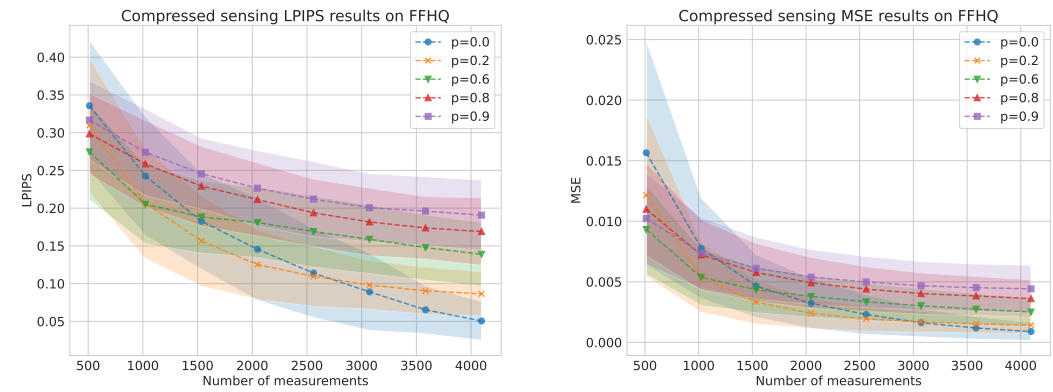
F ADDITIONAL NATURAL IMAGE PERFORMANCE RESULTS



(a) Compressed Sensing with 4000 measurements per Number of Function Evaluations (NFEs).

(b) $2\times$ Super-Resolution per Number of Function Evaluations (NFEs).

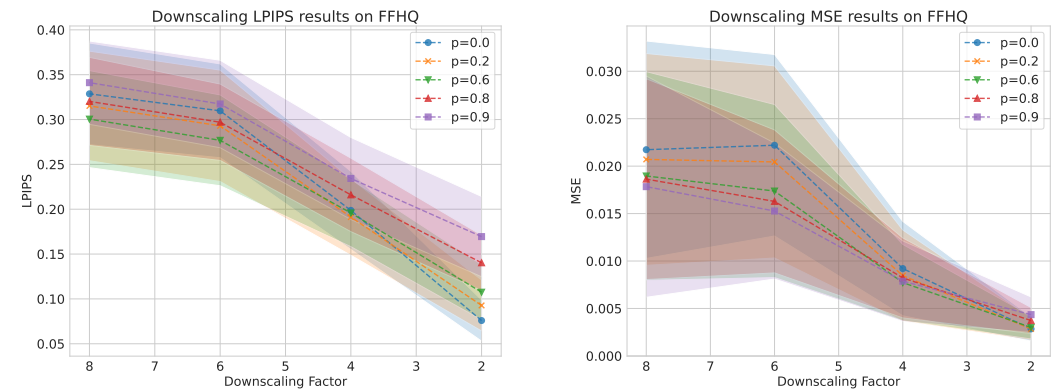
Figure 12: Speed performance plots for AFHQ.



(a) LPIPS per Number of Measurements.

(b) MSE per Number of Measurements.

Figure 13: Compressed Sensing Results for FFHQ.




(a) LPIPS per Downscaling Factor.

(b) MSE per Downscaling Factor.

Figure 14: Downscaling Results for FFHQ.

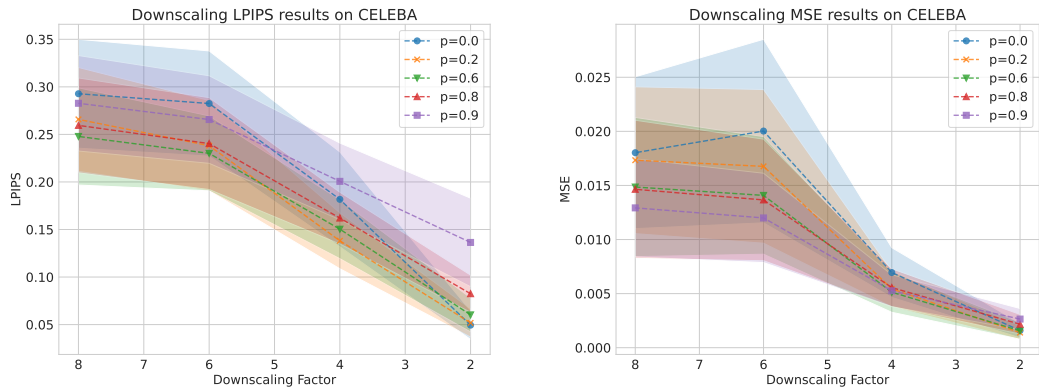
1188
 1189
 1190
 1191
 1192
 1193
 1194
 1195
 1196
 1197
 1198
 1199
 1200
 1201
 1202
 1203
 1204
 1205
 1206
 1207
 1208
 1209
 1210
 1211
 1212
 1213
 1214
 1215
 1216
 1217
 1218
 1219
 1220
 1221
 1222
 1223
 1224
 1225
 1226
 1227
 1228
 1229
 1230
 1231
 1232
 1233
 1234
 1235
 1236
 1237
 1238
 1239
 1240
 1241



(a) LPIPS per Number of Measurements.

(b) MSE per Number of Measurements.

Figure 15: Compressed Sensing Results for Celeb-A.



(a) LPIPS per Downscaling Factor.

(b) MSE per Downscaling Factor.

Figure 16: Downscaling Results for Celeb-A.

For the following experiments, we compare a model trained without corruption to a model trained with randomly inpainted data at $p = 0.6$. The evaluation dataset is FFHQ. The training size and hyperparameters for both models are the same. For both models, we extensively tune the sampling hyperparameters to maximize their performance (starting from the recommended hyperparameters from the DPS work). The finding we get in these additional experiments is that the model trained on corrupted data (inpainting) outperforms the model trained on clean data in the high corruption regime across all evaluation tasks.




Figure 17: Box Inpainting results for FFHQ





Figure 18: Additive Gaussian Noise results for FFHQ

Figure 19: Super Resolution + Additive Gaussian Noise ($\sigma = 0.05$) results for FFHQ

G ADDITIONAL NATURAL IMAGE SPEED RESULTS

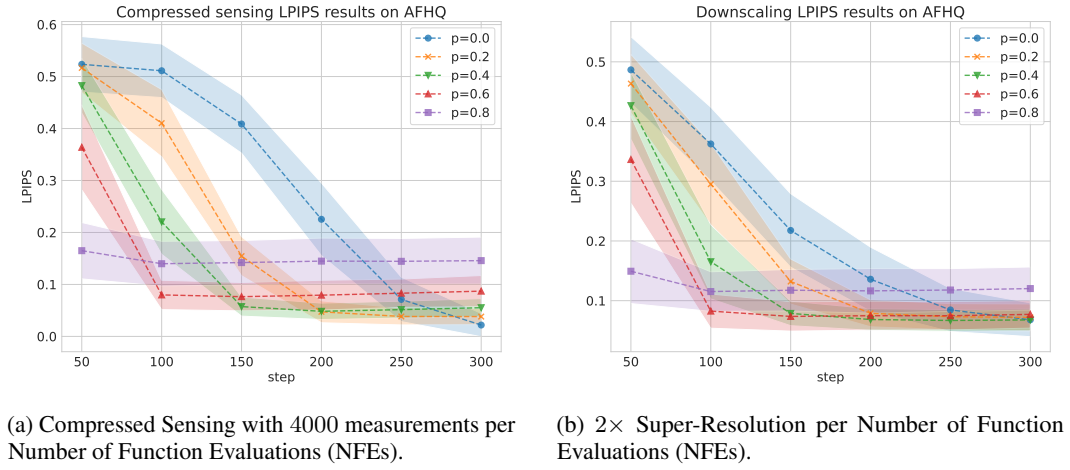


Figure 20: Speed LPIPS performance plots for AFHQ.

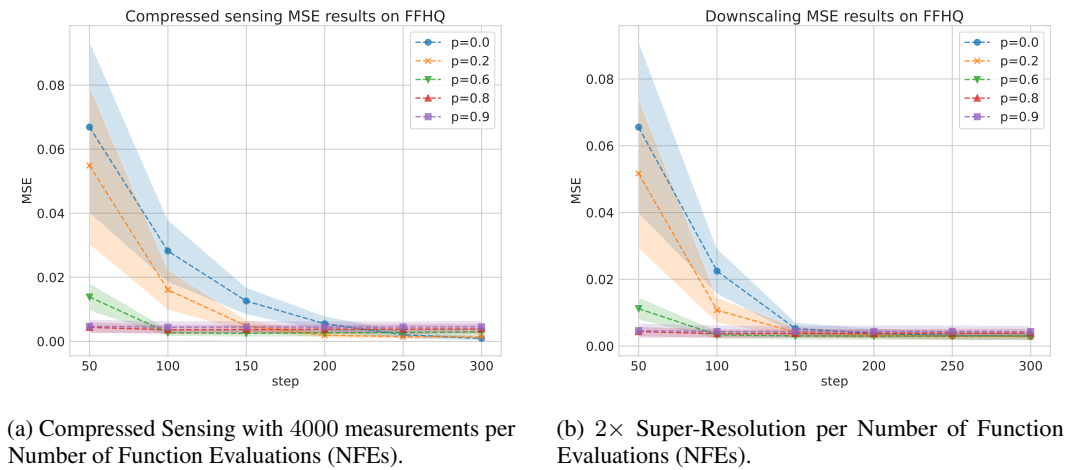



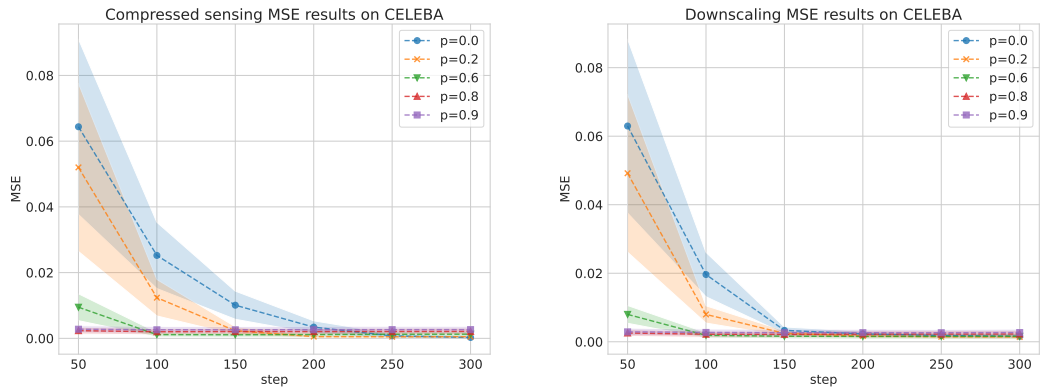
Figure 21: Speed MSE performance plots for FFHQ.

1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380
 1381
 1382
 1383
 1384
 1385
 1386
 1387
 1388
 1389
 1390
 1391
 1392
 1393
 1394
 1395
 1396
 1397
 1398
 1399
 1400
 1401
 1402
 1403



(a) Compressed Sensing with 4000 measurements per Number of Function Evaluations (NFEs). (b) $2\times$ Super-Resolution per Number of Function Evaluations (NFEs).

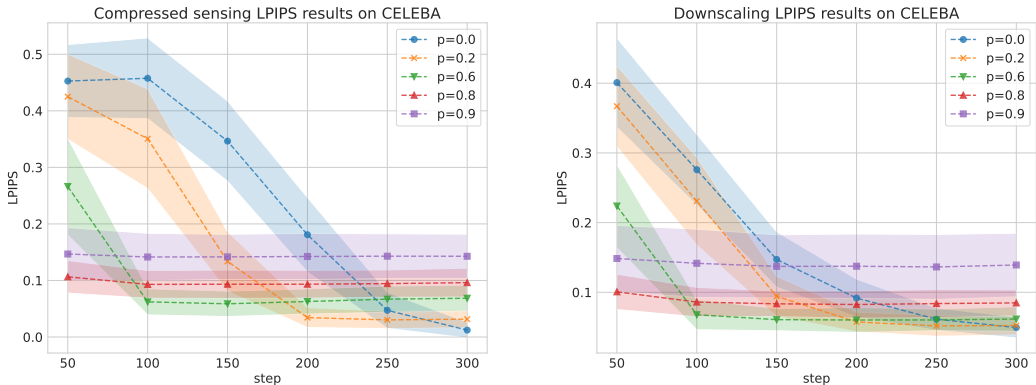
Figure 22: Speed LPIPS performance plots for FFHQ.



(a) Compressed Sensing with 4000 measurements per Number of Function Evaluations (NFEs). (b) $2\times$ Super-Resolution per Number of Function Evaluations (NFEs).

Figure 23: Speed MSE performance plots for Celeb-A.

1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457



(a) Compressed Sensing with 4000 measurements per Number of Function Evaluations (NFEs). (b) $2\times$ Super-Resolution per Number of Function Evaluations (NFEs).

Figure 24: Speed LPIPS performance plots for Celeb-A.

H LIMITATIONS AND FUTURE DIRECTIONS.

Our work has several limitations. First, there is a lack of theoretical understanding of the observed experimental phenomenon: we do not have a good grasp of why models trained on highly corrupted data lead to better priors for inverse problems of another corruption type. Second, we only tested the reconstruction performance of our models using the DPS algorithm. Several other recent algorithms have been developed for solving inverse problems with diffusion models and it is unknown whether our findings generalize for these reconstruction algorithms. Finally, our method relies on the existence of Ambient Diffusion models. There are many more models available trained on clean data than models trained on highly corrupted data.

The auto-calibration signal (ACS) region is shared across training samples, which could lead to an over-representation of certain k-space regions. While our current approach does not explicitly address this, introducing a diagonal weighting matrix in the loss function, as suggested by Millard & Chiew (2023), is a promising direction to mitigate this issue.

Currently, we define A_{train} according to Eq 2.7. This is convenient because it lets us use an image-to-image network architecture such as that used by EDM. An alternative could be to define A_{train} directly as in Eq 2.5, in which case the network would go from multi-coil k-space to an image. This network could be preceded by an IFFT, which would effectively be a multi-coil-to-image network, and could be standardized using coil compression (Zhang et al., 2013). Such a network architecture could be interesting and potentially more expressive, as we currently collapse the multi-coil information through zero-filling adjoint. Another possibility is to define A_{train} as the pseudo-inverse applied to A , which could be solved efficiently with the conjugate gradient algorithm.