

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NOISE-ADAPTIVE DIFFUSION SAMPLING FOR IN- VERSE PROBLEMS WITHOUT TASK-SPECIFIC TUNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion models (DMs) have recently shown remarkable performance on inverse problems (IPs). Optimization-based methods can fast solve IPs using DMs as powerful regularizers, but it is susceptible to local minima and noise overfitting. Although DMs can provide strong priors for Bayesian approaches, enforcing measurement consistency during the denoising process leads to manifold infeasibility issues. We propose Noise-space Hamiltonian Monte Carlo (N-HMC), a posterior sampling method that treats reverse diffusion as a deterministic mapping from initial noise to clean images. N-HMC enables comprehensive exploration of the solution space, avoiding local optima. By moving inference entirely into the initial-noise space, N-HMC keeps proposals on the learned data manifold. We provide a comprehensive theoretical analysis of our approach and extend the framework to a noise-adaptive variant (NA-NHMC) that effectively handles IPs with unknown noise type and level. Extensive experiments across four linear and three nonlinear inverse problems demonstrate that NA-NHMC achieves superior reconstruction quality with robust performance across different hyperparameters and initializations, significantly outperforming recent state-of-the-art methods. Code will be made available on GitHub upon publication.

1 INTRODUCTION

Inverse problems (IPs) have wide applications in many domains, including computer vision (Janai et al., 2021; Quan et al., 2024), protein science (Yi et al., 2023; Ouyang-Zhang et al., 2023; Yang et al., 2019), medical imaging (Song et al., 2022b; Chu et al., 2025; Dao et al., 2024), scientific computing (Zheng et al., 2025; Xia & Zabaras, 2022; Xu et al., 2024). The goal is to reconstruct an unknown $\mathbf{x} \in \mathbb{R}^n$ from noisy measurements $\mathbf{y} \in \mathbb{R}^m$:

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \eta, \quad (1)$$

where \mathcal{A} is a known forward operator, and $\eta \in \mathbb{R}^m$ is additive noise. Diffusion models (DMs) have recently shown powerful capabilities in modeling complex data distributions, which can provide a powerful class of priors for high-dimensional data x in solving IPs. Existing diffusion-based methods have demonstrated remarkable success across diverse inverse problems (Chung et al., 2023; Daras et al., 2024; Zheng et al., 2025; Song et al., 2022b).

Although remarkable progress has been made, as illustrated in Figure 1, current diffusion-based methods suffer from three complementary limitations and issues: (1) Iterative guidance methods such as DPS (Chung et al., 2023), DDRM (Kawar et al., 2022), DDNM (Wang et al., 2023), PIIGDM (Song et al., 2023), and TMPD (Boys et al., 2024) use the likelihood term to shift intermediate images directly, systematically pushing intermediate states off the learned data manifold and violating the training-time noise-conditioning of the denoiser, resulting in various failure reconstructions like like accumulated artifacts as shown in Figure 1 (a). (2) Stochastic MAP methods that optimize in image space, including ReSample (Song et al., 2024), DiffPIR (Zhu et al., 2023), DAPS (Zhang et al., 2024), SITCOM (Alkhouri et al., 2025a), and DIP (Chihaoui & Favaro, 2025a) can match y well with very sharp details but require carefully tuned hyperparameters to not overfit to noise. This limits their effectiveness in high or unknown noise settings (3) Deterministic MAP methods that optimize in the DM noise space (DMPlug, (Wang et al., 2024)) remove randomness but often get stuck in a single mode, especially in severely ill-posed problems like phase retrieval, due to a lack

of exploration. In short, enforcing data consistency mid-diffusion can break prior adherence, while optimizing only for fidelity leads to overfitting or mode collapse. Building a solution that maintains both measurement fidelity and manifold adherence remains open.

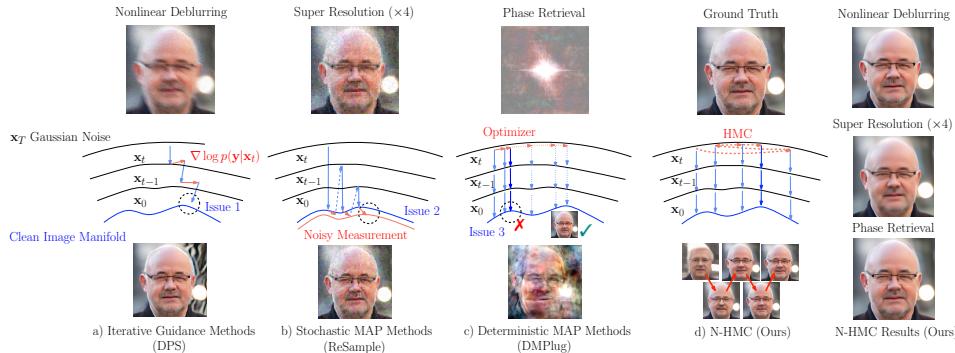


Figure 1: Comparison of existing methods and their limitations with the N-HMC method. (a) Iterative Guidance Methods (*DPS*) lead to *manifold infeasibility*. (b) Stochastic MAP methods (*Resample*) (Song et al., 2024) are susceptible to *overfitting to noise*. (c) Deterministic MAP methods (*DMPlug*) (Wang et al., 2024) become *trapped in a local mode*. (d) Our method performs sampling in the noise space \mathbf{x}_T and maps samples to images via a deterministic mapping $\mathbf{x}_0 = \mathcal{D}(\mathbf{x}_T)$.

Sampling from the full posterior ensures that the learned prior acts automatically as a regularizer, while an annealing schedule for **noise standard deviation** σ_y promotes efficient exploration and prevents the sampler from being trapped in early local modes. Importantly, the method relies only on a fixed set of hyperparameters that remain constant across tasks, datasets, and levels of measurement noise, avoiding the repeated tuning required by many existing approaches.

To address the practical challenges that the measurement noise level is often unknown, we further introduce a Noise-Adaptive N-HMC (NA-NHMC). Instead of requiring a fixed noise level, we take a principled Bayesian approach, placing a non-informative prior on the noise variance and marginalizing it out. This yields a parameter-free likelihood term that automatically adapts to the true underlying noise in the measurements. As shown in experiments, this allows NA-NHMC to achieve robust, high-quality reconstructions across varying and even unknown noise types and levels without any task-specific hyperparameter tuning. In contrast, the performance of other methods depends on the hyperparameters listed in Section A.5, which were specifically tuned for Gaussian noise. Our key *contributions* include: (1) In Section 3.1, we propose N-HMC, a posterior sampling method that addresses the three key limitations of existing state-of-the-art (SOTA) approaches. We further analyze its sampling behavior and provide a theoretical guarantee of its robustness to measurement noise in Section 3.2 and Appendix A.2. (2) In Section 3.3, we extend our method to settings with unknown noise types and levels. We show that it outperforms SOTA methods on most metrics (Section 4.3), especially for non-linear and high noise problems. (3) In extensive experiments, NA-NHMC method solves diverse inverse problem tasks under unknown noise types and levels without any task- or noise-specific hyperparameter tuning, in contrast to many existing methods. (4) We demonstrate in Section 4.1 that the annealing schedule for σ_y helps promote early exploration and prevent local-mode collapse, especially in severely ill-posed tasks like phase retrieval.

2 PRELIMINARIES

2.1 DIFFUSION MODELS FOR INVERSE PROBLEMS

Daras et al. (2024) broadly classifies methods for solving inverse problems (IPs) into two categories. The first is maximum a posteriori (MAP) inference, which aims to find the single most probable \mathbf{x} . An alternative is the Bayesian framework, where the goal becomes generating plausible reconstructions by sampling from the posterior distribution $p(\mathbf{x}|\mathbf{y})$, where $p(\mathbf{x}|\mathbf{y})$ can be decomposed into the prior $p(\mathbf{x})$ and the likelihood $p(\mathbf{y}|\mathbf{x})$. MAP delivers fast optimization, but struggles with high noise and multimodal posteriors, easily converging to local minima. In contrast, the Bayesian approach samples from $p(\mathbf{x}|\mathbf{y})$ to generate plausible reconstructions, quantify uncertainty, and handle mul-

108 timodality. Both approaches critically depend on powerful prior models like DMs that encode the
 109 complex statistical structure of complex data and prior knowledge.
 110

111 Most diffusion-based approaches to IPs are based on the denoising diffusion probabilistic models
 112 (DDPM) framework (Ho et al., 2020; Song & Ermon, 2020). The framework consists of forward and
 113 reverse diffusion processes. The forward process gradually corrupts the clean images \mathbf{x}_0 towards
 114 standard Gaussian noise \mathbf{x}_T . This process can be described by a stochastic differential equation
 115 (SDE), $d\mathbf{x} = -\frac{\beta_t}{2}\mathbf{x}dt + \sqrt{\beta_t}d\mathbf{w}$, where \mathbf{w} is the standard Wiener process. In practice, the process
 116 is discretized via a variance schedule $\{\beta_t\}_{t=1}^T$, forming a Markov chain:
 117

$$118 \quad q(\mathbf{x}_{1:T}|\mathbf{x}_0) := \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t}\mathbf{x}_{t-1}, \beta_t \mathbf{I}) \quad (2)$$

120 In order to generate clean images, the reverse process begins with a noisy sample $\mathbf{x}_T \sim \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$,
 121 and recursively refines it according to the reverse SDE, $d\mathbf{x} = -\frac{\beta_t}{2}\mathbf{x}dt - \beta_t \nabla_{\mathbf{x}} \log p_t(\mathbf{x})dt + \sqrt{\beta_t}d\mathbf{w}$,
 122 where \mathbf{w} is the time-reversed standard wiener process, and $p_t(\mathbf{x})$ is the marginal probability of the
 123 noisy manifold at time t . $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ is called the score function and is usually approximated by a
 124 neural network θ trained through score-matching methods.
 125

126 Using the same discretization, clean images can be generated from the prior using an iterative de-
 127 noising process.
 128

$$129 \quad p_{\theta}(\mathbf{x}_{0:T}) := p(\mathbf{x}_T) \prod_{t=1}^T p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t), \quad p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_{\theta}(\mathbf{x}_t, t), \boldsymbol{\Sigma}_{\theta}(\mathbf{x}_t, t)) \quad (3)$$

131 Building on the DDPM framework, to accelerate the denoising process, Song et al. (2022a) proposes
 132 Denoising Diffusion Implicit Models (DDIM), which define a non-Markovian and fully determin-
 133 istic forward/reverse process ($\beta_t = 0$). Unlike DDPM, which injects stochasticity at each step to
 134 improve robustness, DDIM iteratively maps the initial noise \mathbf{x}_T to a clean sample \mathbf{x}_0 via a deter-
 135 ministic trajectory. For our method, this property is particularly beneficial, as it allows us to consider
 136 the entire reverse process as a deterministic mapping from \mathbf{x}_T to \mathbf{x}_0 .
 137

138 Among successful DM-based methods for inverse problems, DPS and its variants (Chung et al.,
 139 2023; Kawar et al., 2022; Wang et al., 2023; Song et al., 2023; Chung et al., 2022) are best-known
 140 reconstruction algorithms. **But they suffer from approximation errors from Tweedie’s formula cor-**
 141 **rections. To mitigate noise sensitivity, TMPD incorporates second-order information to correct the**
 142 **guidance trajectory; however, like other iterative methods, it relies on modifying intermediate states,**
 143 **which risks drifting off the learned manifold. SITCOM (Alkhouri et al., 2025b) operates on the**
 144 **noisy image at each diffusion step and enforces a triple-consistency constraint: data fidelity, back-**
 145 **ward consistency with the diffusion posterior mean, and forward consistency along the diffusion**
 146 **trajectory. DIP (Chihaoui & Favaro, 2025b) updates the initial noise \mathbf{x}_T with data fidelity gradients**
 147 **after the standard diffusion sampling process. DMPlug (Wang et al., 2024) proposes a noise-space**
 148 **formulation but treats inverse problems as optimization tasks, making it *sensitive to noise*. While**
 149 **early stopping can mitigate this, its criterion is task- and noise-dependent. At the high noise levels**
 150 **considered here, the optimizer often becomes *trapped in a local mode*, rendering early stopping**
 151 **ineffective. Similar behavior is observed in other Maximum a Posteriori (MAP) methods such as**
 152 **ReSample (Song et al., 2024), which optimizes directly in clean-image space (leading to noisy or**
 153 **blurry images under early stopping). DAPS (Zhang et al., 2024), despite being formulated as a pos-**
 154 **terior sampling method, uses a heuristic $\hat{\sigma}_y$ that is much smaller than its true value to strengthen the**
 155 **consistency of the measurement. This deviation from true posterior sampling makes DAPS effec-**
 156 **tively MAP-like, inheriting the same sensitivity to noise.**

157 2.2 HAMILTONIAN MONTE CARLO (HMC)

158 Hamiltonian Monte Carlo (HMC) (Duane et al., 1987) is an MCMC (Metropolis et al., 1953) sam-
 159 pling method that utilizes a fictitious momentum variable and simulates Hamiltonian dynamics to
 160 efficiently explore distant regions. Due to its superior scaling properties in high dimensions com-
 161 pared to other simpler Metropolis methods Brooks et al. (2011), HMC is particularly well suited for
 162 sampling in high-dimensional space, such as the $3 \times 256 \times 256$ pixel space of images.

162 The Hamiltonian is defined as $H = U + V$, where $U = -\log p(\mathbf{x})$ and $V = \frac{1}{2}\mathbf{v}^\top \mathbf{M}^{-1}\mathbf{v}$. Then,
 163 we discretize the trajectory using the leapfrog integrator. For a single leapfrog step with step size δ ,
 164 we have

$$166 \quad \mathbf{v}(t + \delta/2) = \mathbf{v}(t) - \frac{\delta}{2} \frac{\partial U}{\partial \mathbf{x}} \Big|_{\mathbf{x}(t)}, \quad (4)$$

$$168 \quad \mathbf{x}(t + \delta) = \mathbf{x}(t) + \delta \mathbf{M}^{-1} \mathbf{v}(t + \delta/2), \quad (5)$$

$$170 \quad \mathbf{v}(t + \delta) = \mathbf{v}(t + \delta/2) - \frac{\delta}{2} \frac{\partial U}{\partial \mathbf{x}} \Big|_{\mathbf{x}(t+\delta)}, \quad (6)$$

172 where $\mathbf{v}(0) \sim \mathcal{N}(\mathbf{v}; \mathbf{0}, \mathbf{M})$. This process is repeated L times to form a full trajectory. Due to a
 173 discretization error, the Hamiltonian is no longer preserved, which introduces bias and violates the
 174 detailed balance. To correct for this, a Metropolis-Hastings (MH) correction step is applied at the
 175 end of each trajectory with acceptance probability of $\alpha = \min(1, \exp(-H_1 + H_0))$, where H_0, H_1
 176 denotes the initial and proposed Hamiltonian, respectively.

3 METHODOLOGY

180 In this section, we propose a posterior sampling method, Noise-space Hamiltonian Monte Carlo
 181 (N-HMC), to solve IPs with pretrained DMs. We show its derivation in Section 3.1 and discuss its
 182 robustness to measurement noise in Section 3.2. In Section 3.3, our method is modified to allow for
 183 unknown types and levels of measurement noise.

3.1 NOISE-SPACE HAMILTONIAN MONTE CARLO (N-HMC)

187 The goal in solving inverse problems is to sample from the posterior distribution $p(\mathbf{x}_0|\mathbf{y}) \propto$
 188 $p(\mathbf{x}_0)p(\mathbf{y}|\mathbf{x}_0)$. Since direct sampling from $p(\mathbf{x}_0)$ is intractable, pretrained diffusion models are
 189 employed to provide a powerful prior. Standard diffusion-based approaches draw $\mathbf{x}_T \sim p(\mathbf{x}_T)$
 190 from a Gaussian noise prior and iteratively denoise through intermediate timesteps, aiming to sample
 191 from $p(\mathbf{x}_T | \mathbf{y}), p(\mathbf{x}_{T-1} | \mathbf{y}), \dots, p(\mathbf{x}_0 | \mathbf{y})$ in sequence. The key challenge is evaluating the
 192 intractable likelihood $p(\mathbf{y}|\mathbf{x}_t)$ at each intermediate timestep t . To address this, iterative guidance
 193 methods (Kawar et al., 2022; Wang et al., 2023; Chung et al., 2023; Song et al., 2023; Rozet et al.,
 194 2024; Song et al., 2024; Zhang et al., 2024) introduce approximations and apply likelihood corrections
 195 of $\mathbf{x}_t \leftarrow \mathbf{x}_t + \eta \nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$. However, these gradient-based corrections systematically
 196 push intermediate states \mathbf{x}_t away from the distribution on which the denoiser is trained, leading to
 197 what we refer to as the *manifold feasibility problem*. Following SITCOM (Alkhouri et al., 2025b),
 198 we formalize this issue as:

199 **Definition 3.1** (Manifold Feasibility). *For a pretrained diffusion model, let $p_t(\mathbf{x}_t)$ denote the
 200 marginal distribution at noise level t , and let \mathcal{M}_t be its high-probability generative manifold. An
 201 inverse-problem solver maintains manifold feasibility if the intermediate states $\{\mathbf{x}_t\}$ fed into the
 202 denoiser remain close to \mathcal{M}_t for all t , ensuring the final reconstruction \mathbf{x}_0 lies on the learned data
 203 manifold.*

204 Geometrically, standard guidance methods update \mathbf{x}_t using the likelihood gradient $\nabla_{\mathbf{x}_t} \log p(\mathbf{y}|\mathbf{x}_t)$.
 205 In high-dimensional spaces, this gradient vector often contains components orthogonal to the local
 206 tangent space of the data manifold \mathcal{M}_t . Consequently, adding this gradient systematically pushes
 207 the state \mathbf{x}_t into low-probability regions (off-manifold), feeding out-of-distribution inputs to the
 208 denoiser and causing accumulated artifacts as shown in Figure 1 (a). To avoid such approximations,
 209 we propose posterior sampling by drawing from the initial noise space. The sampled noise is then
 210 unconditionally denoised to a clean image. We adopt unconditional DDIM for the denoising process,
 211 which treats the entire denoising trajectory as a deterministic mapping $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{x}_T)$, so the problem
 212 becomes evaluating the posterior distribution of noise (Xia et al., 2023), i.e., $p(\mathbf{x}_T|\mathbf{y})$. We refer
 213 to our approach as *noise-space sampling* because HMC updates are performed exclusively on the
 214 initial noise $\mathbf{x}_T \sim \mathcal{N}(0, I)$. This differs from image-space and iterative guidance methods that
 215 directly modify intermediate states \mathbf{x}_t using measurement-consistency gradients. Sampling from
 the noise space offers two advantages: (i) the prior $p(\mathbf{x}_T)$ is a simple Gaussian distribution, and (ii)
 the likelihood $p(\mathbf{y}|\mathbf{x}_T) = p(\mathbf{y}|\mathcal{D}(\mathbf{x}_T))$ is directly accessible without intermediate approximations.

We use HMC for efficient posterior sampling in the noise space. To strictly justify our sampling objective, we formulate the inference process as a latent variable model where the initial noise \mathbf{x}_T is the sole latent variable. We treat the unconditional DDIM process with N steps as a deterministic parameterized generator function, denoted as $\mathcal{D} : \mathbb{R}^n \rightarrow \mathbb{R}^n$, which maps \mathbf{x}_T to a clean image $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{x}_T)$. Under this formulation, the measurement generation process is defined by $\mathbf{x}_T \xrightarrow{\mathcal{D}} \hat{\mathbf{x}}_0 \xrightarrow{\mathcal{A}, \eta} \mathbf{y}$. Consequently, the conditional distribution of \mathbf{y} given \mathbf{x}_T depends entirely on the generated image $\hat{\mathbf{x}}_0$. The likelihood term is thus mathematically exact: $p(\mathbf{y}|\mathbf{x}_T) = p(\mathbf{y}|\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{x}_T)) = \mathcal{N}(\mathbf{y}; \mathcal{A}(\mathcal{D}(\mathbf{x}_T)), \sigma_y^2 \mathbf{I})$. This allows us to perform posterior sampling directly in the noise space using the exact gradient of the likelihood with respect to \mathbf{x}_T . Then we can compute the conditional score using Bayes' rule:

$$\nabla_{\mathbf{x}_T} \log p(\mathbf{x}_T|\mathbf{y}) = \nabla_{\mathbf{x}_T} \log p(\mathbf{x}_T) + \nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathbf{x}_T). \quad (7)$$

Since \mathbf{x}_T is Gaussian noise in the DDIM framework, the first term is simply

$$\nabla_{\mathbf{x}_T} \log p(\mathbf{x}_T) = -\nabla_{\mathbf{x}_T} \frac{\|\mathbf{x}_T\|^2}{2} = -\mathbf{x}_T. \quad (8)$$

For the case of Gaussian measurement noise, if the noise level σ_y^2 is known, the likelihood term becomes

$$\nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathbf{x}_T) = \nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathcal{D}(\mathbf{x}_T)) = -\nabla_{\mathbf{x}_T} \frac{\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2}{2\sigma_y^2}. \quad (9)$$

We define $p(\mathbf{y}|\mathbf{x}_T) = p(\mathbf{y}|\mathcal{D}(\mathbf{x}_T))$ by viewing the denoising trajectory as a deterministic mapping $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{x}_T)$. This term can be computed directly using automatic differentiation. Because $\mathcal{D}(\mathbf{x}_T)$ results from a multi-step denoising process, backpropagating through multiple score networks can be computationally expensive. Following Wang et al. (2024), we illustrate in Appendix A.9 that accurate samples can still be obtained with as few as two denoising steps.

Algorithm 1: N-HMC

Require: # HMC iterations K , # leapfrog steps L , initial integration step size δ , measurement noise schedule $\{\sigma_{y,k}\}$, \mathbf{x}_T , \mathbf{y} , \mathcal{A} , γ

1: **for** $k = 0$ to $K - 1$ **do**

2: **repeat**

3: $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ // Initial momentum

4: $\hat{\mathbf{x}}_0 = \text{DDIM}(\mathbf{x}_T)$

5: $H_0 = \frac{1}{2}\|\mathbf{x}_T\|^2 + \frac{1}{2\sigma_{y,k}^2}\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2 + \frac{1}{2}\mathbf{p}^\top \mathbf{p}$ // Current Hamiltonian

6: $\mathbf{x}_T^* \leftarrow \mathbf{x}_T$ // Initialize proposal \mathbf{x}_T

7: **for** $l = 0$ to $L - 1$ **do**

8: $\mathbf{p} \leftarrow \mathbf{p} - \frac{\delta}{2} \left(\mathbf{x}_T^* + \frac{1}{2\sigma_{y,k}^2} \nabla_{\mathbf{x}_T^*} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2 \right)$ // Update momentum

9: $\mathbf{x}_T^* \leftarrow \mathbf{x}_T^* + \delta \mathbf{p}$ // Update \mathbf{x}_T^*

10: $\hat{\mathbf{x}}_0^* = \text{DDIM}(\mathbf{x}_T^*)$

11: $\mathbf{p} \leftarrow \mathbf{p} - \frac{\delta}{2} \left(\mathbf{x}_T^* + \frac{1}{2\sigma_{y,k}^2} \nabla_{\mathbf{x}_T^*} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2 \right)$ // Update momentum

12: **end for**

13: $H_1 = \frac{1}{2}\|\mathbf{x}_T^*\|^2 + \frac{1}{2\sigma_{y,k}^2}\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2 + \frac{1}{2}\mathbf{p}^\top \mathbf{p}$ // Proposal Hamiltonian

14: $u \sim \text{Unif}(0, 1)$

15: **if** $u < \exp(H_0 - H_1)$ **then**

16: Accept proposal

17: **else**

18: $\delta \leftarrow \gamma \delta$ // Anneal step size δ

19: **end if**

20: **until** Proposal accepted

21: $\mathbf{x}_T \leftarrow \mathbf{x}_T^*$ // Accept the proposal

22: **end for**

23: **return** \mathbf{x}_T

Once the conditional score $\nabla_{\mathbf{x}_T} \log p(\mathbf{x}_T | \mathbf{y})$ is computed, our method proceeds with standard Hamiltonian Monte Carlo (HMC) sampling. We use the identity matrix as the mass matrix for momentum sampling. During implementation, we observed that the initial noise may lie in regions of very low posterior probability, which forces HMC to adopt a tiny step size in order to maintain a proper acceptance rate. To address this issue, we use an annealing schedule for σ_y , allowing \mathbf{x}_T to explore the noise space more freely with a larger step size in the start-up stage. Once σ_y gradually declines to the target level, posterior samples are collected. The complete procedure is summarized in Algorithm 1, along with the unconditional DDIM denoising process in Algorithm 2.

3.2 ROBUSTNESS TO MEASUREMENT NOISE

An additional benefit of N-HMC over MAP methods is that the Gaussian prior acts as a regularization term in the noise space, keeping the noise vector \mathbf{x}_T close to the hypersphere of radius \sqrt{n} . Therefore, N-HMC produces samples that are *robust to measurement noise*, as justified by Proposition 1. For simplicity, we assume Gaussian measurement noise and that the forward operator \mathcal{A} is approximately linear along the clean image manifold.

Proposition 1. Assume that the distribution of the decoded sample \mathbf{x}_0 around the ground truth \mathbf{x}_0^* is well-approximated by a Gaussian distribution $p_\theta(\hat{\mathbf{x}}_0) \approx \mathcal{N}(\hat{\mathbf{x}}_0; \mathbf{x}_0^*, \sigma_0^2 \mathbf{I}_n)$. Then, the residual $\mathbf{y} - \mathcal{A}\hat{\mathbf{x}}_0$ satisfies

$$\mathbb{E}_{(\hat{\mathbf{x}}_0, \mathbf{y}) \sim p_\theta(\hat{\mathbf{x}}_0, \mathbf{y} | \mathbf{x}_0^*)} \|\mathbf{y} - \mathcal{A}\hat{\mathbf{x}}_0\|^2 = \sigma_y^2 \text{tr}(\mathbf{B}\mathbf{B}^\top) + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top),$$

where

$$\Sigma_{\text{post}} = \left(\frac{\mathbf{A}^\top \mathbf{A}}{\sigma_y^2} + \frac{\mathbf{I}_n}{\sigma_0^2} \right)^{-1}, \quad \mathbf{B} = \left(\mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2} \right),$$

and \mathbf{I}_m is the $m \times m$ identity matrix. m denotes the dimension of \mathbf{y} .

The expected residual decomposes into two contributions: a noise-dependent term that appears only when measurement noise is present, and a second term that persists in all settings due to intrinsic uncertainty of prior diffusion models. In Corollary 1.1 below, we show that both terms behave in a way that yields a residual whose magnitude matches the true measurement noise.

Corollary 1.1 Under the assumptions of Proposition 1, if $\sigma_0/\sigma_y \ll 1$, the residual $\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)$ satisfies

$$\mathbb{E}_{(\hat{\mathbf{x}}_0, \mathbf{y}) \sim p_\theta(\hat{\mathbf{x}}_0, \mathbf{y} | \mathbf{x}_0^*)} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2 \rightarrow m\sigma_y^2.$$

In other words, the magnitude of residual aligns with the true known level of measurement noise, indicating that N-HMC remains robust and does not overfit to noise.

3.3 NOISE-ADAPTIVE NHMC

In practice, the type and level of measurement noise are often unknown, making the likelihood term $p(\mathbf{y} | \mathbf{x}_T)$ intractable. To address this, other methods usually have tunable hyperparameters that control the strength of the likelihood term or use task-specific early stopping criteria. Instead of this heuristic approach, we introduce a noise-adaptive sampling method, *NA-NHMC*, which extends N-HMC to the *unknown noise setting without any additional hyperparameter tuning*.

We treat the noise variance as a latent variable and adopt the Jeffreys prior, a principled noninformative choice due to its parameterization invariance. It is scale-invariant and represents maximal uncertainty about the noise level, making it appropriate when no prior information about σ_y is available. It can also be viewed as the limiting case of an Inverse-Gamma prior $\sigma_y^2 \sim \text{Inv-}\Gamma(\alpha, \beta)$ as $\alpha, \beta \rightarrow 0$. The Inverse-Gamma distribution is the conjugate prior for the variance of a Gaussian likelihood. Proposition 2 characterizes the resulting behavior under additional assumptions.

$$p(\sigma_y^2) \sim \frac{1}{\sigma_y^2}.$$

324
 325 **Proposition 2** Under the assumptions of Proposition 1 and that the pretrained diffusion model
 326 unconditionally generates images that lie on the high-quality manifold ($\sigma_0/\sigma_y \ll 1$), then the update
 327 rule of NA-NHMC follows:

328
$$\nabla_{\mathbf{x}_T} \log p(\mathbf{y} | \mathbf{x}_T)_{\text{NA-NHMC}} = -\frac{1}{2\sigma_y^2} \nabla_{\mathbf{x}_T} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2.$$

 329
 330

331 By marginalizing σ_y^2 , the likelihood term becomes
 332

333
$$p(\mathbf{y} | \mathbf{x}_T) = \int_0^\infty p(\mathbf{y} | \mathbf{x}_T, \sigma_y^2) p(\sigma_y^2) d\sigma_y^2 \quad (10)$$

 334
 335

336
$$\propto \left(\frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2 \right)^{-m/2}. \quad (11)$$

 337
 338

339 where m denotes the dimensionality of the measurement
 340 space. The derivation of this expression is provided in
 341 Appendix A.1. Substituting this marginalized likelihood
 342 into the N-HMC framework yields our proposed noise-
 343 adaptive Algorithm 3. Proposition 2 below shows that,
 344 with an appropriate measurement noise prior, the likeli-
 345 hood term $\nabla_{\mathbf{x}_T} \log p(\mathbf{y} | \mathbf{x}_T)$ of NA-NHMC is identical
 346 to that of N-HMC (with known noise level).

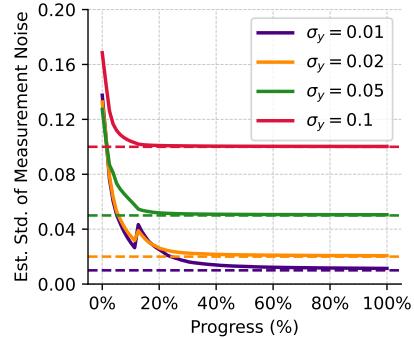
347 Figure 2 demonstrates Proposition 2 in practice. All ex-
 348 periments use an identical setup across different noise
 349 levels, highlighting that our method does not require any
 350 hyperparameter tuning for a specific noise level. Despite
 351 the absence of such tunable parameters, the estimated
 352 standard deviation of the measurement noise, computed
 353 as $\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|/\sqrt{m}$, closely matches the true, unknown
 354 noise level σ_y . This confirms that NA-NHMC effectively
 355 adapts to varying noise without specific tuning. The com-
 356 plete NA-NHMC is summarized in Algorithm 3.

357 The flexibility of NA-NHMC goes beyond noise-level ro-
 358 bustness. Although NA-NHMC is formulated assuming
 359 Gaussian measurement noise, experiments with alter-
 360 native noise types (Section 4.3) show that the method re-
 361 mains effective across other common noise distributions,
 362 demonstrating its broader robustness.

363 4 EXPERIMENTS

364 Following previous approaches (Wang et al., 2024) (Zhang et al., 2024), we evaluate our method
 365 on two datasets: FFHQ 256 × 256 (Karras et al., 2019) and ImageNet 256 × 256 (Deng et al.,
 366 2009), using 100 images from the validation set of each dataset. We utilize the same pretrained DM
 367 trained by Chung et al. (2023) for FFHQ and by Dhariwal & Nichol (2021) for ImageNet except for
 368 ReSample. For ReSample, we use a pretrained LDM by Rombach et al. (2022). All measurements
 369 are corrupted by additive Gaussian noise with standard deviation σ_y .

370 We compare our method against several representative baselines, including DiffPIR (Zhu et al.,
 371 2023), RED-diff (Mardani et al., 2023), DPS (Chung et al., 2023), DAPS (Zhang et al., 2024),
 372 ReSample (Song et al., 2024), **SITCOM** (Alkhouri et al., 2025b), and DMPlug (Wang et al., 2024).
 373 The implementation details for all the baseline methods are provided in Appendix A.5. We evaluate
 374 reconstruction quality using three standard metrics: peak signal-to-noise ratio (PSNR), structural
 375 similarity index measure (SSIM) (Wang et al., 2004), and Learned Perceptual Image Patch Similarity
 376 (LPIPS) (Zhang et al., 2018).



377 Figure 2: Gaussian deblurring task on
 378 FFHQ (256 × 256) with varying
 379 measurement noise levels σ_y . The esti-
 380 mated standard deviation of mea-
 381 surement noise $\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|/\sqrt{m}$ demon-
 382 strates that our noise-adaptive method
 383 accurately recovers the true σ_y (indi-
 384 cated by dashed line) without overfitting
 385 across different noise levels.

378
379

4.1 EXPERIMENT RESULTS

380
381
382
383
384

Linear IPs. We evaluate our approach on four linear inverse problems. For super-resolution tasks, we consider both $4\times$ and $16\times$ downsampling using 4×4 and 16×16 average pooling operations, respectively. We also examine random inpainting with 92% of pixels randomly masked, and anisotropic Gaussian deblurring using blur kernels with standard deviations of 20 and 1 in orthogonal directions. The results for linear IPs are presented in Tables 10-13.

385
386
387
388
389

Nonlinear IPs. We further assess performance on three challenging nonlinear inverse problems. The first is nonlinear deblurring using encoded blur kernels from Tran et al. (2021). The second is phase retrieval, where only the Fourier magnitude is observed as measurements. Finally, we consider HDR reconstruction, which aims to recover images with a higher dynamic range by a factor of 2 from tone-mapped observations. The results for nonlinear IPs are presented in Tables 1, 2, 14, 15.

390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415

Main Results. Our method achieves comparable or superior performance across most tasks, as measured by PSNR and SSIM on both the FFHQ and ImageNet datasets. Notably, the improvement over SOTA methods is more pronounced for nonlinear tasks, which are substantially more challenging than linear IPs. Many existing SOTA approaches are MAP-based by design (e.g., ReSample, SITCOM, and DMPlug) or become MAP-like heuristically (e.g., DAPS). While these methods perform well in low-noise regimes, they often overfit when the noise level is higher. Since the noise levels used in our experiments ($\sigma_y = 0.05, 0.20$) exceed those commonly reported in prior work, our results further demonstrate that the proposed noise-adaptive method is more robust and consistently outperforms alternatives across most tasks and metrics without any hyperparameter tuning. Figure 4 contains visual examples for the nonlinear deblurring problem. See Appendix A.12 for more examples. A fundamental distinction lies in the generalization capability across diverse degradation conditions. Standard guidance-based methods (e.g., DPS) inherently rely on manual hyperparameter calibration to balance measurement fidelity against the diffusion prior. As evidenced in Table 4, the optimal step size is highly task-dependent (ranging from $\zeta = 0.4$ to $\zeta = 10.0$), meaning a static configuration fails to generalize. In contrast, NA-NHMC derives its dynamics from the marginalized posterior, which effectively acts as an automatic gradient normalization mechanism. This structural advantage allows a single configuration to robustly generalize across varying tasks and noise levels without task-specific recalibration.

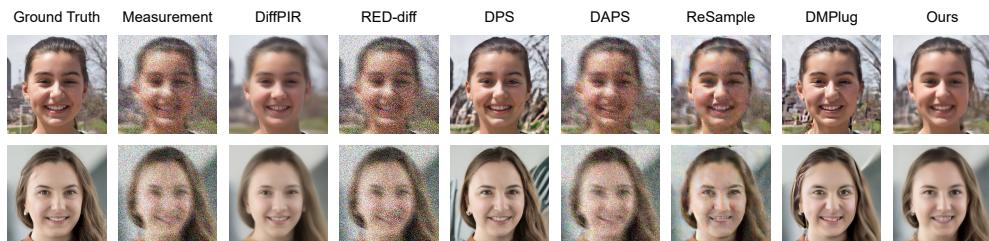
416
417
418
419
420
421
422
423424
425
426
427
428

Figure 4: Nonlinear deblurring results on FFHQ (256 \times 256) dataset with $\sigma_y = 0.2$. Visual comparison across state-of-the-art methods shows our approach produces high-quality reconstructions with sharp details and minimal artifacts.

429
430
431

4.2 HIGHLY ILL-POSED IPs: PHASE RETRIEVAL

Another challenge commonly faced by both MAP and sampling-based methods is becoming trapped in a local mode, particularly in highly multimodal IPs such as phase retrieval. Figure 5 illustrates this

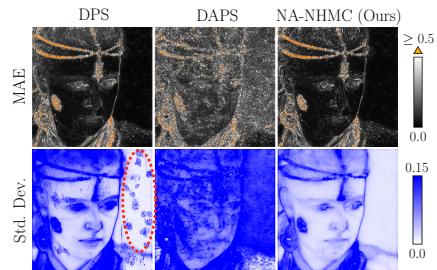


Figure 3: Comparative results are averaged over 100 independent runs. (Top) Mean absolute error (MAE) heatmaps. (Bottom) Standard deviation heatmaps across runs. Our method achieves the lowest standard deviation compared to DPS and DAPS, indicating reduced sensitivity to initialization.

manual hyperparameter calibration to balance measurement fidelity against the diffusion prior. As evidenced in Table 4, the optimal step size is highly task-dependent (ranging from $\zeta = 0.4$ to $\zeta = 10.0$), meaning a static configuration fails to generalize. In contrast, NA-NHMC derives its dynamics from the marginalized posterior, which effectively acts as an automatic gradient normalization mechanism. This structural advantage allows a single configuration to robustly generalize across varying tasks and noise levels without task-specific recalibration.

issue. While DPS and DMPlug occasionally recover the correct solution, most initializations converge to spurious local modes and are thus counted as failures. In contrast, our method incorporates early exploration through σ_y scheduling, making it more robust to initialization and substantially more likely to recover the global solution.

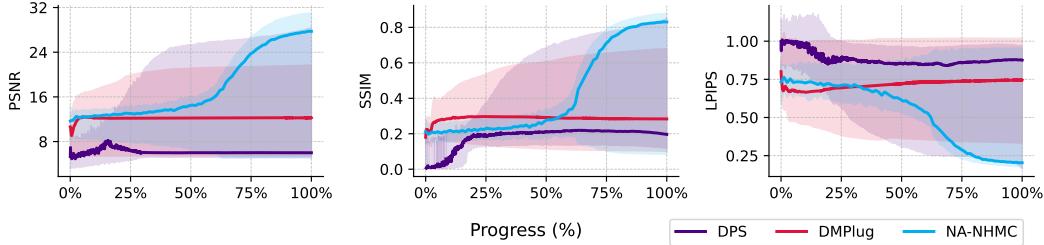


Figure 5: Phase retrieval task on FFHQ (256×256) with $\sigma_y = 0.01$. Each curve shows the median performance, with shaded areas denoting the 5th–95th percentile interval. Our method successfully solves the IP at a much higher rate than DPS and DMPlug. This is due to the annealing schedule of σ_y that allows for initial exploration of the noise space, resulting in a lower probability of being stuck on a local mode.

We quantify robustness to initialization using the standard deviation map in Figure 3. Our method achieves a mean absolute error (MAE) comparable to that of DPS, but with substantially lower pixel-wise standard deviation. While DPS can, on average, produce accurate reconstructions, its performance is sensitive to initialization and may introduce artifacts in both the face and background (red circle). In contrast, such artifacts never appear in any of the 100 runs with our method. Notably, despite exhibiting lower overall uncertainty, our method still assigns uncertainty in complex regions, which aligns with areas of high MAE (orange).

Table 1: Non-linear IPs Results on FFHQ (256×256) with Gaussian Noise $\sigma_y = 0.05$. (**Bold**: best, underline: second best)

	Nonlinear Deblurring			Phase Retrieval			HDR Reconstruction		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	26.12	0.743	0.289	16.77	<u>0.482</u>	0.543	25.20	0.814	0.223
RED-diff	18.12	0.217	0.680	11.83	0.213	0.769	21.44	0.525	0.458
DPS	23.26	0.672	0.300	10.87	0.296	0.714	<u>27.46</u>	0.849	0.168
DAPS	27.00	0.736	0.283	<u>18.52</u>	0.414	0.528	26.03	0.758	0.259
ReSample	24.57	0.637	0.432	13.95	0.377	0.677	23.65	0.722	0.386
SITCOM	24.97	0.569	0.328	11.89	0.216	0.723	26.97	0.753	0.256
DMPlug	27.15	0.784	0.266	-	-	-	25.17	0.783	0.260
NA-NHMC (ours)	27.66	<u>0.792</u>	<u>0.249</u>	19.30	<u>0.554</u>	0.482	28.45	0.849	<u>0.217</u>

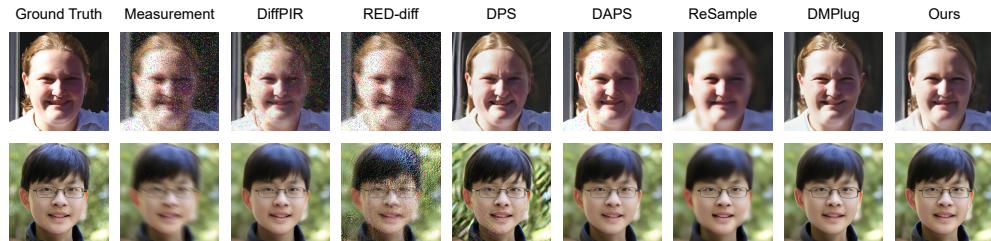
Table 2: Non-linear IPs on ImageNet (256×256) with Gaussian Noise $\sigma_y = 0.05$. (**Bold**: best, underline: second best)

	Nonlinear Deblurring			HDR Reconstruction		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	24.24	<u>0.638</u>	0.381	23.29	0.730	0.273
RED-diff	17.94	0.244	0.623	20.98	0.524	0.415
DPS	17.60	0.427	0.482	<u>25.31</u>	<u>0.763</u>	0.248
DAPS	<u>24.28</u>	0.632	0.404	23.57	0.709	0.283
SITCOM	24.00	0.556	<u>0.355</u>	24.76	0.708	0.276
DMPlug	22.30	0.576	0.421	20.61	0.562	0.431
NA-NHMC (ours)	24.98	0.694	0.308	25.86	0.779	<u>0.253</u>

4.3 ROBUSTNESS TO UNKNOWN MEASUREMENT NOISE

In practice, the measurement noise may be unknown and its type may not be Gaussian. This can pose a problem as many methods require multiple hyperparameters that are tuned for a specific level

486 and type of measurement noise. In this section, we evaluate our method’s robustness to unknown
 487 measurement noise on two tasks and two noise types: impulse and speckle. For impulse noise,
 488 each pixel in each channel is randomly replaced by 0 or 1 with a probability $p/2$ each, where $p \sim$
 489 $\text{Unif}(0, 0.2)$. For speckle noise, the noise takes the form $y(1 + \epsilon)$, where y is the measurement
 490 tensor and $\epsilon \sim \text{Unif}(0, 0.4)$. Table 3 shows that our method achieves superior performance on most
 491 metrics while using the exact same hyperparameters as the Gaussian noise experiment. As illustrated
 492 in Figure 6, methods such as DiffPIR, which do not suffer from noise overfitting in the Gaussian
 493 setting, now struggle with impulse noise. In contrast, even under high noise levels, NA-NHMC
 494 remains robust, producing high-quality reconstructions.



504 Figure 6: Nonlinear deblurring results on FFHQ (256 × 256) dataset under different noise conditions.
 505 (Top) Impulse noise. (Bottom) Speckle noise.

509 Table 3: The Performance Comparison with Different Types and Levels of Measurement Noise on
 510 FFHQ (256 × 256). (**Bold**: best, underline: second best)

	Super Resolution ($\times 4$)						Nonlinear Deblurring					
	Impulse			Speckle			Impulse			Speckle		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	19.54	0.492	0.549	25.91	0.733	0.324	21.00	0.402	0.526	25.96	0.731	0.299
RED-diff	15.15	0.341	0.692	21.81	0.481	0.516	13.82	0.109	0.781	18.59	0.269	0.650
DPS	21.99	0.581	<u>0.395</u>	<u>27.00</u>	<u>0.761</u>	0.246	21.64	0.595	<u>0.322</u>	23.42	0.678	0.302
DAPS	15.00	0.361	0.702	24.48	0.597	0.442	17.94	0.259	0.657	26.46	0.698	0.308
ReSample	22.98	0.639	0.483	26.17	0.733	0.387	22.74	0.616	0.471	24.64	0.692	0.409
SITCOM	16.56	0.392	0.628	23.16	0.600	0.425	17.92	0.259	0.612	26.49	0.667	<u>0.295</u>
DMPlug	19.52	0.358	0.562	26.79	0.689	0.336	<u>23.79</u>	<u>0.662</u>	0.335	25.82	<u>0.740</u>	0.308
NA-NHMC (ours)	23.42	<u>0.631</u>	0.382	27.36	0.768	<u>0.290</u>	24.16	0.677	<u>0.319</u>	27.97	0.796	0.253

5 CONCLUSION

524 In this work, we introduce N-HMC, a posterior sampler that operates in the noise space using reverse
 525 diffusion as a deterministic mapping from initial noise to a clean image, enabling posterior exploration
 526 while keeping proposals on the learned data manifold. The developed noise-adaptive variant,
 527 NA-NHMC, eliminates task-specific hyperparameter tuning by automatically adapting to unknown
 528 noise types and levels, which is a significant practical advantage over existing approaches. Theory
 529 establishes the correctness and efficiency of noise-space sampling, and experiments across diverse
 530 linear and nonlinear inverse problems on FFHQ and ImageNet show state-of-the-art reconstructions,
 531 robustness to initialization and noise, competitive runtimes with a few denoising steps, and
 532 uncertainty-aware estimates. The provided analysis and experiments also show that our method can
 533 mitigate measurement-consistency drift, noise overfitting, and local-mode collapse without relying
 534 on any task-specific hyperparameter tuning.

535 While NA-NHMC shows promising results, it incurs higher computational cost compared to other
 536 methods such as DPS due to HMC sampling. Additionally, its reliance on a small number of diffu-
 537 sion steps may limit its immediate applicability to more complex applications. Moreover, the high
 538 dimensionality of the posterior leads to long warmup phases before reaching stationarity. Future
 539 work could address these challenges by developing more efficient gradient estimation techniques
 and incorporating faster warmup strategies that relax the requirement of exact detailed balance.

540 REFERENCES
541

542 Ismail Alkhouri, Shijun Liang, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar,
543 and Rongrong Wang. Sitcom: Step-wise triple-consistent diffusion sampling for inverse prob-
544 lems, 2025a. URL <https://arxiv.org/abs/2410.04479>.

545 Ismail Alkhouri, Shijun Linag, Cheng-Han Huang, Jimmy Dai, Qing Qu, Saiprasad Ravishankar,
546 and Rongrong Wang. Sitcom: Step-wise triple-consistent diffusion sampling for inverse prob-
547 lems. In *International Conference on Machine Learning (ICML)*, 2025b.

548 Benjamin Boys, Mark Girolami, Jakiw Pidstrigach, Sebastian Reich, Alan Mosca, and O. Deniz
549 Akyildiz. Tweedie moment projected diffusions for inverse problems, 2024. URL <https://arxiv.org/abs/2310.06721>.

550 Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte*
551 *Carlo*. Chapman and Hall/CRC, May 2011. ISBN 9780429138508. doi: 10.1201/b10905. URL
552 <http://dx.doi.org/10.1201/b10905>.

553 Hamadi Chihoui and Paolo Favaro. Diffusion image prior. *arXiv preprint arXiv:2503.21410v1*,
554 2025a. URL <https://arxiv.org/abs/2503.21410>. cs.CV.

555 Hamadi Chihoui and Paolo Favaro. Diffusion image prior. In *Proceedings of the IEEE/CVF Inter-
556 national Conference on Computer Vision (ICCV)*, pp. 24636–24644, October 2025b.

557 Jiaoyue Chu, Chenhe Du, Xiyue Lin, Xiaoqun Zhang, Lihui Wang, Yuyao Zhang, and Hongjiang
558 Wei. Highly accelerated mri via implicit neural representation guided posterior sampling
559 of diffusion models. *Medical Image Analysis*, 100:103398, 2025. ISSN 1361-8415. doi:
560 <https://doi.org/10.1016/j.media.2024.103398>. URL <https://www.sciencedirect.com/science/article/pii/S1361841524003232>.

561 Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models
562 for inverse problems using manifold constraints. *Advances in Neural Information Processing
563 Systems*, 35:25683–25696, 2022.

564 Hyungjin Chung, Jeongsol Kim, Michael Thompson Mccann, Marc Louis Klasky, and Jong Chul
565 Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh Interna-
566 tional Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=OnD9zGAGT0k>.

567 Duy-Phuong Dao, Hyung-Jeong Yang, and Jahae Kim. Conditional diffusion model for longitudinal
568 medical image generation, 2024. URL <https://arxiv.org/abs/2411.05860>.

569 Giannis Daras, Hyungjin Chung, Chieh-Hsin Lai, Yuki Mitsufuji, Jong Chul Ye, Peyman Milan-
570 far, Alexandros G. Dimakis, and Mauricio Delbracio. A survey on diffusion models for inverse
571 problems, 2024. URL <https://arxiv.org/abs/2410.00083>.

572 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
573 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,
574 pp. 248–255, 2009. doi: 10.1109/CVPR.2009.5206848.

575 Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021. URL
576 <https://arxiv.org/abs/2105.05233>.

577 Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo.
578 *Physics letters B*, 195(2):216–222, 1987.

579 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models, 2020. URL
580 <https://arxiv.org/abs/2006.11239>.

581 Joel Janai, Fatma Güney, Aseem Behl, and Andreas Geiger. Computer vision for autonomous vehi-
582 cles: Problems, datasets and state of the art, 2021. URL <https://arxiv.org/abs/1704.05519>.

594 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
 595 adversarial networks, 2019. URL <https://arxiv.org/abs/1812.04948>.
 596

597 Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Ana-
 598 lyzing and improving the image quality of stylegan, 2020. URL <https://arxiv.org/abs/1912.04958>.
 599

600 Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration
 601 models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Ad-*
 602 *vances in Neural Information Processing Systems*, 2022. URL [https://openreview.net/](https://openreview.net/forum?id=kxXvopt9pWK)
 603 [forum?id=kxXvopt9pWK](https://openreview.net/forum?id=kxXvopt9pWK).

604 Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. A variational perspective on solving
 605 inverse problems with diffusion models, 2023. URL <https://arxiv.org/abs/2305.04391>.
 606

607 Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward
 608 Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*,
 609 21(6):1087–1092, 1953.

610 Jeffrey Ouyang-Zhang, Daniel J. Diaz, Adam R. Klivans, and Philipp Krähenbühl. Predicting a
 611 protein’s stability under a million mutations, 2023. URL <https://arxiv.org/abs/2310.12979>.
 612

613 Weize Quan, Jiaxi Chen, Yanli Liu, Dong-Ming Yan, and Peter Wonka. Deep learning-based image
 614 and video inpainting: A survey, 2024. URL <https://arxiv.org/abs/2401.03395>.
 615

616 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
 617 resolution image synthesis with latent diffusion models, 2022. URL <https://arxiv.org/abs/2112.10752>.
 618

619 François Rozet, Gérôme Andry, Francois Lanusse, and Gilles Louppe. Learning diffusion priors
 620 from observations by expectation maximization. In *The Thirty-eighth Annual Conference on*
 621 *Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=7v88Fh6iSM>.
 622

623 Bowen Song, Soo Min Kwon, Zecheng Zhang, Xinyu Hu, Qing Qu, and Liyue Shen. Solving inverse
 624 problems with latent diffusion models via hard data consistency. In *The Twelfth International*
 625 *Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=j8hdRqOUhN>.
 626

627 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models, 2022a.
 628 URL <https://arxiv.org/abs/2010.02502>.
 629

630 Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion
 631 models for inverse problems. In *International Conference on Learning Representations*, 2023.
 632 URL https://openreview.net/forum?id=9_gsMA8MRKQ.
 633

634 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution,
 635 2020. URL <https://arxiv.org/abs/1907.05600>.
 636

637 Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imag-
 638 ing with score-based generative models, 2022b. URL <https://arxiv.org/abs/2111.08005>.
 639

640 Phong Tran, Anh Tuan Tran, Quynh Phung, and Minh Hoai. Explore image deblurring via encoded
 641 blur kernel space. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*
 642 (*CVPR*), pp. 11951–11960, 2021. doi: 10.1109/CVPR46437.2021.01178.
 643

644 Hengkang Wang, Xu Zhang, Taihui Li, Yuxiang Wan, Tiancong Chen, and Ju Sun. DMPlug: A
 645 plug-in method for solving inverse problems with diffusion models. In *The Thirty-eighth Annual*
 646 *Conference on Neural Information Processing Systems*, 2024. URL <https://openreview.net/forum?id=81IFFsfQUj>.
 647

648 Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion
 649 null-space model. In *The Eleventh International Conference on Learning Representations*, 2023.
 650 URL <https://openreview.net/forum?id=mRieQgMtNTQ>.

651 Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error
 652 visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
 653 doi: 10.1109/TIP.2003.819861.

654 Yingzhi Xia and Nicholas Zabaras. Bayesian multiscale deep generative model for the solution of
 655 high-dimensional inverse problems. *Journal of Computational Physics*, 455:111008, 2022.

656 Yingzhi Xia, Qifeng Liao, and Jinglai Li. Vi-dgp: A variational inference method with deep genera-
 657 tive prior for solving high-dimensional inverse problems. *Journal of Scientific Computing*, 97(1):
 658 16, 2023.

659 Zhihang Xu, Yingzhi Xia, and Qifeng Liao. A domain-decomposed vae method for bayesian inverse
 660 problems. *International Journal for Uncertainty Quantification*, 14(3), 2024.

661 Kevin K. Yang, Zachary Wu, and Frances H. Arnold. Machine learning-guided directed evolution
 662 for protein engineering, 2019. URL <https://arxiv.org/abs/1811.10775>.

663 Kai Yi, Bingxin Zhou, Yiqing Shen, Pietro Liò, and Yu Guang Wang. Graph denoising diffusion for
 664 inverse protein folding, 2023. URL <https://arxiv.org/abs/2306.16819>.

665 Bingliang Zhang, Wenda Chu, Julius Berner, Chenlin Meng, Anima Anandkumar, and Yang
 666 Song. Improving diffusion inverse problem solving with decoupled noise annealing. *CoRR*,
 667 abs/2407.01521, 2024. URL <https://doi.org/10.48550/arXiv.2407.01521>.

668 Jiawei Zhang, Ziyuan Liu, Leon Yan, Gen Li, and Yuantao Gu. Improving diffusion-based inverse
 669 algorithms under few-step constraint via learnable linear extrapolation, 2025. URL <https://arxiv.org/abs/2503.10103>.

670 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 671 effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF Conference on Computer
 672 Vision and Pattern Recognition*, pp. 586–595, 2018. doi: 10.1109/CVPR.2018.00068.

673 Hongkai Zheng, Wenda Chu, Bingliang Zhang, Zihui Wu, Austin Wang, Berthy T Feng, Caifeng
 674 Zou, Yu Sun, Nikola Kovachki, Zachary E Ross, et al. Inversebench: Benchmarking plug-and-
 675 play diffusion priors for inverse problems in physical sciences. *arXiv preprint arXiv:2503.11043*,
 676 2025.

677 Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van
 678 Gool. Denoising diffusion models for plug-and-play image restoration, 2023. URL <https://arxiv.org/abs/2305.08995>.

686 ETHICS STATEMENT

687 We affirm adherence to the ICLR Code of Ethics. This work uses publicly available datasets and
 688 standard benchmarks; no human subjects, personal data, or sensitive attributes are collected or pro-
 689 cessed. We assessed potential risks (misuse, bias, fairness, privacy, and security) and found none
 690 beyond those commonly associated with generic image/model evaluation. All institutional and le-
 691 gal requirements were respected, and we disclose that we have no conflicts of interest or external
 692 sponsorship that could unduly influence the results.

694 REPRODUCIBILITY STATEMENT

695 We provide implementation details, hyperparameters, and training/evaluation protocols in the main
 696 text and appendix; key equations and assumptions are stated with complete proofs in the appendix.
 697 All datasets and preprocessing steps are described and referenced; code and configuration files will
 698 be released in an anonymous repository as supplementary material to enable exact reproduction
 699 of results. Where applicable, we reference the specific sections, tables, and figures that document
 700 model architectures, inference procedures, and evaluation metrics.

702 A APPENDIX
703704 A.1 ROBUST-N-HMC DERIVATION
705706 Assumptions
707708 1. Measurement noise $\eta \in \mathbb{R}^m$ follows gaussian distribution with unknown σ_y^2 :
709

710
$$p(\mathbf{y}|\mathbf{x}_T, \sigma_y^2) = \frac{1}{(2\pi\sigma_y^2)^{m/2}} \exp\left[-\frac{\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2}{2\sigma_y^2}\right]. \quad (12)$$

711

712 2. σ_y follows a Jeffreys prior distribution:
713

714
$$p(\sigma_y^2) \propto \frac{1}{\sigma_y^2}. \quad (13)$$

715

716 Marginalizing σ_y^2 yields
717

718
$$p(\mathbf{y}|\mathbf{x}_T) = \int_0^\infty p(\mathbf{y}|\mathbf{x}_T, \sigma_y^2) p(\sigma_y^2) d\sigma_y^2 \quad (14)$$

719

720
$$\propto \int_0^\infty \frac{1}{(2\pi\sigma_y^2)^{m/2}} \exp\left[-\frac{\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2}{2\sigma_y^2}\right] \frac{1}{\sigma_y^2} d\sigma_y^2 \quad (15)$$

721

722
$$\propto \int_0^\infty (\sigma_y^2)^{-\frac{m}{2}-1} \exp\left[-\frac{(1/2)\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2}{\sigma_y^2}\right] d\sigma_y^2 \quad (16)$$

723

724
$$\propto \left(\frac{1}{2}\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2\right)^{-m/2}. \quad (17)$$

725

726 And then, we have
727

728
$$\log p(\mathbf{y}|\mathbf{x}_T) = \left(-\frac{m}{2}\right) \log\left(\frac{1}{2}\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2\right). \quad (18)$$

729

730 A.2 PROOFS
731732 **Measurement model.** Let $\mathbf{x}_0^* \in \mathbb{R}^m$ denote the ground truth signal. The measurement is given by
733

734
$$\mathbf{y} = \mathcal{A}(\mathbf{x}_0^*) + \eta, \quad \eta \sim \mathcal{N}(0, \sigma_y^2 \mathbf{I}_m), \quad (19)$$

735

736 where $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ is the measurement operator and η represents Gaussian measurement noise.
737 In the following proofs, \mathcal{A} is assumed to be approximately linear around \mathbf{x}_0^* . Thus, $\mathcal{A}(\mathbf{x}_0) = \mathbf{A}\mathbf{x}_0$.
738739 **Generative model.** Consider the DDIM sampler defined by
740

741
$$\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{x}_T), \quad \mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I}_n), \quad (20)$$

742

743 where \mathcal{D} denotes the deterministic decoder via the diffusion model.
744745 **Lemma 1.** Product of two Gaussian probability density functions (PDFs).
746

747
$$q_1(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mu_1, \Sigma_1), \quad q_2(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mu_2, \Sigma_2).$$

748

749 Then, the product of $q_1(\mathbf{x})$ and $q_2(\mathbf{x})$ is proportional to a Gaussian PDF $\mathcal{N}(\mathbf{x}, \mu, \Sigma)$, where
750

751
$$\mu = \Sigma (\Sigma_1^{-1} \mu_1 + \Sigma_2^{-1} \mu_2), \quad \Sigma = (\Sigma_1^{-1} + \Sigma_2^{-1})^{-1}.$$

752

753 **Lemma 2.** Bias-variance decomposition of a random variable \mathbf{x} with mean μ and covariance ma-
754 trix Σ , i.e.,
755

$$\mathbb{E}[\|\mathbf{x} - \mathbf{a}\|^2] = \|\mu - \mathbf{a}\|^2 + \text{tr}(\Sigma).$$

756 **Proposition 1.** Assume that the distribution of the decoded sample $\hat{\mathbf{x}}_0$ around the ground truth
 757 \mathbf{x}_0^* is well-approximated by a Gaussian distribution $p_\theta(\hat{\mathbf{x}}_0) \approx \mathcal{N}(\hat{\mathbf{x}}_0; \mathbf{x}_0^*, \sigma_0^2 \mathbf{I}_n)$. Then, the residual
 758 $\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0$ satisfies

$$759 \quad 760 \quad \mathbb{E}_{(\hat{\mathbf{x}}_0, \mathbf{y}) \sim p_\theta(\hat{\mathbf{x}}_0, \mathbf{y} | \mathbf{x}_0^*)} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|^2 = \sigma_y^2 \text{tr}(\mathbf{B}\mathbf{B}^\top) + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top),$$

761 where

$$762 \quad 763 \quad \Sigma_{\text{post}} = \left(\frac{\mathbf{A}^\top \mathbf{A}}{\sigma_y^2} + \frac{\mathbf{I}_n}{\sigma_0^2} \right)^{-1}, \quad \mathbf{B} = \left(\mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2} \right).$$

764 *Proof* First, consider the distribution of $\hat{\mathbf{x}}_0$ given fixed \mathbf{x}_0^* and \mathbf{y}

$$765 \quad 766 \quad p_\theta(\hat{\mathbf{x}}_0 | \mathbf{y}) \stackrel{(a)}{\propto} p_\theta(\mathbf{y} | \hat{\mathbf{x}}_0) p_\theta(\hat{\mathbf{x}}_0) \quad (21)$$

$$767 \quad 768 \quad = \mathcal{N}(\mathbf{y}; \mathcal{A}(\hat{\mathbf{x}}_0), \sigma_y^2 \mathbf{I}_m) \mathcal{N}(\hat{\mathbf{x}}_0; \mathbf{x}_0^*, \sigma_0^2 \mathbf{I}_n) \quad (22)$$

$$769 \quad 770 \quad \stackrel{(b)}{=} \mathcal{N}(\hat{\mathbf{x}}_0; (\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top \mathbf{y}, \sigma_y^2 (\mathbf{A}^\top \mathbf{A})^{-1}) \mathcal{N}(\hat{\mathbf{x}}_0; \mathbf{x}_0^*, \sigma_0^2 \mathbf{I}_n) \quad (23)$$

$$771 \quad 772 \quad \stackrel{(c)}{=} \mathcal{N}(\hat{\mathbf{x}}_0; \boldsymbol{\mu}_{\text{post}}(\mathbf{y}), \Sigma_{\text{post}}) \quad (24)$$

$$773 \quad 774 \quad p_\theta(\mathbf{A}\hat{\mathbf{x}}_0 | \mathbf{y}) = \mathcal{N}(\hat{\mathbf{x}}_0; \mathbf{A}\boldsymbol{\mu}_{\text{post}}(\mathbf{y}), \mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top), \quad (25)$$

775 where (a) follows from Bayes' theorem. (b) is from local linearity of \mathcal{A} . (c) is the result of Lemma
 776 1 with

$$777 \quad 778 \quad \boldsymbol{\mu}_{\text{post}}(\mathbf{y}) = \Sigma_{\text{post}} \left(\frac{\mathbf{A}^\top \mathbf{y}}{\sigma_y^2} + \frac{\mathbf{x}_0^*}{\sigma_0^2} \right), \quad \Sigma_{\text{post}} = \left(\frac{\mathbf{A}^\top \mathbf{A}}{\sigma_y^2} + \frac{\mathbf{I}_n}{\sigma_0^2} \right)^{-1}.$$

780 The expected squared residual conditioned on \mathbf{y} is

$$781 \quad 782 \quad \mathbb{E}_{\hat{\mathbf{x}}_0 \sim p_\theta(\hat{\mathbf{x}}_0 | \mathbf{y})} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|^2 = \|\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_{\text{post}}(\mathbf{y})\|^2 + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top), \quad (26)$$

783 which is the result of Lemma 2. Then, integrate over \mathbf{y} conditioned on \mathbf{x}_0^*

$$784 \quad 785 \quad \mathbb{E}_{\hat{\mathbf{x}}_0, \mathbf{y} \sim p_\theta(\hat{\mathbf{x}}_0, \mathbf{y} | \mathbf{x}_0^*)} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|^2 \quad (27)$$

$$786 \quad 787 \quad = \mathbb{E}_{\mathbf{y} \sim q(\mathbf{y} | \mathbf{x}_0^*)} [\mathbb{E}_{\hat{\mathbf{x}}_0 \sim p_\theta(\hat{\mathbf{x}}_0 | \mathbf{y})} \|\mathbf{y} - \mathbf{A}\hat{\mathbf{x}}_0\|^2] \quad (28)$$

$$788 \quad 789 \quad = \mathbb{E}_{\mathbf{y} \sim q(\mathbf{y} | \mathbf{x}_0^*)} [\|\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_{\text{post}}(\mathbf{y})\|^2 + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top)] \quad (29)$$

$$790 \quad 791 \quad = \mathbb{E}_{\mathbf{y} \sim q(\mathbf{y} | \mathbf{x}_0^*)} \left[\left\| \left(\mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2} \right) \mathbf{y} - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{x}_0^*}{\sigma_0^2} \right\|^2 \right] + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top) \quad (30)$$

$$792 \quad 793 \quad \stackrel{(a)}{=} \left\| \left(\mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2} \right) \mathbf{A}\mathbf{x}_0^* - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{x}_0^*}{\sigma_0^2} \right\|^2 + \text{tr}(\mathbf{B}(\sigma_y^2 \mathbf{I}_m) \mathbf{B}^\top) + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top) \quad (31)$$

$$794 \quad 795 \quad = \sigma_y^2 \text{tr}(\mathbf{B}\mathbf{B}^\top) + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top), \quad (32)$$

796 where (a) is the result of Lemma 2, and $\mathbf{B} = \mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2}$.

797 **Corollary 1.1** Under the assumptions of Proposition 1, if $\sigma_0/\sigma_y \ll 1$, the residual $\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)$
 798 satisfies

$$800 \quad 801 \quad \mathbb{E}_{(\hat{\mathbf{x}}_0, \mathbf{y}) \sim p_\theta(\hat{\mathbf{x}}_0, \mathbf{y} | \mathbf{x}_0^*)} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2 \rightarrow m\sigma_y^2.$$

803 *Proof* If $\sigma_0/\sigma_y \ll 1$,

$$804 \quad 805 \quad \Sigma_{\text{post}} = \left(\frac{\mathbf{A}^\top \mathbf{A}}{\sigma_y^2} + \frac{\mathbf{I}_n}{\sigma_0^2} \right)^{-1} = \sigma_0^2 \left(\frac{\sigma_0^2 \mathbf{A}^\top \mathbf{A}}{\sigma_y^2} + \mathbf{I}_n \right)^{-1} \rightarrow \sigma_0^2 \mathbf{I}_n \quad (33)$$

$$806 \quad 807 \quad \mathbf{B} = \mathbf{I}_m - \frac{\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top}{\sigma_y^2} \rightarrow \mathbf{I}_m - \frac{\sigma_0^2 \mathbf{A}\mathbf{A}^\top}{\sigma_y^2} \rightarrow \mathbf{I}_m \quad (34)$$

$$808 \quad 809 \quad \mathbb{E} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2 = \sigma_y^2 \text{tr}(\mathbf{B}\mathbf{B}^\top) + \text{tr}(\mathbf{A}\Sigma_{\text{post}}\mathbf{A}^\top) \rightarrow m\sigma_y^2 \quad (35)$$

810
 811 **Proposition 2** Under the assumptions of Proposition 1 and that the pre-trained diffusion model
 812 unconditionally generates images that lie on the high quality manifold ($\sigma_0/\sigma_y \ll 1$), then the
 813 update rule of NA-NHMC follows:

814

$$\nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathbf{x}_T)_{\text{NA-NHMC}} = -\frac{1}{2\sigma_y^2} \nabla_{\mathbf{x}_T} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2$$

815

816
 817 *Proof* We can consider the likelihood term of NA-NHMC

818

$$\log p(\mathbf{y}|\mathbf{x}_T) = \left(-\frac{m}{2}\right) \log \left(\frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2\right) \quad (36)$$

819

820

$$\nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathbf{x}_T) = \left(-\frac{m}{2\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2}\right) \nabla_{\mathbf{x}_T} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2 \quad (37)$$

821

822

$$= \left(-\frac{1}{2\sigma_y^2}\right) \nabla_{\mathbf{x}_T} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2, \quad (38)$$

823

824
 825 which follows from Corollary 1.1. This likelihood term is exactly the same as that of N-HMC, where
 826 the true noise level σ_y is known.

827
 828
 829
 830
 831
 832
 833
 834
 835
 836
 837
 838
 839
 840
 841
 842
 843
 844
 845
 846
 847
 848
 849
 850
 851
 852
 853
 854
 855
 856
 857
 858
 859
 860
 861
 862
 863

864 A.3 PSEUDOCODE OF UNCONDITIONAL DDIM
865866 The DDIM method (Song et al., 2022a) we used in our experiment follows the Algorithm 2 below:
867868 **Algorithm 2:** DDIM
869

Require: # diffusion steps T , diffusion model s_θ , initial seed \mathbf{x}_T
1: **for** $t = T - 1$ to 0 **do**
2: $\hat{\epsilon}_{t+1} = s_\theta(\mathbf{x}_{t+1}, t+1)$ // Compute the score
3: $\hat{\mathbf{x}}_0(\mathbf{x}_{t+1}) = \frac{1}{\sqrt{\bar{\alpha}_{t+1}}} (\mathbf{x}_{t+1} - \sqrt{1 - \bar{\alpha}_{t+1}} \hat{\epsilon}_{t+1})$ // Predict $\hat{\mathbf{x}}_0$ with Tweedie's formula
4: $\hat{\mathbf{x}}_t = \sqrt{\bar{\alpha}_t} \hat{\mathbf{x}}_0(\mathbf{x}_{t+1}) + \sqrt{1 - \bar{\alpha}_t} \hat{\epsilon}_{t+1}$ // Unconditional DDIM step
5: **end for**
6: **return** \mathbf{x}

877

878

879 A.4 PSEUDOCODE OF THE NOISE-ADAPTIVE NHMC
880881 Following the reasoning in Proposition 2, we assume an uninformative prior on σ_y . Under this
882 assumption, the likelihood term can be written as
883

884
$$\log p(\mathbf{y}|\mathbf{x}_T) = \left(-\frac{m}{2} \right) \log \left(\frac{1}{2} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2 \right).$$

885
886

887 where $m = \dim(\mathbf{y})$. The factor 1/2 inside the logarithm can be omitted for the Hamiltonian and
888 gradient computations. The corresponding gradient is
889

890
$$\nabla_{\mathbf{x}_T} \log p(\mathbf{y}|\mathbf{x}_T) = \left(-\frac{m}{2\|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2} \right) \nabla_{\mathbf{x}_T} \|\mathbf{y} - \mathcal{A}(\mathcal{D}(\mathbf{x}_T))\|^2$$

891
892

893 **Algorithm 3:** NA-NHMC
894

Require: # HMC iterations K , # leapfrog steps L , initial integration step size δ , $\mathbf{x}_T, \mathbf{y}, \mathcal{A}, \gamma$
1: **for** $k = 0$ to $K - 1$ **do**
2: **repeat**
3: $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ // Initial momentum
4: $\hat{\mathbf{x}}_0 = \text{DDIM}(\mathbf{x}_T)$
5: $H_0 = \frac{1}{2} \|\mathbf{x}_T\|^2 + \frac{m}{2} \log (\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|^2) + \frac{1}{2} \mathbf{p}^\top \mathbf{p}$ // Current Hamiltonian
6: $\mathbf{x}_T^* \leftarrow \mathbf{x}_T$ // Initialize proposal \mathbf{x}_T
7: **for** $l = 0$ to $L - 1$ **do**
8: $\mathbf{p} \leftarrow \mathbf{p} - \frac{\delta}{2} \left(\mathbf{x}_T^* + \frac{m}{2\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2} \nabla_{\mathbf{x}_T^*} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2 \right)$ // Update momentum
9: $\mathbf{x}_T^* \leftarrow \mathbf{x}_T^* + \delta \mathbf{p}$ // Update \mathbf{x}_T^*
10: $\hat{\mathbf{x}}_0^* = \text{DDIM}(\mathbf{x}_T^*)$
11: $\mathbf{p} \leftarrow \mathbf{p} - \frac{\delta}{2} \left(\mathbf{x}_T^* + \frac{m}{2\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2} \nabla_{\mathbf{x}_T^*} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2 \right)$ // Update momentum
12: **end for**
13: $H_1 = \frac{1}{2} \|\mathbf{x}_T^*\|^2 + \frac{m}{2} \log (\|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0^*)\|^2) + \frac{1}{2} \mathbf{p}^\top \mathbf{p}$ // Proposal Hamiltonian
14: $u \sim \text{Unif}(0, 1)$
15: **if** $u < \exp(H_0 - H_1)$ **then**
16: Accept proposal
17: **else**
18: $\delta \leftarrow \gamma \delta$ // Anneal step size δ
19: **end if**
20: **until** Proposal accepted
21: $\mathbf{x}_T \leftarrow \mathbf{x}_T^*$ // Accept the proposal
22: **end for**
23: **return** \mathbf{x}_T

918 A.5 IMPLEMENTATION DETAILS FOR BASELINE METHODS
919920 **DiffPIR**921 Number of diffusion steps: 100
922 Number of optimization steps: 50
923924 We follow the recommended hyperparameter $\eta = 1.0$ and $\lambda = 7.0$ from Zhang et al. (2025). The
925 learning rate of the schedule-free AdamW optimizer is set to 0.1.926 **RED-diff**927 Number of optimization steps: 1000
928929 We follow the recommended the hyperparameter $\lambda = 0.25$ and an Adam optimizer with $lr = 0.5$ as
930 in Zhang et al. (2025)931 **DPS**932 Number of diffusion steps: 1000
933934 We follow the learning rate form in Chung et al. (2023) with ζ_i adjusted for different tasks, as shown
935 in Table 4.937 Table 4: Tuned learning rate ζ_i for DPS
938

	SR ($\times 4$)	SR ($\times 16$)	Inpainting (92%)	Gaussian Deblurring	Nonlinear Deblurring	Phase Retrieval	HDR
FFHQ	1.0	0.6	1.0	1.0	1.0	0.4	1.0
ImageNet	1.0	0.6	1.0	0.4	0.5	-	1.0

943 **DAPS**944 Number of diffusion steps: 250
945946 Number of ODE solver steps: 4
947948 We follow the hyperparameter settings of Zhang et al. (2024), as listed in Table 5, and adopt their
949 heuristic $\sigma_y = 0.01$ in place of the actual value. $\delta = 0.01$ for all tasks. The number of MCMC
950 sampling steps $N = 100$ for FFHQ (256×256) and $N = 40$ for ImageNet (256×256). Otherwise,
951 the hyperparameter for each task is the same for both datasets.952 Table 5: η_0 for DAPS
953

	SR ($\times 4$)	SR ($\times 16$)	Inpainting (92%)	Gaussian Deblurring	Nonlinear Deblurring	Phase Retrieval	HDR
	1e-4	1e-4	1e-4	1e-4	5e-5	5e-5	2e-5

954 **ReSample**955 Number of diffusion steps: 500
956957 For both noise levels ($\sigma_y = 0.05, 0.20$) tested in this paper, the recommended optimization steps
958 lead to overfitting to noise and poor performance. Instead, we used 50 steps for pixel optimization
959 and 25 steps for latent optimization.960 **SITCOM**961 Number of diffusion steps: 20
962963 We follow the hyperparameter settings of Alkhouri et al. (2025a), as listed in Table 6. The stopping
964 criterion δ for $\sigma_y \in \{0.05, 0.2\}$ is chosen as $0.051\sqrt{m}$ and $0.201\sqrt{m}$ respectively, with m denoting
965 the dimension of y .966 **DMPlug**967 Number of diffusion steps: 3
968

972
973
974
975 Table 6: Optimization Steps K for SITCOM
976
977

SR ($\times 4$)	SR ($\times 16$)	Inpainting (92%)	Gaussian Deblurring	Nonlinear Deblurring	Phase Retrieval	HDR
20	20	30	30	30	30	40

978
979 $t = [250, 500, 750]$

980 We set the Adam optimizer learning rate to 0.01. We follow the recommended stopping criteria in
981 Wang et al. (2024). For linear tasks, we use a window size = 10, patience = 100, and a maximum
982 iterations = 5000. For nonlinear tasks, we use the window size = 50 and patience = 300 with
983 maximum iterations = 10000.984
985

A.6 IMPLEMENTATION DETAILS FOR OUR METHOD

986 Number of diffusion steps: 2

987
988 $t = [375, 750]$

989 We implement NA-NHMC with the same hyperparameter configuration $L = 20$, $\delta_0 = 0.05$, $\gamma =$
990 0.95 for both datasets. For all tasks except phase retrieval, an initial step size $\delta_0 = 0.05$ and an
991 annealing schedule $\sigma_{y,k} = 0.5 + 2(1 - k/10)$ is applied during the first 10 HMC iterations, after
992 which, the noise-adaptive scheme (Algorithm 3) is used.993 For phase retrieval, we use initial step size $\delta_0 = 0.2$ and keep other hyperparameters unchanged. An
994 annealing schedule $\sigma_{y,k} = 1.0 + 20\sqrt{1 - k/50}$ is applied during the first 50 HMC iterations, after
995 which, the noise-adaptive scheme (Algorithm 3) is used.996 These annealing schedules are chosen to encourage sufficient exploration of the posterior in the early
997 stage. Our results are not sensitive to the exact schedule: using a slower schedule does not degrade
998 performance.1000
1001
1002

A.7 HYPERPARAMETER SENSITIVITY ANALYSIS

1003 We evaluate several hyperparameter choices for the Hamiltonian Monte Carlo sampler on the super-
1004 resolution ($\times 4$) inverse problem using the FFHQ 256×256 dataset. When varying a particular
1005 hyperparameter, all remaining hyperparameters are kept at their default values used in all other
1006 experiments. For different choices of the number of leapfrog steps L , we also adjust the number of
1007 HMC iterations to ensure that each setting uses the same amount of computational resources. The
1008 resulting performance is summarized in the tables below. The results indicate that the performance
1009 of NA-NHMC exhibits little sensitivity to the choice of hyperparameters.1010
1011 Table 7: Different step sizes ϵ for Super Resolution ($\times 4$) on FFHQ (256×256) with Gaussian
1012 Noise $\sigma_y = 0.05$. (**Bold**: best)

	0.02	0.05	0.10	0.15	0.20
PSNR \uparrow	27.12	27.29	27.31	27.31	27.31
SSIM \uparrow	0.745	0.770	0.771	0.771	0.772
LPIPS \downarrow	0.299	0.291	0.288	0.288	0.286

1013
1014 Table 8: Different number of leapfrog steps L for Super Resolution ($\times 4$) on FFHQ (256×256)
1015 with Gaussian Noise $\sigma_y = 0.05$. (**Bold**: best)

	10	15	20	25	30
PSNR \uparrow	26.86	27.18	27.29	27.36	27.34
SSIM \uparrow	0.749	0.765	0.770	0.771	0.777
LPIPS \downarrow	0.318	0.299	0.291	0.286	0.281

1026 Table 9: Step size decay factor γ for Super Resolution ($\times 4$) on FFHQ (256×256) with Gaussian
 1027 Noise $\sigma_y = 0.05$. (**Bold**: best)

	0.91	0.93	0.95	0.97	0.99
PSNR \uparrow	27.29	27.31	27.29	27.31	27.30
SSIM \uparrow	0.770	0.769	0.770	0.771	0.770
LPIPS \downarrow	0.291	0.289	0.291	0.288	0.289

1034 A.8 ADDITIONAL EXPERIMENT RESULTS

1036 Linear IPs results

1038 Table 10: Linear IPs on FFHQ (256×256) with Gaussian Noise $\sigma_y = 0.05$. (**Bold**: best, underline:
 1039 second best)

	Super Resolution ($\times 4$)			Super Resolution ($\times 16$)			Random Inpainting (92%)			Gaussian Deblurring		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	25.96	0.735	0.322	19.84	<u>0.541</u>	0.444	20.93	0.595	0.405	27.48	0.778	0.287
RED-diff	21.58	0.390	0.602	17.60	0.391	0.567	23.70	0.651	0.344	17.07	0.213	0.692
DPS	26.84	0.762	0.239	20.06	0.522	0.380	25.74	0.745	0.245	26.88	0.761	0.234
DAPS	24.58	0.559	0.514	17.28	0.420	0.541	25.68	0.685	0.331	23.34	0.478	0.474
ReSample	26.18	0.737	0.382	20.01	0.532	0.576	24.12	0.599	0.442	25.98	0.728	0.385
SITCOM	27.35	0.787	<u>0.268</u>	<u>20.82</u>	0.574	<u>0.400</u>	<u>26.56</u>	0.785	<u>0.266</u>	<u>27.94</u>	<u>0.796</u>	0.266
DMPlug	26.73	0.697	0.321	17.42	0.280	0.607	26.15	<u>0.769</u>	0.270	27.81	0.769	0.289
NA-NHMC (ours)	27.29	<u>0.770</u>	0.291	20.85	0.531	0.452	26.72	0.785	0.268	28.36	0.798	<u>0.259</u>

1049 Table 11: Linear IPs ImageNet (256×256) with Gaussian Noise $\sigma_y = 0.05$. (**Bold**: best, underline:
 1050 second best)

	Super Resolution ($\times 4$)			Super Resolution ($\times 16$)			Random Inpainting (92%)			Gaussian Deblurring		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	23.99	0.626	0.426	18.48	<u>0.387</u>	0.626	19.30	0.443	0.583	25.24	0.678	0.378
RED-diff	17.67	0.266	0.613	12.45	<u>0.151</u>	0.726	17.25	0.360	0.541	13.99	0.170	0.688
DPS	23.36	0.623	<u>0.345</u>	17.15	0.339	0.524	22.31	0.593	0.347	22.55	0.555	0.401
DAPS	23.86	0.568	0.461	14.29	0.139	0.753	23.23	0.585	0.432	24.52	0.558	0.423
SITCOM	<u>24.93</u>	0.684	0.318	18.58	0.404	0.525	<u>23.89</u>	0.684	0.320	<u>25.63</u>	0.712	0.311
DMPlug	24.52	<u>0.667</u>	0.378	16.74	0.311	0.590	23.49	0.668	0.358	23.55	0.605	0.433
NA-NHMC (ours)	24.99	0.665	0.355	19.09	0.396	0.580	24.10	0.676	0.324	25.76	0.699	0.327

1059 All experiments in Section 3.3 are repeated with a **higher level of Gaussian measurement noise**
 1060 ($\sigma_y = 0.20$). The results are shown below.

1062 Table 12: Linear IPs on FFHQ (256×256) with Gaussian Noise $\sigma_y = 0.20$. (**Bold**: best, underline:
 1063 second best)

	Super Resolution ($\times 4$)			Super Resolution ($\times 16$)			Random Inpainting (92%)			Gaussian Deblurring		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
DiffPIR	21.22	0.591	0.417	16.09	0.420	0.551	17.83	0.470	0.509	24.29	0.683	0.355
RED-diff	12.64	0.101	0.824	11.49	0.152	0.799	15.33	0.168	0.731	8.47	0.037	0.868
DPS	21.80	0.556	<u>0.385</u>	16.13	0.377	<u>0.507</u>	21.60	0.531	<u>0.404</u>	24.45	0.678	0.290
DAPS	13.48	0.121	0.792	17.08	0.420	0.549	20.64	0.353	0.587	8.12	0.031	0.862
ReSample	22.95	0.632	0.501	17.93	0.468	0.661	22.62	0.615	0.535	24.60	0.682	0.438
SITCOM	<u>23.04</u>	0.647	0.362	<u>17.49</u>	0.469	0.495	<u>23.23</u>	<u>0.653</u>	0.359	<u>24.97</u>	<u>0.709</u>	<u>0.323</u>
DMPlug	15.95	0.140	0.706	11.69	0.093	0.773	19.65	0.347	0.564	17.33	0.181	0.660
NA-NHMC (ours)	23.29	<u>0.636</u>	0.391	17.36	0.393	0.552	23.69	0.670	<u>0.365</u>	25.57	0.710	0.327

1073 A.9 ABLATION STUDIES

1076 Number of HMC iterations

1077 Sampling with HMC requires a warmup phase, since the initial noise x_T may be far from the solution.
 1078 As shown in Figure 7, the quality of sampled images improves monotonically with the number
 1079 of iterations, as expected. Performance begins to plateau after roughly 120 iterations. Unlike MAP-
 based methods such as ReSample (Song et al., 2024) and DMPlug (Wang et al., 2024), it does not

1080
 1081
 1082
 1083
Table 13: Linear IPs on ImageNet (256 × 256) with Gaussian Noise $\sigma_y = 0.20$. (**Bold**: best,
underline: second best)

	Super Resolution (×4)			Super Resolution (×16)			Random Inpainting (92%)			Gaussian Deblurring		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
DiffPIR	19.90	0.448	0.577	18.48	0.387	0.626	16.91	0.295	0.684	22.38	0.550	0.493
RED-diff	11.35	0.097	0.776	9.29	0.098	0.811	11.37	0.079	0.782	8.08	0.043	0.817
DPS	19.29	0.406	<u>0.481</u>	11.25	0.145	0.728	18.64	0.361	0.505	20.33	0.455	0.468
DAPS	13.71	0.151	0.760	14.27	0.139	0.755	18.89	0.248	0.598	9.06	0.060	0.789
SITCOM	<u>20.87</u>	0.490	0.458	16.16	<u>0.325</u>	<u>0.628</u>	<u>20.90</u>	<u>0.494</u>	<u>0.457</u>	22.73	0.584	0.402
DMPlug	18.58	0.412	0.500	9.53	0.125	0.782	19.15	0.447	0.474	<u>23.05</u>	<u>0.591</u>	0.423
NA-NHMC (ours)	21.53	0.510	0.492	<u>16.43</u>	0.276	0.675	21.80	0.536	0.456	23.45	0.597	<u>0.405</u>

1093
 1094
 1095
 1096
 1097
 1098
 1099
Table 14: Non-linear IPs on FFHQ (256 × 256) with Gaussian Noise $\sigma_y = 0.20$. (**Bold**: best,
underline: second best)

	Nonlinear Deblurring			Phase Retrieval			HDR Reconstruction		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
DiffPIR	<u>23.34</u>	0.641	0.374	16.76	0.482	0.543	21.85	0.694	0.344
RED-diff	12.85	0.063	0.816	10.07	0.061	0.855	16.73	0.222	0.649
DPS	22.83	<u>0.643</u>	0.307	10.60	0.267	0.701	<u>24.92</u>	<u>0.703</u>	<u>0.321</u>
DAPS	17.38	0.154	0.728	12.93	0.103	0.797	18.04	0.299	0.607
ReSample	23.30	0.635	0.477	12.51	0.335	0.712	22.51	0.677	0.428
SITCOM	16.26	0.173	0.656	10.19	0.082	0.810	20.11	0.346	0.534
DMPlug	22.08	0.544	0.437	-	-	-	16.17	0.473	0.481
NA-NHMC (ours)	24.89	0.705	<u>0.317</u>	<u>16.17</u>	<u>0.434</u>	<u>0.570</u>	26.61	0.793	0.271

1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
Table 15: Non-linear IPs on ImageNet (256 × 256) with Gaussian Noise $\sigma_y = 0.20$. (**Bold**: best,
underline: second best)

	Nonlinear Deblurring			HDR Reconstruction		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
DiffPIR	21.52	0.492	0.526	19.94	0.556	0.418
RED-diff	12.47	0.071	0.759	16.46	0.267	0.593
DPS	16.11	0.340	0.551	21.92	0.519	0.448
DAPS	17.84	0.201	0.617	18.08	0.351	0.536
SITCOM	14.49	0.156	0.668	19.82	0.441	0.500
DMPlug	21.80	<u>0.561</u>	0.420	20.54	0.562	0.430
NA-NHMC (ours)	22.61	0.585	0.382	24.12	0.701	0.320

1131
 1132
 1133

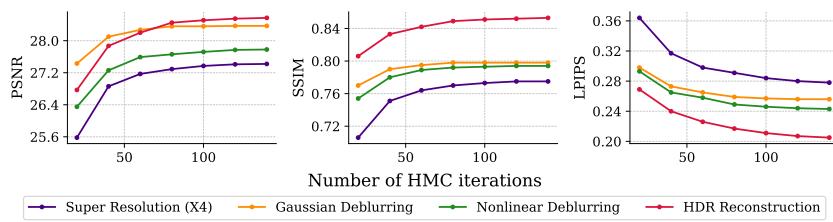


Figure 7: Performance of NA-NHMC across four tasks for FFHQ (256 × 256) as a function of the number of HMC iterations K . For all tasks, performance increases monotonically with more steps, but with diminishing improvements.

deteriorate beyond this point. This stability provides evidence that the prior term acts as an effective regularizer, preventing overfitting.

Number of diffusion steps and memory usage

In this section, we evaluate NA-NHMC with varying numbers of diffusion steps, using fixed parameters $K = 80$ and $L = 20$. Both runtime (in seconds) and memory usage (in GB) increase linearly with the number of steps. We ran all experiments on NVIDIA H200 GPU. The baseline cost is 90 seconds and 3.63 GB for two steps, with each additional step adding roughly 45 seconds and 1.84 GB. The quantitative evaluations are shown in Figure 8. While three diffusion steps achieve the highest PSNR and lowest LPIPS, the improvement over two steps is marginal. To avoid incurring roughly 50% additional runtime and memory overhead, we use two diffusion steps in all experiments.

Note that performance appears to decline when using more than three diffusion steps. This effect arises because the sampler converges more slowly to its stationary distribution as the number of diffusion steps increases. While increasing the number of HMC iterations could offset this effect, it would further amplify runtime costs to an impractical level.

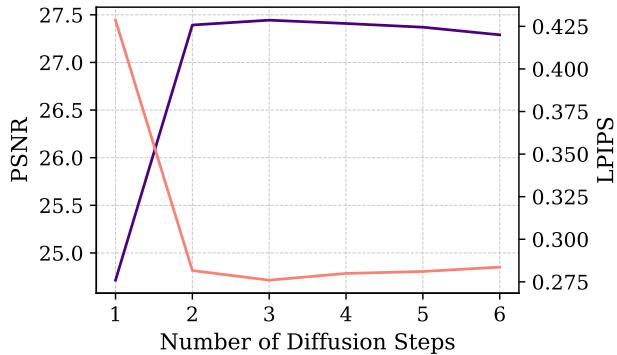


Figure 8: Performance of NA-NHMC on SR (×4) task for FFHQ (256 × 256) as a function of the number of diffusion steps. The initial step is fixed at $T = 750$ for all cases to avoid numerical instability, and the remaining steps are evenly spaced in $[0, 750]$.

Diffusion schedule

The pre-trained DMs used in this paper have 1000 diffusion steps. While other methods usually use evenly-spaced schedule with the first step being pure Gaussian noise ($\bar{a}_t = 0$), we found this choice to be numerically unstable for our few-step setting. Since we are using two steps for unconditional DDIM, the natural choice is to use timesteps in the middle. Thus, we choose $t = [375, 750]$, which is spread evenly and avoids numerical stability. Table 16 confirms that this schedule yields superior performance in PSNR and SSIM while being close to optimal for LPIPS. We adopt this diffusion schedule for all main experiments.

1188 Table 16: Performance of NA-NHMC on SR ($\times 4$) for FFHQ (256×256). Each schedule is defined
 1189 by two parameters: (i) the first timestep (rows: 600, 750, 900) and (ii) the final timestep (columns:
 1190 250, 375, 500).

1191

Metrics	PSNR			SSIM			LPIPS		
Schedule	600	750	900	600	750	900	600	750	900
250	27.12	27.24	26.82	0.744	0.767	0.718	0.290	0.287	0.335
375	27.03	27.29	27.07	0.736	0.770	0.738	0.304	0.291	0.319
500	26.87	27.17	27.01	0.723	0.763	0.741	0.330	0.305	0.317

1196

1197

A.10 ALTERNATIVE SAMPLING SCHEMES

1199

1200



1211

1212

Figure 9: Unadjusted Langevin Algorithm (ULA) with different step sizes. Larger step sizes accelerate convergence but introduce greater discretization error, substantially degrading sample quality.

1215

1216

Sampling in a space as high-dimensional as ($3 \times 256 \times 256$) is a challenging task. Many standard sampling algorithms are not suitable in this setting. A key requirement for efficiency is the use of gradient information to accelerate convergence. The simplest such method is the Unadjusted Langevin Algorithm (ULA).

1217

1218

However, because ULA lacks a Metropolis–Hastings (MH) correction, its step size must be tuned carefully. Figure 9 illustrates this trade-off: large step sizes enable rapid exploration but cause significant discretization error as σ_y approaches the target value, resulting in poor samples; conversely, small step sizes reduce error but lead to very slow exploration and long runtime.

1219

1220

1221

1222

1223

Since different stages of the sampling chain require different effective step sizes, algorithms with a Metropolis–Hastings (MH) correction are more attractive, as the acceptance test provides a natural criterion for adapting step size. We therefore consider the Metropolis-Adjusted Langevin Algorithm (MALA), the No-U-Turn Sampler (NUTS), and Hamiltonian Monte Carlo (HMC). In practice, however, both MALA and NUTS tend to settle on excessively small step sizes in this high-dimensional setting, resulting in impractically long runtimes. By contrast, HMC accommodates larger step sizes and achieves a more favorable trade-off between accuracy and efficiency, making it the most suitable choice for our framework.

1231

1232

A.11 ALTERNATIVE PRIOR MODEL: GAN-BASED INFERENCE

1233

1234

1235

1236

In place of the diffusion models, we experimented with StyleGAN2 (Karras et al., 2020) as the prior model for FFHQ 256×256 dataset. The quantitative results are presented in Table 17. The quality of image samples are significantly inferior to diffusion models across all tasks.

1237

1238

1239

A.12 ADDITIONAL QUALITATIVE RESULTS

1240

1241

In this section, we present additional qualitative results. Since we don’t have access to an LDM for ImageNet (256×256), ReSample cannot be applied to this dataset.

1242

1243

1244 Table 17: GAN-Based Inference for Linear IPs on FFHQ (256 × 256) with Gaussian Noise $\sigma_y =$
1245 0.05.

1246

1247

	Super Resolution (×4)			Random Inpainting (92%)		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
NA-NHMC (GAN)	18.27	0.454	0.513	17.80	0.440	0.528
NA-NHMC (DDIM)	27.29	0.770	0.291	26.72	0.785	0.268

1252

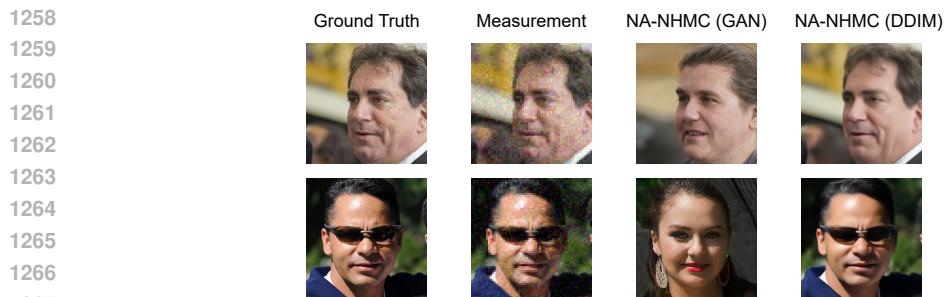
1253

1254

1255

1256

1257

1268 Figure 10: Comparison between GAN and Diffusion Model (DDIM) for SR(×4). FFHQ (256 × 256). $\sigma_y = 0.05$.

1269

1270

1271

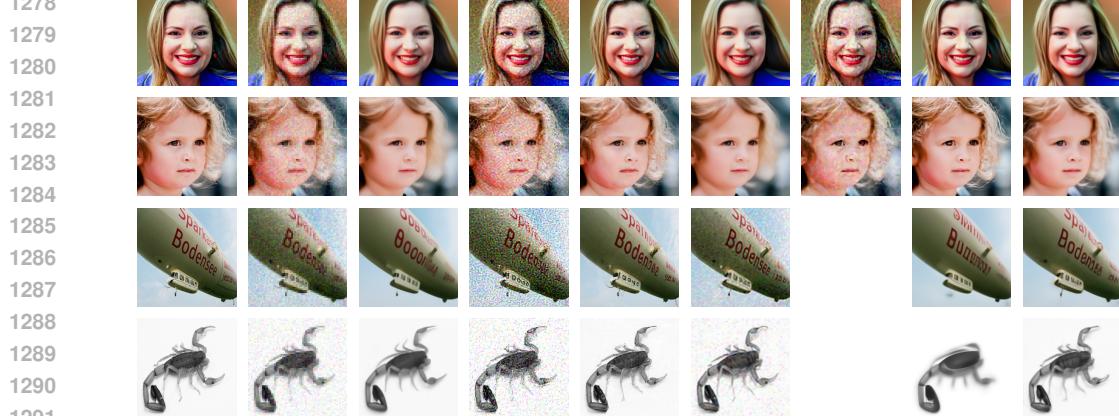
1272

1273

1274

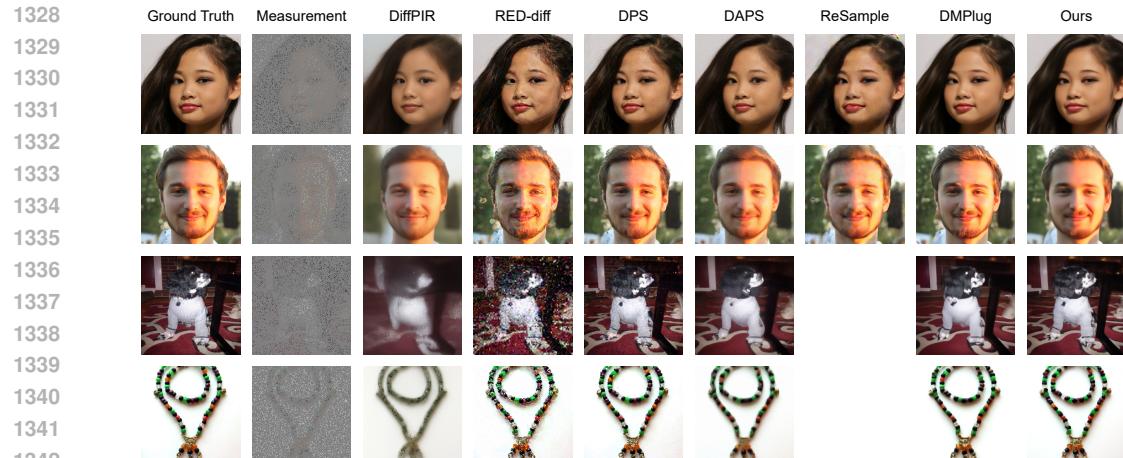
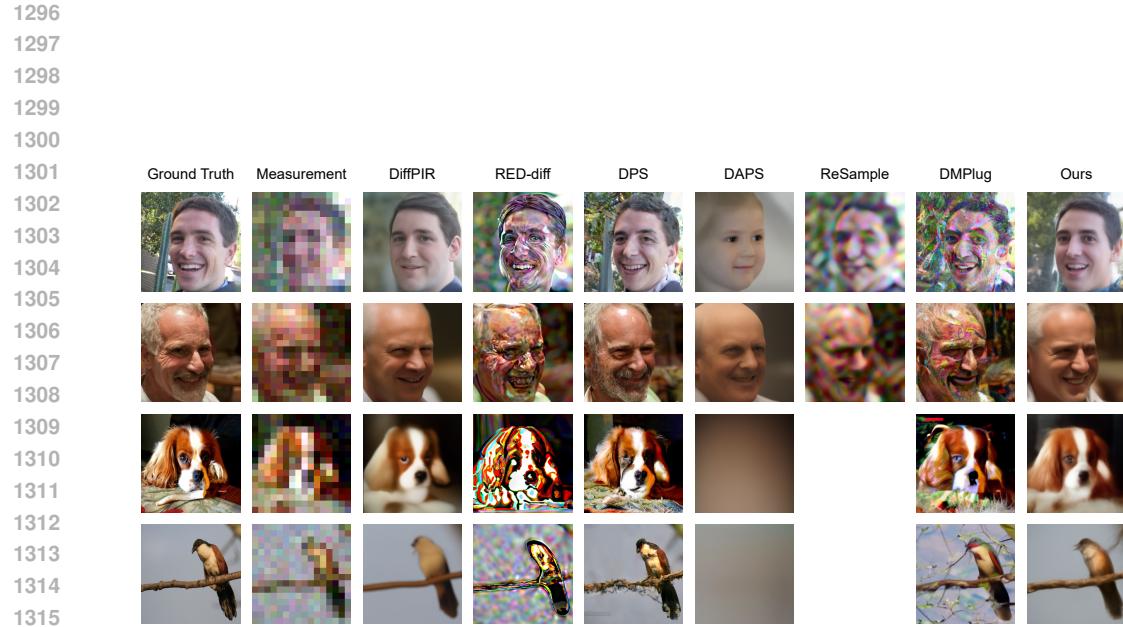
1275

1276

1293 Figure 11: SR(×4). (Top) FFHQ (256 × 256). (Bottom) ImageNet (256 × 256). $\sigma_y = 0.05$

1294

1295



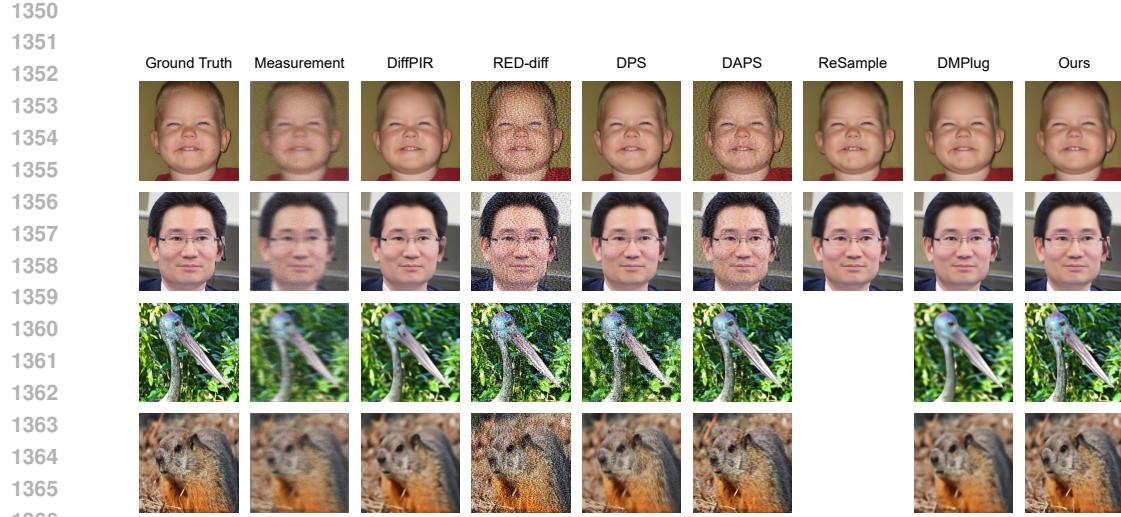


Figure 14: Gaussian deblurring. (Top) FFHQ (256 × 256). (Bottom) ImageNet (256 × 256). $\sigma_y = 0.05$

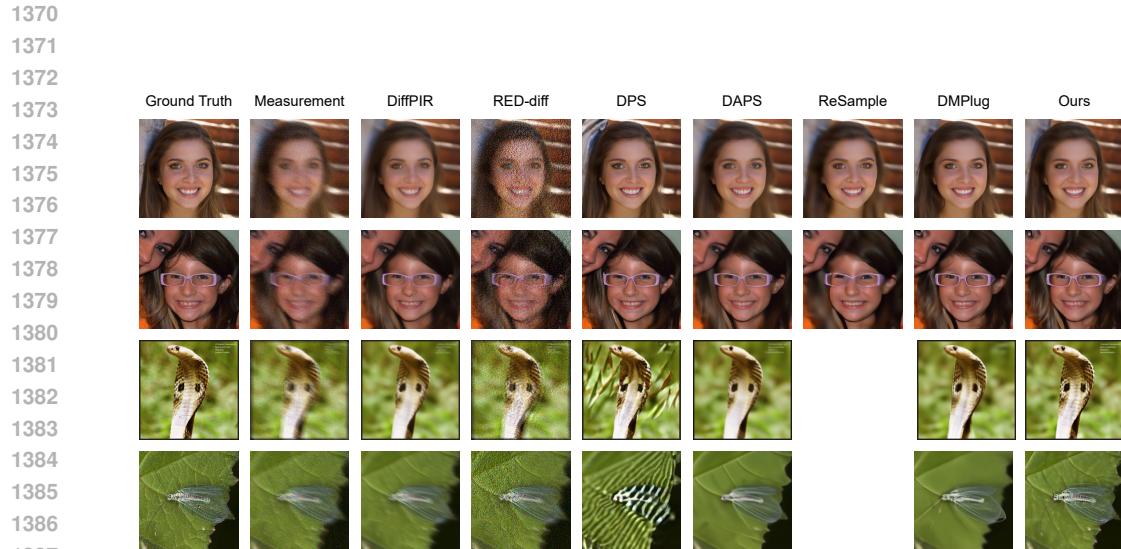


Figure 15: Nonlinear deblurring. (Top) FFHQ (256 × 256). (Bottom) ImageNet (256 × 256). $\sigma_y = 0.05$

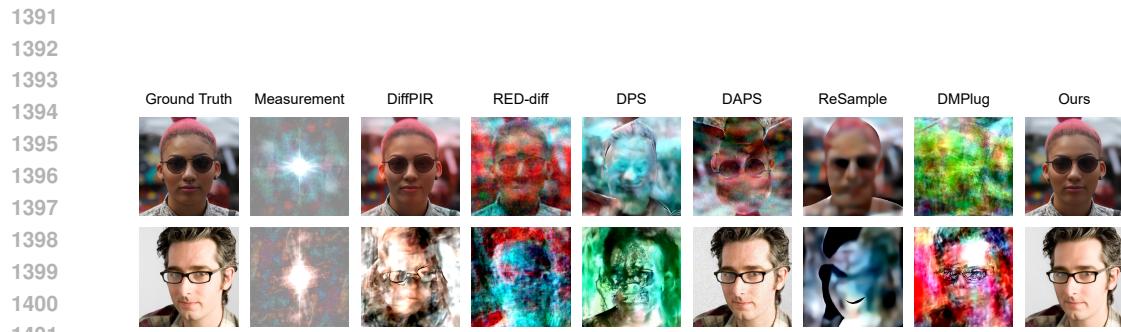


Figure 16: Phase retrieval. FFHQ (256 × 256).

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

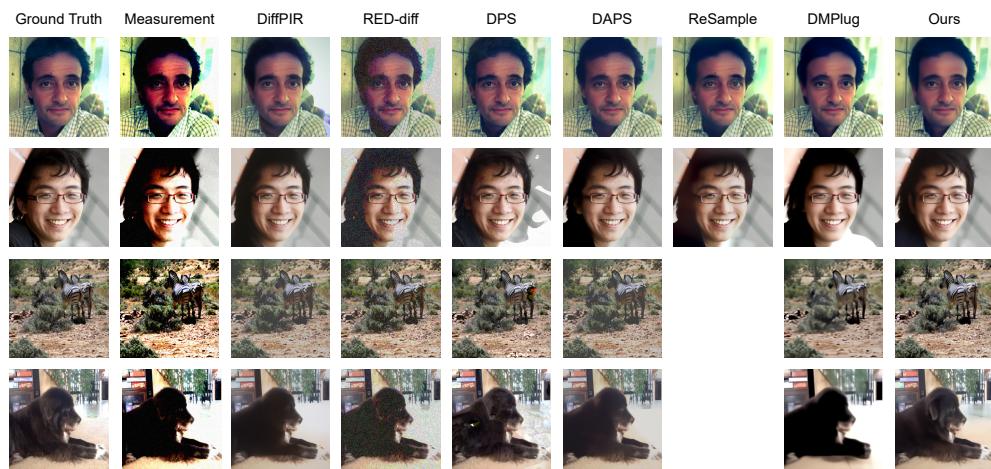


Figure 17: HDR reconstruction. (Top) FFHQ (256 × 256). (Bottom) ImageNet (256 × 256).
 $\sigma_y = 0.05$